Theta-Graphs and Other Constrained Spanners

by

André van Renssen

A thesis submitted to the Faculty of Graduate and Postdoctoral Affairs
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Computer Science

Carleton University
Ottawa, Ontario

Abstract

We provide improved upper bounds on the spanning ratio of various geometric graphs, one of which being θ-graphs. Given a set of points in the plane, a θ-graph partitions the plane around each vertex into m disjoint cones, each having aperture $\theta = 2\pi/m$, and adds an edge to the 'closest' vertex in each cone. We provide tight bounds on a large number of these graphs, for different values of m, and improve the upper bounds on the other θ-graphs.

We also study the ordered setting, where the θ-graph is built by inserting vertices one at a time and we consider only previously-inserted vertices when determining the 'closest' vertex in each cone. We improve some of the upper bounds in this setting, but our main contribution is that we show that a number of θ-graphs that are spanners in the unordered setting are not spanners in the ordered setting.

Our main topic, however, is the constrained setting: We introduce line segment constraints that the edges of the graph are not allowed to cross and show that the upper bounds shown for θ-graphs carry over to constrained θ-graphs. We also construct a bounded-degree plane spanner based on the constrained half-θ_6-graph (the constrained Delaunay graph whose empty convex shape is an equilateral triangle) and we provide a local competitive routing algorithm for the constrained θ_6-graph.

Next, we look at constrained Yao-graphs, which are comparable to constrained θ-graphs, but use a different distance function, and show that these graphs are spanners.

Finally, we look at constrained generalized Delaunay graphs: Delaunay graphs where the empty convex shape is not necessarily a circle, but can be any convex shape. We show that regardless of the convex shape, these graphs are connected, plane spanners. We then proceed to improve the spanning ratio for a subclass of these graphs, where the empty convex shape is an arbitrary rectangle.

Acknowledgements

I would like to express my appreciation and thanks to my supervisors Dr. Prosenjit Bose, Dr. Vida Dujmović, and Dr. Pat Morin. You have been great mentors for me, offering support and advice when I needed it, while allowing me to pursue my research interests. This has allowed me to grow as a researcher, which has been and will continue to be priceless during my future career.

I would also like to express my appreciation to Dr. Nicolas Bonichon, Dr. Paola Flocchini, Dr. Michiel Smid, and Dr. Brett Stevens for serving as committee members during my defence, as well as for all your helpful comments and suggestions.

I would also like to thank all professors, students, and other members of the Computational Geometry Lab for making my stay interesting and eventful. Your openness has helped me gain insight into what being a professor entails. The atmosphere provided by the group as a whole has been very pleasant, which is in no small part because everyone treats each other as colleagues as well as friends. The administration of the School of Computer Science has also been a tremendous support and has helped me make sense of all the guidelines, forms, and requirements associated with obtaining my PhD.

An important part of my PhD has involved getting to know and work with numerous researchers, either while they were visiting Carleton University or during the various conferences and workshops I attended. Special thanks to Dr. Rolf Fagerberg, who visited during the initial phase of my research and was a tremendous help while working on the first results which grew to become this thesis, and Dr. Bettina Speckmann, for encouraging me to apply to do my PhD at Carleton University, which turned out to be the best choice I could have made.

I also want to thank the School of Linguistics and Language Studies, especially the Japanese language section, for allowing me to attend their language courses. This helped me both to pick up a new language and to occasionally force my mind not to think about my research, which gave rise to some of the best parts of this thesis.

Special thanks to my family for all their support during some of the harder periods

and making me feel right back at home during my yearly visits. I cannot express the peace of mind it gave me to know that you were always there for me.

I would also like to thank all my friends at Anime@Carleton for making sure I never had a dull weekend and exposing me (for better or worse) to all sorts of shows I never would have considered watching otherwise. Additional thanks go to my friends at N2U and Anime North, both staff and attendees, for giving me a few unforgettable weeks at both events. I would also like to thank my friends at activeAnime for giving me the opportunity to improve my non-scientific writing, and having a lot of fun in the meantime.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

In this thesis we look at networks or graphs. Graphs are widely used to model the structure of information in many fields, including economics, genetics, social sciences, mathematics, and computer science. Our main goal is to uncover the underlying properties of these graphs and use these properties to show that the various graphs that we study contain short paths between every pair of vertices. We also aim to find efficient routing algorithms on these graphs, i.e. algorithms that are able to follow these short paths without knowing the full structure of the graph.

In the remainder of this chapter, we introduce some of the definitions used throughout this thesis. We also review the related work and relevant results, as well as give an overview of the contributions of this thesis.

1.1 Spanners

A geometric graph G is a graph whose vertices are points in the plane and whose edges are line segments between pairs of points. A graph G is called plane if no two edges intersect properly, i.e. if any two edges intersect only at a shared endpoint. Every edge is weighted by the Euclidean distance between its endpoints. The distance between two vertices u and v in G, denoted by $\delta_G(u, v)$, or simply $\delta(u, v)$ 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, $\delta_H(u, v) \leq t \cdot \delta_G(u, v)$. 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 [28, 53] for a

comprehensive overview of the topic).

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 points 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* or *are visible to 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 [36] 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 [38] showed how to construct a spanner of $Vis(P, S)$ with constant spanning ratio and constant degree. Bose and Keil [23] showed that the Constrained Delaunay Triangulation, which contains an edge between two vertices u and v if and only if there exists a circle with u and v on its boundary that does not contain any vertices visible to both u and v, is a 2.42-spanner of $Vis(P, S)$.

1.2 Constrained Bounded-Degree Spanners

One area of research focuses on building spanners in such a way that the maximum degree of a vertex is bounded by a constant and the resulting graph is plane. The first spanner satisfying these requirements was introduced by Bose *et al.* [22]. This spanner was constructed from the Delaunay triangulation and has degree at most 27 and has spanning ratio $(\pi+1)\cdot SR_D \approx 8.275$, where SR_D is the spanning ratio of the Delaunay triangulation (currently the best upper bound on SR_D is 1.998 [58]). This result was later improved by Li and Wang [49], who reduced the maximum degree to 23 and the spanning ratio to $(1 + \pi \sin(\pi/4)) \cdot SR_D \approx 6.437$. Bose *et al.* [29] reduced the maximum degree further to 17 but increased the spanning ratio to $(2+2\sqrt{3}+3\pi/2+2\pi\sin(\pi/12))\cdot SR_D \approx 23.582$. Kanj and Perković [44] showed how to construct a spanner of maximum degree 14 with spanning ratio $(1+2\pi/(14\cos(\pi/14)))\cdot SR_D \approx 2.918$.

Kanj and Xia [45] later improved this further to maximum degree 11 with spanning ratio $2\sin(2\pi/5)\cos(\pi/5)/(2\sin(2\pi/5)\cos(\pi/5)-1)\cdot SR_D \approx 5.706$. Recently, Bose *et al.* [11] decreased the maximum degree to 6 while increasing the spanning ratio to $(1/(1-\tan(\pi/7)(1+1/\cos(\pi/14))))\cdot SR_D \approx 81.664$, which is currently the lowest maximum degree when starting from the Delaunay triangulation.

When starting from the half-θ_6-graph (which is identical to the Delaunay triangulation whose empty region is an equilateral triangle), Bonichon *et al.* [7] showed how to construct a 6-spanner of maximum degree 6. Recently, Bonichon *et al.* [9] showed that it is possible to construct a spanner of maximum degree 4 by starting from the L_1-Delaunay triangulation, which uses an empty square that is rotated $\pi/4$ radians as its empty region. The resulting graph has a spanning ratio of $\sqrt{4+2\sqrt{2}}(1+\sqrt{2})^2(3+\sqrt{2})^6 \approx 112676$. On the other hand, a lower bound on the degree is 3, since a Hamiltonian path through a grid of n vertices has spanning ratio $\Omega(\sqrt{n})$. All of the above results, however, focus on the unconstrained setting.

In Chapter 3, 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.* [18]. A key difficulty in proving the latter stems from the fact that the constrained Delaunay graph is **not** necessarily a triangulation (see Figure 1.1). We then generalize the elegant construction of Bonichon *et al.* [7] 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.

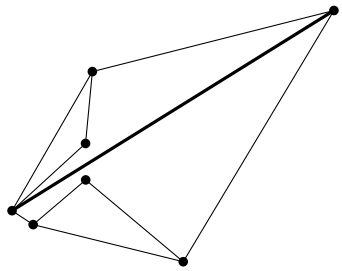

Figure 1.1: The constrained half-θ_6-graph is not necessarily a triangulation. The thick line segment represents a constraint

1.3 Routing in the Constrained Setting

Another fundamental problem in any graph is the question of how to route a message from one vertex to another. What makes this more challenging is that often in a network the routing strategy must be *local*. Informally, a routing strategy is *k-local* when the routing algorithm must decide which vertex to forward a message to based solely on knowledge of the source and destination vertex, the current vertex and all vertices that can be reached from the current vertex by following at most k edges. Routing algorithms are considered *geometric* when the underlying graph is embedded in the plane, with edges being straight line segments connecting pairs of points and weighted by the Euclidean distance between their endpoints. Geometric routing algorithms are important in wireless sensor networks (see [52] and [54] for surveys of the area) since they offer routing strategies that use the coordinates of the vertices to help guide the search as opposed to using the more traditional routing tables.

As was the case with constructing spanners, most of the research related to routing has focused on the situation where the underlying graph is the complete Euclidean graph. In Chapter 4, we study this problem in the constrained setting. While it is known that some constrained graphs are spanners and thus contain short paths, it is not known how to follow these paths in a local fashion. It seems very common to run a shortest path algorithm on the graph and send messages along the returned path, but this becomes less and less feasible as the graph grows. To overcome this issue, we look at local routing algorithms in the constrained setting. Furthermore, we require our algorithms to be *competitive*, i.e. the length of the returned path needs to be related to the length of the shortest path in the graph (a more precise definition follows in Chapter 2). We first show that no deterministic 1-local routing algorithm is $o(\sqrt{n})$-competitive on all pairs of non-visible vertices of the constrained θ_6-graph (the union of two constrained half-θ_6-graphs, formally defined in Chapter 2), regardless of the amount of memory it is allowed to use. Two other lower bounds on the competitiveness were provided by Bose *et al.* [18] for the unconstrained setting: The first states that no deterministic k-local 0-memory routing algorithm can be better than 2-competitive on the θ_6-graph. The second states that no deterministic k-local 0-memory routing algorithm can be better than $5/\sqrt{3}$-competitive on the half-θ_6-graph. This second lower bound applies specifically to the situation where the destination lies in a negative cone of the source.

Following this negative result, we describe a 1-local routing algorithm between any two visible vertices of the constrained θ_6-graph that guarantees that the length

of the returned path is at most 2 times the Euclidean distance between the source and destination. In order to do this, we recall that the constrained θ_6-graph is the union of two constrained half-θ_6-graphs and provide a 1-local routing algorithm on the constrained half-θ_6-graph that provides the same guarantees on the length of the returned path. Routing on the constrained θ_6-graph now becomes a matter of determining which constrained half-θ_6-graph to route on, followed by applying the routing algorithm on that graph. Additionally, we provide a 1-local 18-competitive routing algorithm on the constrained half-θ_6-graph. To the best of our knowledge, these are the first local routing algorithms in the constrained setting with guarantees on the length of the returned path.

1.4 Theta-Graphs

We then take a step back and look at θ-graphs in the unconstrained setting, i.e. the situation where the set of constraints S is empty and $Vis(P, S)$ is the complete graph. Introduced independently by Clarkson [36] and Keil [46], θ-graphs are constructed as follows (a more precise definition follows in Chapter 2): for each vertex u, we partition the plane into m disjoint cones with apex u, each having aperture $\theta = 2\pi/m$. When m cones are used, we denote the resulting θ-graph by the θ_m-graph. The θ-graph is constructed by, for each cone with apex u, connecting u to the vertex v whose projection onto the bisector of the cone is closest. Ruppert and Seidel [55] showed that the spanning ratio of these graphs is at most $1/(1 - 2\sin(\theta/2))$, when $\theta < \pi/3$, i.e. there are at least seven cones.

Recently, Bonichon *et al.* [6] showed that the θ_6-graph has spanning ratio 2. This was done by dividing the cones into two sets, positive and negative cones, such that each positive cone is adjacent to two negative cones and vice versa. It was shown that when edges are added only in the positive cones, in which case the graph is called the half-θ_6-graph, the resulting graph is equivalent to the Delaunay triangulation where the empty region is an equilateral triangle. The spanning ratio of this graph is 2, as shown by Chew [35]. An alternative, inductive proof of the spanning ratio of the half-θ_6-graph was presented by Bose *et al.* [18], along with an optimal local competitive routing algorithm on the half-θ_6-graph.

Tight bounds, i.e. matching upper and lower bounds, on spanning ratios are notoriously hard to obtain. The standard Delaunay triangulation (where the empty region is a circle) is a good example. Its spanning ratio has been studied for over 20 years and the upper and lower bounds still do not match. Also, even though it was

introduced about 25 years ago, the spanning ratio of the θ_6-graph has only recently been shown to be finite and tight, making it the first and, until now, only θ-graph for which tight bounds are known.

In Chapter 5, we improve on the existing upper bounds on the spanning ratio of all θ-graphs with at least seven cones, though the proof can also be applied to the θ_6-graph. First, we generalize the spanning proof of the half-θ_6-graph given by Bose *et al.* [18] to a large family of θ-graphs: the $\theta_{(4k+2)}$-graph, where $k \geq 1$ is an integer. We show that the $\theta_{(4k+2)}$-graph has a tight spanning ratio of $1+2\sin(\theta/2)$. We continue by looking at upper bounds on the spanning ratio of the other three families of θ-graphs: the $\theta_{(4k+3)}$-graph, the $\theta_{(4k+4)}$-graph, and the $\theta_{(4k+5)}$-graph, where k is an integer and at least 1. We show that the $\theta_{(4k+4)}$-graph has a spanning ratio of at most $1 + 2\sin(\theta/2)/(\cos(\theta/2) - \sin(\theta/2))$. We also show that the $\theta_{(4k+3)}$-graph and the $\theta_{(4k+5)}$-graph have spanning ratio at most $\cos(\theta/4)/(\cos(\theta/2) - \sin(3\theta/4))$.

We also consider routing on θ-graphs. Due to the structure and approach of the spanning proof of Ruppert and Seidel [55], their proof also shows that the *θ-routing* algorithm (defined in Section 2.1) is $1/(1 - 2\sin(\theta/2))$-competitive on these graphs. The θ-routing algorithm is the standard routing algorithm on all θ-graphs having at least seven cones. As was the case for Ruppert and Seidel, the structure of our spanning proofs implies that the upper bounds also apply to the competitiveness of θ-routing on these graphs. These results are summarized in Table 1.1.

	Current Spanning	Current Routing	Previous Spanning & Routing
$\theta_{(4k+2)}$-graph	$1 + 2\sin\left(\frac{\theta}{2}\right)$	$\frac{1}{1-2\sin\left(\frac{\theta}{2}\right)}$ [55]	$\frac{1}{1-2\sin\left(\frac{\theta}{2}\right)}$ [55]
$\theta_{(4k+3)}$-graph	$\frac{\cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right)-\sin\left(\frac{3\theta}{4}\right)}$	$1 + \frac{2\sin\left(\frac{\theta}{2}\right)\cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right)-\sin\left(\frac{\theta}{2}\right)}$	$\frac{1}{1-2\sin\left(\frac{\theta}{2}\right)}$ [55]
$\theta_{(4k+4)}$-graph	$1 + \frac{2\sin\left(\frac{\theta}{2}\right)}{\cos\left(\frac{\theta}{2}\right)-\sin\left(\frac{\theta}{2}\right)}$	$1 + \frac{2\sin\left(\frac{\theta}{2}\right)}{\cos\left(\frac{\theta}{2}\right)-\sin\left(\frac{\theta}{2}\right)}$	$\frac{1}{1-2\sin\left(\frac{\theta}{2}\right)}$ [55]
$\theta_{(4k+5)}$-graph	$\frac{\cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right)-\sin\left(\frac{3\theta}{4}\right)}$	$1 + \frac{2\sin\left(\frac{\theta}{2}\right)\cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right)-\sin\left(\frac{\theta}{2}\right)}$	$\frac{1}{1-2\sin\left(\frac{\theta}{2}\right)}$ [55]

Table 1.1: An overview of current and previous spanning and routing ratios of θ-graphs

Finally, we present improved lower bounds on the spanning ratio of these graphs (see Section 5.2) and we provide a partial order on the upper and lower bounds of

these families (see Section 5.3). In particular, we show that θ-graphs with $4k+4$ cones have spanning ratio at least $1+2\tan(\theta/2)+2\tan^2(\theta/2)$. This is somewhat surprising since, for equal values of k, the spanning ratio of θ-graphs with $4k+4$ cones is greater than that of θ-graphs with $4k+2$ cones, showing that increasing the number of cones can make the spanning ratio worse.

1.5 Constrained Theta-Graphs

After looking at the unconstrained setting, we return our attention to the constrained setting. In Chapter 6, we generalize the results on unconstrained θ-graphs to the constrained setting. In order to do this, there are two main obstacles that need to be overcome. First, the main difficulty with the constrained setting is that induction cannot be applied directly, as the destination need not be visible from the vertex closest to the source. Second, when the graph does not have $4k+2$ cones, the cones do not line up as nicely as in Chapter 3, making it more difficult to apply induction.

We overcome these two difficulties and show that constrained θ-graphs with $4k+2$ cones have a spanning ratio of at most $1+2\sin(\theta/2)$, where θ is $2\pi/(4k+2)$. Since the lower bounds of the unconstrained θ-graphs carry over to the constrained setting, this shows that this spanning ratio is tight. We also show that constrained θ-graphs with $4k+4$ cones have a spanning ratio of at most $1+2\sin(\theta/2)/\left(\cos(\theta/2)-\sin(\theta/2)\right)$, where θ is $2\pi/(4k+4)$. Finally, we show that constrained θ-graphs with $4k+3$ or $4k+5$ cones have a spanning ratio of at most $\cos(\theta/4)/\left(\cos(\theta/2)-\sin(3\theta/4)\right)$, where θ is $2\pi/(4k+3)$ or $2\pi/(4k+5)$.

1.6 Ordered Theta-Graphs

In Chapter 7, we look at the ordered variant of θ-graphs. The ordered θ-graph is constructed by inserting the vertices one by one and connecting each vertex to the closest previously-inserted vertex in each cone (a more precise definition follows in Chapter 7). These graphs were introduced by Bose *et al.* [21] in order to construct spanners with nice additional properties, such as logarithmic maximum degree and logarithmic diameter for specific insertion orders. The current upper bound on the spanning ratio of these graphs is $1/(1-2\sin(\theta/2))$ [21], when $\theta < \pi/3$, i.e. there are at least seven cones.

By applying techniques similar to the ones used in Chapters 5 and 6, we improve the spanning ratio of the ordered $\theta_{(4k+4)}$-graph to $1+2\sin(\theta/2)/(\cos(\theta/2)-\sin(\theta/2))$ and show that this spanning ratio is tight in the worst case. Unfortunately, this inductive proof cannot be applied to ordered θ-graphs with an odd number of cones, as the triangle we apply induction on can become larger, depending on the order in which the vertices are inserted. We also show that the ordered $\theta_{(4k+2)}$-graph ($k \geq 2$) has a tight spanning ratio of $1/(1-2\sin(\theta/2))$.

We also provide lower bounds for ordered θ-graphs with $4k+3$ and $4k+5$ cones (see Table 1.2). For ordered θ-graphs with $4k+2$ and $4k+5$ cones these lower bounds are strictly greater than the worst case spanning ratios of their unordered counterparts. Finally, we show that there exist orderings such that the ordered θ-graphs with 3, 4, 5, and 6 cones are not spanners. For the ordered θ_3-graph this is not surprising, as its unordered counterpart is connected [1], but not a spanner [41]. For the ordered θ_4, θ_5, and θ_6-graph, however, this is a bit surprising since their unordered counterparts have recently been shown to be spanners [5, 6, 27]. In other words, we show, for the first time, that obtaining the nice additional properties of the ordered θ-graphs comes at a price.

	Upper Bound	Lower Bound
θ_3, θ_4, θ_5, and θ_6-graph	-	Not spanners.
$\theta_{(4k+2)}$-graph	$\frac{1}{1-2\sin\left(\frac{\theta}{2}\right)}$, for $k \geq 2$ [21]	$\frac{1}{1-2\sin\left(\frac{\theta}{2}\right)}$
$\theta_{(4k+3)}$-graph	$\frac{1}{1-2\sin\left(\frac{\theta}{2}\right)}$ [21]	$\frac{\cos\left(\frac{\theta}{4}\right)+\sin\theta}{\cos\left(\frac{3\theta}{4}\right)}$
$\theta_{(4k+4)}$-graph	$1+\frac{2\sin\left(\frac{\theta}{2}\right)}{\cos\left(\frac{\theta}{2}\right)-\sin\left(\frac{\theta}{2}\right)}$	$1+\frac{2\sin\left(\frac{\theta}{2}\right)}{\cos\left(\frac{\theta}{2}\right)-\sin\left(\frac{\theta}{2}\right)}$
$\theta_{(4k+5)}$-graph	$\frac{1}{1-2\sin\left(\frac{\theta}{2}\right)}$ [21]	$1+\frac{2\sin\left(\frac{\theta}{2}\right)\cdot\cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right)-\sin\left(\frac{3\theta}{4}\right)}$

Table 1.2: An overview of upper and lower bounds on the spanning ratio of ordered θ-graphs

1.7 Constrained Yao-Graphs

Next, we move away from θ-graphs and look at other constrained graphs. In Chapter 8, we look at Yao-graphs. Introduced independently by Flinchbaugh and Jones [42] and Yao [60], Yao-graphs partition the plane around each vertex into k disjoint cones, each having aperture $\theta = 2\pi/k$. The Y_k-graph is constructed by, for each cone of each vertex u, connecting u to the vertex v that is closest to u (a more precise definition follows in Chapter 8). However, neither Flinchbaugh and Jones nor Yao proved that these graphs are spanners. To the best of our knowledge, the first such proof was given by Althöfer *et al.* [2], who proved that for every spanning ratio $t > 1$, there exists a k such that the Y_k-graph is a t-spanner. It appears that a similar result was already known by that time, since Clarkson [36] remarked in 1987 that the Y_{12}-graph is a $1 + \sqrt{3}$-spanner, though without providing a proof or reference.

In 2004, Bose *et al.* [25] provided a more precise bound on the spanning ratio. They showed that Yao-graphs with at least 9 cones have spanning ratio at most $1/(\cos\theta - \sin\theta)$. This was later strengthened to show that Yao-graphs with at least 7 cones are $1/\left(1 - 2\sin(\theta/2)\right)$-spanners [13]. Recently, Damian and Raudonis [37] showed that the Y_6-graph is a 17.64-spanner and Bose *et al.* [14] showed that the Y_4-graph has spanning ratio at most 663. Barba *et al.* [4] showed that the Y_5-graph is a $\left(2 + \sqrt{3}\right)$-spanner. In the same paper, they also improved the upper bound on the spanning ratio of the Y_6-graph to 5.8 and that of Yao-graphs with an odd number of cones to $1/\left(1 - 2\sin(3\theta/8)\right)$. On the other hand, when a Yao-graph has less than 4 cones, El Molla [41] showed that there is no constant t such that it is a t-spanner.

The above results, however, focus on Yao-graphs where the underlying graph is the complete Euclidean geometric graph. We study this problem in the more general setting with the introduction of line segment *constraints*. Surprisingly, Yao-graphs have not been considered in the constrained setting. As such, it is unknown whether they are spanners of $Vis(P, S)$. In Chapter 8, we set an important first step towards answering this question by showing that constrained Yao-graphs with at least 7 cones are spanners. In particular, we prove that constrained Yao-graphs with an even number of cones have spanning ratio at most $1/\left(1 - 2\sin(\theta/2)\right)$. When the constrained Yao-graph has an odd number of cones, we can improve on this result and show an upper bound of $1/\left(1 - 2\sin(3\theta/8)\right)$. These bounds match the current upper bounds in the unconstrained setting.

1.8 Constrained Generalized Delaunay Graphs

Finally, we look at the constrained generalized Delaunay graph, where the empty convex shape can be any convex polygon. Given an arbitrary convex shape C, an unconstrained Delaunay graph is constructed by adding an edge between two vertices p and q if and only if there exists a homothet (scaled translate) of C with p and q on its boundary that does not contain any other vertices. In the unconstrained setting, it is known that generalized Delaunay graphs are spanners [12], regardless of the convex shape used to construct it. These bounds are very general, but unfortunately not tight.

In special cases, better bounds are known. For example, when the empty convex shape is a circle, Dobkin *et al.* [40] showed that the spanning ratio is at most $\pi(1+\sqrt{5})/2 \approx 5.09$. Improving on this, Keil and Gutwin [47] reduced the spanning ratio to $4\pi/3\sqrt{3} \approx 2.42$. Recently, Xia showed that the spanning ratio is at most 1.998 [58]. On the other hand, Bose *et al.* [17] showed a lower bound of 1.58, which is greater than $\pi/2$, which was conjectured to be the tight spanning ratio up to that point. Later, Xia and Zhang [59] improved this to 1.59.

Chew [35] showed that if an equilateral triangle is used instead, the spanning ratio is 2 and this ratio is tight. In the case of squares, Chew [34] showed that the spanning ratio is at most $\sqrt{10} \approx 3.16$. This was later improved by Bonichon *et al.* [8], who showed a tight spanning ratio of $\sqrt{4+2\sqrt{2}} \approx 2.61$.

In Chapter 9, we show that the constrained generalized Delaunay graph G is a spanner whose spanning ratio depends solely on the properties of the empty convex shape C used to create it: We show that G satisfies the α_C-diamond property and the visible-pair κ_C-spanner property (defined in Section 9.2.3), which implies that it is a t-spanner for:

$$t = \begin{cases} 2\kappa_C \cdot \max\left(\frac{3}{\sin(\alpha_C/2)}, \kappa_C\right), & \textit{if } G \textit{ is a triangulation} \\ 2\kappa_C^2 \cdot \max\left(\frac{3}{\sin(\alpha_C/2)}, \kappa_C\right), & \textit{otherwise.} \end{cases}$$

To this end, we first provide two simple proofs showing that these graphs are plane and connected. Though the spanning proof is very general, since it holds for arbitrary convex shapes, its implied spanning ratio is far from tight. To improve on this, we also consider the special case where the empty convex shape C is a rectangle and show that it has spanning ratio at most $\sqrt{2} \cdot (2l/s + 1)$, where l and s are the length of the long and short side of C.

Chapter 2

Preliminaries

Let a *cone* be a region in the plane between two rays originating from the same vertex referred to as the apex of the cone. When constructing a (constrained) θ_m-graph, for each vertex u consider the rays originating from u with the angle between consecutive rays being $\theta = 2\pi/m$ (see Figure 2.1). Each pair of consecutive rays defines a cone. The cones are oriented such that the bisector of some cone coincides with the vertical halfline through u that lies above u. We refer to this cone as C_0 and number the cones in clockwise order around u. The cones around the other vertices have the same orientation as the ones around u. We write C_i^u to indicate the i-th cone of a vertex u, or C_i if the apex is clear from the context.

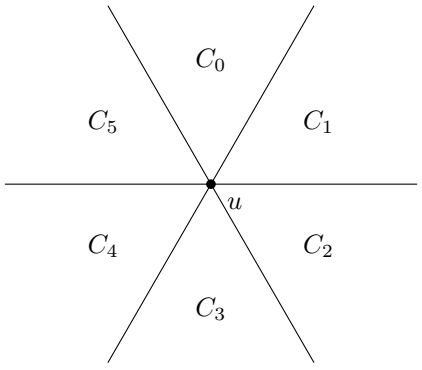

Figure 2.1: The cones having apex u in the θ_6-graph

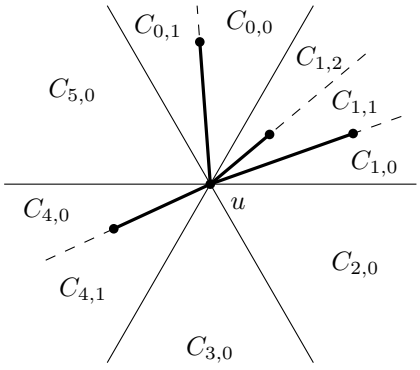

Figure 2.2: The subcones having apex u in the θ_6-graph. Constraints are shown as thick line segments

We only consider point sets in general position: no two vertices lie on a line parallel

to one of the rays that define the cones, no two vertices lie on a line perpendicular to the bisector of one of the cones, and no three points are collinear.

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 *subcones* (see Figure 2.2). We use $C_{i,j}^u$ to denote the j-th subcone of C_i^u. 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 consider a cone that is not split to be a *single subcone*.

The constrained θ_m-graph is constructed as follows: for each subcone $C_{i,j}$ of each vertex u, add an edge from u to the closest vertex in that cone, where distance is measured along the bisector of the original cone (*not the subcone*) (see Figure 2.3). More formally, we add an edge between two vertices u and v if v can see u, $v \in C_{i,j}$, and for all vertices $w \in C_{i,j}$ that can see u, $|uv'| \leq |uw'|$, where v' and w' denote the orthogonal projection of v and w on the bisector of C_i and $|xy|$ denotes the length of the line segment between two points x and y. Note that our general position assumptions imply that each vertex adds at most one edge per subcone to the graph. The θ_m-graph is a special case of the constrained θ_m-graph where no constraints are present, i.e. $S = \emptyset$.

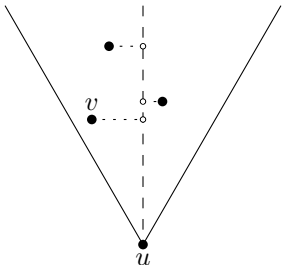

Figure 2.3: Three vertices are projected onto the bisector of a cone of u. Vertex v is the closest vertex

Next, we define the constrained half-θ_6-graph. The constrained half-θ_6-graph is similar to the constrained θ_6-graph, but there is one major difference: its cones are categorized as positive and negative. Let $(C_0, \overline{C}_2, C_1, \overline{C}_0, C_2, \overline{C}_1)$ be the sequence of cones in counterclockwise order starting from C_0 (see Figure 2.4). 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 $\overline{C}_i^u$ to denote cones C_i and $\overline{C}_i$ with apex u. Note that for any two vertices u and v, $v \in C_i^u$ if and

only if $u \in \overline{C}_i^v$. Analogous to the subcones defined for θ-graphs, constraints can split cones into subcones. We call a subcone of a positive cone a *positive subcone* and a subcone of a negative cone a *negative subcone* (see Figure 2.5).

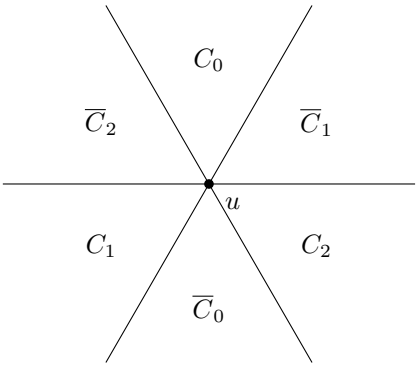

Figure 2.4: The cones having apex u in the half-θ_6-graph

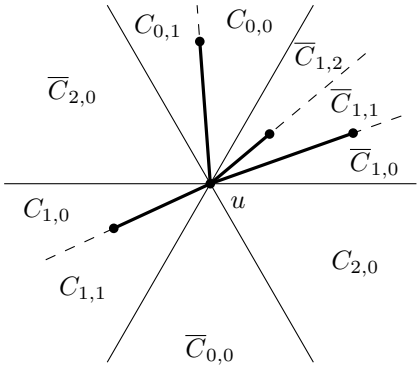

Figure 2.5: The subcones having apex u in the half-θ_6-graph. Constraints are shown as thick line segments

The constrained half-θ_6-graph is constructed as follows: for each **positive** subcone of each vertex u, add an edge from u to the closest vertex in that subcone that can see u, where distance is measured along the bisector of the original cone (not the subcone). 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 points $w \in C_{i,j}^u$ 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. Again, note that our general position assumption implies that each vertex adds at most one edge per subcone to the graph.

Finally, given a vertex w in positive cone C of a vertex u, we define the *canonical triangle* T_{uw} to be the triangle defined by the borders of C and the line through w perpendicular to the bisector of C. We use m to denote the midpoint of the side of T_{uw} opposite u and α to denote the smaller unsigned angle between uw and um (see Figure 2.6). Note that for any pair of vertices u and w in the (constrained/ordered) θ_m-graph, there exist two canonical triangles: T_{uw} and T_{wu}. Also note that for each pair of vertices u and w in the (constrained) half-θ_6-graph, there exists a unique canonical triangle.

As we are dealing with constraints, some vertices may not be visible from a certain vertex u. To avoid confusion later on, we therefore explicitly define what it means for a region to be *empty*. We say that a region is *empty* if it does not contain any

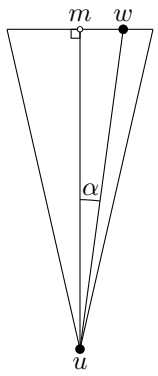

Figure 2.6: The canonical triangle T_{uw}

vertices (not just no vertices visible to some vertex u).

2.1 Routing

Next, we define our routing model. A routing algorithm is a deterministic k-local, m-memory routing algorithm, if the vertex to which a message is forwarded from the current vertex s is a function of s, t, $N_k(s)$, and M, where t is the destination vertex, $N_k(s)$ is the k-neighborhood of s and M is a memory of size m, stored with the message. The k-neighborhood of a vertex s is the set of vertices in the graph that can be reached from s by following at most k edges. For our purposes, we consider a unit of memory to consist of $\log_2 n$ bits or a point in $\mathbb{R}^2$ whose x- and y-coordinate consist of $\log_2 n$ bits. Our model also assumes that the only information stored at each vertex of the graph is $N_k(s)$ and the constraints the vertex is an endpoint of. The routing algorithms we consider are all 1-local and hence storing $N_k(s)$ as a list of vertices provides the same information as storing the k-local graph structure. This is not the case for $k > 1$. Finally, since our graphs are geometric, we identify each vertex by its coordinates in the plane. Unless otherwise noted, all routing algorithms we consider are deterministic 0-memory algorithms.

There are essentially two notions of *competitiveness* of a routing algorithm. One is to look at the Euclidean distance between the two vertices: A routing algorithm is *c-competitive with respect to the Euclidean distance* provided that the total distance traveled by the message is not more than c times the Euclidean distance between source and destination. Analogous to the spanning ratio, the *routing ratio* of an

algorithm *with respect to the Euclidean distance* is the smallest c for which it is c-competitive with respect to the Euclidean distance.

The other is to compare the routing path to the shortest path in the graph: A routing algorithm is *c-competitive with respect to the shortest path* provided that the total distance traveled by the message is not more than c times the length of the shortest path between source and destination. Analogous to the spanning ratio, the *routing ratio* of an algorithm *with respect to the shortest path* is the smallest c for which it is c-competitive with respect to the shortest path.

Since the length of the shortest path between two vertices is at least the Euclidean distance between the two vertices, an algorithm that is c-competitive with respect to the Euclidean distance is also c-competitive with respect to the shortest path. We use the stronger notion of competitiveness with respect to the Euclidean distance when proving upper bounds, since these imply competitiveness with respect to the shortest path. Similarly, we use the weaker notion of competitiveness with respect to the shortest path when proving lower bounds, since these imply competitiveness with respect to the Euclidean distance.

Finally, we define the θ-routing algorithm, which is commonly used to route on θ-graphs. Using the structure of the θ_m-graph, *θ-routing* is defined as follows. Let t be the destination of the routing algorithm and let u be the current vertex. If there exists a direct edge to t, follow this edge. Otherwise, follow the edge to the closest vertex in the cone of u that contains t.

2.2 Some Geometric Lemmas

Next, we prove a few geometric lemmas that are useful when bounding the spanning ratios of the various graphs. We start with a property of visibility graphs.

Lemma 2.1 *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.

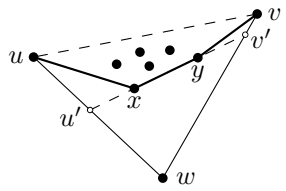

Figure 2.7: A convex chain from u to v and intersections u' and v' of the triangle and the line through x and y

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 2.7). 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 then this vertex would have been included in the convex hull and x and y would not be consecutive vertices on the convex hull. 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. □

Next, we proceed to prove a nice geometric property of the $\theta_{(4k+2)}$-graph that is used implicitly when bounding the spanning ratio of the various types of this graph in Chapters 3, 5, 6, and 7.

Lemma 2.2 *In the $\theta_{(4k+2)}$-graph, any line perpendicular to the bisector of a cone is parallel to the boundary of some cone.*

Proof. The angle between the bisector of a cone and the boundary of that cone is $\theta/2$. In the $\theta_{(4k+2)}$-graph, since $\theta = 2\pi/(4k+2)$, the angle between the bisector and the line perpendicular to this bisector is $\pi/2 = ((4k+2)/4) \cdot \theta = k \cdot \theta + \theta/2$. Thus the angle between the line perpendicular to the bisector and the boundary of the cone is $\pi/2 - \theta/2 = k \cdot \theta$. Since a cone boundary is placed at every multiple

of θ, the line perpendicular to the bisector is parallel to the boundary of some cone. $\square$

Finally, we prove a few other geometric lemmas that help to bound the spanning ratios of the various θ-graphs. In particular, Lemmas 2.3 and 2.4 reduce the number of cases we need to consider in the spanning proofs of these graphs in Chapters 5, 6, and 7 by essentially grouping cones $1, ..., k-1$ together. We use $\angle xyz$ to denote the smaller (unsigned) angle between line segments xy and yz.

Lemma 2.3 *Let a, b, c, and d be four points in clockwise order on a circle such that $\angle cad \leq \angle bad \leq \angle adc$. It holds that $|ac| + |cd| \leq |ab| + |bd|$ and $|cd| \leq |bd|$.*

Proof. This situation is illustrated in Figure 2.8. Without loss of generality, we assume that $|ad| = 1$. Since b and c lie on the same circle and $\angle abd$ and $\angle acd$ are the angle opposite to the same chord ad, the inscribed angle theorem implies that $\angle abd = \angle acd$. Furthermore, since $\angle cad \leq \angle adc$, c lies to the right of the perpendicular bisector of ad.

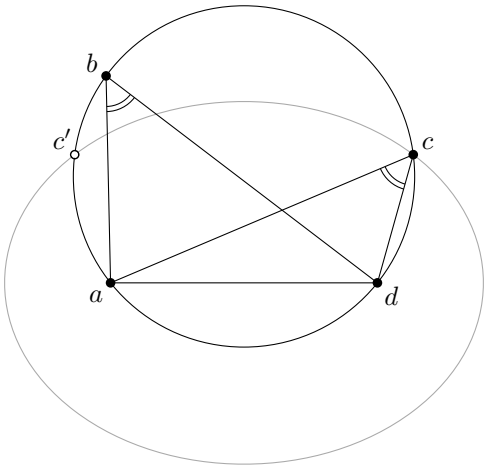

Figure 2.8: Illustration of the proof of Lemma 2.3

First, we show that $|ac| + |cd| \leq |ab| + |bd|$ by showing that $|ac| + |cd| + |ad| \leq |ab| + |bd| + |ad|$. Let c' be the point on the circle when we mirror c along the perpendicular bisector of ad. Points c and c' partition the circle into two arcs. Since $\angle cad \leq \angle bad \leq \angle adc$, b lies on the upper arc of the circle. We focus on triangle acd. The locus of the point c such that the perimeter of acd is constant defines an ellipse. This ellipse has major axis ad and goes through c and c'. Since this major axis is

horizontal, the ellipse does not intersect the upper arc of the circle. Hence, since b lies on the upper arc of the circle, which is outside of the ellipse, the perimeter of abd is greater than that of acd, completing the first half of the proof.

Next, we show that $|cd| \leq |bd|$. Using the sine law, we have that $|cd| = \sin\angle cad / \sin\angle acd$ and $|bd| = \sin\angle bad / \sin\angle abd$. Since $\angle cad \leq \angle bad \leq \angle adc \leq \pi - \angle cad$, we have that $\sin\angle cad \leq \sin\angle bad$. Hence, since $\angle abd = \angle acd$, we have that $|cd| \leq |bd|$. □

Lemma 2.4 *Let u, v and w be three vertices in the $\theta_{(4k+x)}$-graph, where $x \in \{2,3,4,5\}$, such that $w \in C_0^u$ and $v \in T_{uw}$, to the left of w. Let a be the intersection of the side of T_{uw} opposite to u with the left boundary of C_0^v. Let C_i^v denote the cone of v that contains w and let c and d be the upper and lower corner of T_{vw}. If $1 \leq i \leq k-1$, or $i = k$ and $|cw| \leq |dw|$, then $\max\{|vc| + |cw|, |vd| + |dw|\} \leq |va| + |aw|$ and $\max\{|cw|, |dw|\} \leq |aw|$.*

Proof. This situation is illustrated in Figure 2.9. We perform case distinction on $\max\{|cw|, |dw|\}$.

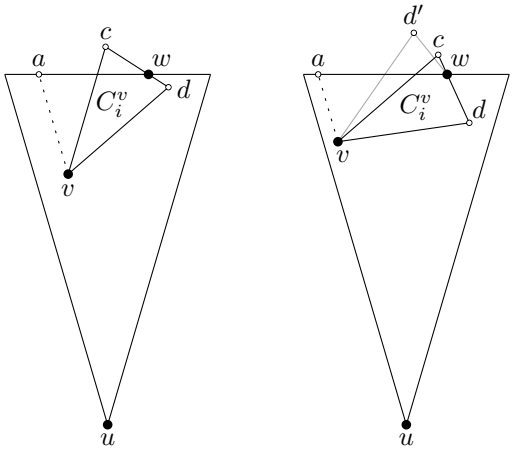

Figure 2.9: The two cases for the situation where we apply Lemma 2.3: (a) $|cw| > |dw|$, (b) $|cw| \leq |dw|$

Case 1: If $|cw| > |dw|$ (see Figure 2.9a), we need to show that when $1 \leq i \leq k-1$, we have that $|vc| + |cw| \leq |va| + |aw|$ and $|cw| \leq |aw|$. Since angles $\angle vaw$ and $\angle vcw$ are both angles between the boundary of a cone and the line perpendicular to its bisector, we have that $\angle vaw = \angle vcw$. Thus, c lies on the circle through a, v, and w. Therefore, if we can show that $\angle cvw \leq \angle avw \leq \angle vwc$, Lemma 2.3 proves this case.

We show $\angle cvw \leq \angle avw \leq \angle vwc$ in two steps. Since $w \in C_i^v$ and $i \geq 1$, we have that $\angle avc = i \cdot \theta \geq \theta$. Hence, since $\angle avw = \angle avc + \angle cvw$, we have that $\angle cvw \leq \angle avw$. It remains to show that $\angle avw \leq \angle vwc$. We note that $\angle avw \leq (i+1) \cdot \theta$ and $(\pi - \theta)/2 \leq \angle vwc$, since $|cw| > |dw|$. Using that $\theta = 2\pi/(4k+x)$ and $x \in \{2,3,4,5\}$, we have the following.

$$\begin{aligned} i &\leq k-1 \\ i &\leq k + \frac{x}{4} - \frac{3}{2} \\ i &\leq \frac{\pi \cdot (4k+x)}{4\pi} - \frac{3}{2} \\ i &\leq \frac{\pi}{2\theta} - \frac{3}{2} \\ (i+1) \cdot \theta &\leq \frac{\pi - \theta}{2} \\ \angle avw &\leq \angle vwc \end{aligned}$$

Case 2: If $|cw| \leq |dw|$ (see Figure 2.9b), we need to show that when $1 \leq i \leq k$, we have that $|vd| + |dw| \leq |va| + |aw|$ and $|dw| \leq |aw|$. Since angles $\angle vaw$ and $\angle vdw$ are both angles between the boundary of a cone and the line perpendicular to its bisector, we have that $\angle vaw = \angle vdw$. Thus, when we reflect d in the line through vw, the resulting point d' lies on the circle through a, v, and w. Therefore, if we can show that $\angle d'vw \leq \angle avw \leq \angle vwd'$, Lemma 2.3 proves this case.

We show $\angle d'vw \leq \angle avw \leq \angle vwd'$ in two steps. Since $w \in C_i^v$ and $i \geq 1$, we have that $\angle avw \geq \angle avc = i \cdot \theta \geq \theta$. Hence, since $\angle d'vw \leq \theta$, we have that $\angle d'vw \leq \angle avw$. It remains to show that $\angle avw \leq \angle vwd'$. We note that $\angle vwd' = \angle dwv = \pi - (\pi - \theta)/2 - \angle dvw$ and $\angle avw = \angle avd - \angle dvw = (i+1) \cdot \theta - \angle dvw$. Using that $\theta = 2\pi/(4k+x)$ and $x \in \{2,3,4,5\}$, we have the following.

$$\begin{aligned} i &\leq k \\ i &\leq k + \frac{x}{4} - \frac{1}{2} \\ i &\leq \frac{\pi \cdot (4k+x)}{4\pi} - \frac{1}{2} \\ i &\leq \frac{\pi}{2\theta} - \frac{1}{2} \\ (i+1) \cdot \theta - \angle dvw &\leq \frac{\pi + \theta}{2} - \angle dvw \\ \angle avw &\leq \angle vwd' \end{aligned}$$

□

The following lemma is used to reduce the number of analogous calculations we need to perform when bounding the spanning ratio of the various θ-graphs in Chapters 5, 6, and 7.

Lemma 2.5 *Let u, v and w be three vertices in the θ_m-graph with $m \geq 6$, such that $w \in C_0^u$, $v \in T_{uw}$ to the left of w, and $w \notin C_0^v$. Let a be the intersection of the side of T_{uw} opposite to u with the left boundary of C_0^v. Let c and d be the corners of T_{vw} opposite to v. Let $\beta = \angle awv$ and let γ be the unsigned angle between vw and the bisector of T_{vw}. Let $\boldsymbol{c}$ be a positive constant. If*

$$\boldsymbol{c} \geq \frac{\cos\gamma - \sin\beta}{\cos\left(\frac{\theta}{2} - \beta\right) - \sin\left(\frac{\theta}{2} + \gamma\right)}, \tag{2.1}$$

then

$$\max\{|vc| + \boldsymbol{c} \cdot |cw|, |vd| + \boldsymbol{c} \cdot |dw|\} \leq |va| + \boldsymbol{c} \cdot |aw|. \tag{2.2}$$

Proof. This situation is illustrated in Figure 2.10. Since the angle between the bisector of a cone and its boundary is $\theta/2$, by the sine law, we have the following.

$$\begin{aligned}
|vc| = |vd| &= |vw| \cdot \frac{\cos\gamma}{\cos\left(\frac{\theta}{2}\right)} \\
\max\{|cw|, |dw|\} &= |vw| \cdot \left(\sin\gamma + \cos\gamma \tan\left(\frac{\theta}{2}\right)\right) \\
|va| &= |vw| \cdot \frac{\sin\beta}{\cos\left(\frac{\theta}{2}\right)} \\
|aw| &= |vw| \cdot \left(\cos\beta + \sin\beta \tan\left(\frac{\theta}{2}\right)\right)
\end{aligned}$$

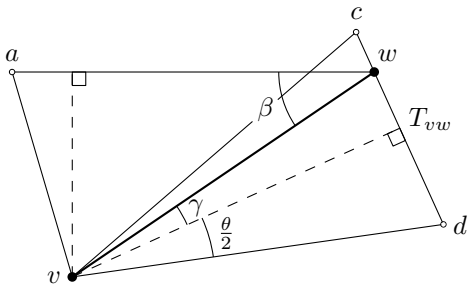

Figure 2.10: Finding a constant $\boldsymbol{c}$ such that $|vd| + \boldsymbol{c} \cdot |dw| \leq |va| + \boldsymbol{c} \cdot |aw|$

To show that (2.2) holds, we first multiply both sides by $\cos(\theta/2)/|vw|$ and rewrite as follows.

$$\begin{aligned}\frac{\cos\left(\frac{\theta}{2}\right)}{|vw|} \cdot \max\{|vc| + \boldsymbol{c} \cdot |cw|, |vd| + \boldsymbol{c} \cdot |dw|\} & \\ &= \cos\gamma + \boldsymbol{c} \cdot \left(\sin\gamma\cos\left(\frac{\theta}{2}\right) + \cos\gamma\sin\left(\frac{\theta}{2}\right)\right) \\ &= \cos\gamma + \boldsymbol{c} \cdot \sin\left(\frac{\theta}{2} + \gamma\right)\end{aligned}$$

$$\begin{aligned}\frac{\cos\left(\frac{\theta}{2}\right)}{|vw|} \cdot (|va| + \boldsymbol{c} \cdot |aw|) &= \sin\beta + \boldsymbol{c} \cdot \left(\cos\beta\cos\left(\frac{\theta}{2}\right) + \sin\beta\sin\left(\frac{\theta}{2}\right)\right) \\ &= \sin\beta + \boldsymbol{c} \cdot \cos\left(\frac{\theta}{2} - \beta\right)\end{aligned}$$

Therefore, to prove that (2.1) implies (2.2), we rewrite (2.1) as follows.

$$\begin{aligned}\boldsymbol{c} &\geq \frac{\cos\gamma - \sin\beta}{\cos\left(\frac{\theta}{2} - \beta\right) - \sin\left(\frac{\theta}{2} + \gamma\right)} \\ \cos\gamma - \sin\beta &\leq \boldsymbol{c} \cdot \left(\cos\left(\frac{\theta}{2} - \beta\right) - \sin\left(\frac{\theta}{2} + \gamma\right)\right) \\ \cos\gamma + \boldsymbol{c} \cdot \sin\left(\frac{\theta}{2} + \gamma\right) &\leq \sin\beta + \boldsymbol{c} \cdot \cos\left(\frac{\theta}{2} - \beta\right)\end{aligned}$$

It remains to show that $\boldsymbol{c} > 0$. Since $w \notin C_0^v$, we have that $\beta \in (0, (\pi - \theta)/2)$. Moreover, we have that $\gamma \in [0, \theta/2)$, by definition. This implies that $\sin(\pi/2 + \gamma) > \sin\beta$, or equivalently, $\cos\gamma - \sin\beta > 0$. Thus, we need to show that $\cos(\theta/2 - \beta) - \sin(\theta/2 + \gamma) > 0$, or equivalently, $\sin(\pi/2 + \theta/2 - \beta) > \sin(\theta/2 + \gamma)$. It suffices to show that $\theta/2 + \gamma < \pi/2 + \theta/2 - \beta < \pi - \theta/2 - \gamma$. This follows from $\beta \in (0, (\pi - \theta)/2)$, $\gamma \in [0, \theta/2)$, and the fact that $\theta \leq 2\pi/7$. □

2.3 Conclusion

In this chapter, we formally introduced the various θ-graphs that are studied in the remainder of this thesis. We also defined two routing models, one in terms of the Euclidean distance between the source and destination and one in terms of the shortest path in the graph. This distinction allows us to prove stronger upper and lower

bounds on the routing ratio in Chapter 4. Finally, we proved a number of geometric lemmas that are used to bound the spanning ratio of the various θ-graphs discussed in Chapters 5, 6, and 7.

Chapter 3

Constrained Bounded-Degree Spanners

In this chapter, 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.* [18]. A key difficulty in proving the latter stems from the fact that the constrained Delaunay graph is **not** necessarily a triangulation. We then generalize the elegant construction of Bonichon *et al.* [7] 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.

The results presented in this chapter are joint work with Prosenjit Bose, Rolf Fagerberg and Sander Verdonschot and appear in the Proceedings of the 10th Latin American Symposium on Theoretical Informatics (LATIN 2012) [19].

3.1 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 $Vis(P, S)$.

Theorem 3.1 *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 3.1). 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 3.2).

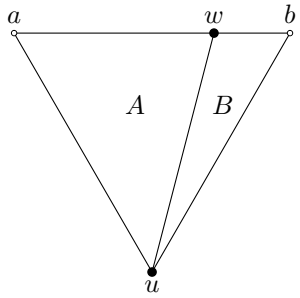

Figure 3.1: Triangles A and B

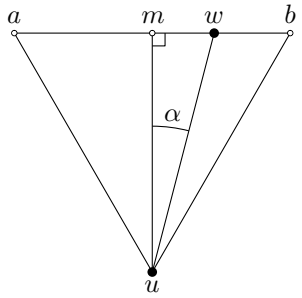

Figure 3.2: Canonical triangle T_{uw}

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.

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 3.3). 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.

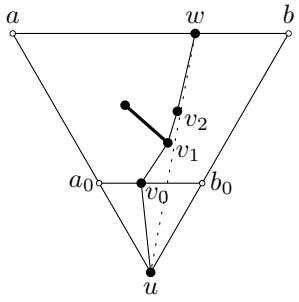

Figure 3.3: Convex chain from v_0 to w

Since uw and uv_0 are visibility edges, by applying Lemma 2.1 to triangle v_0uw, a convex chain $v_0, ..., v_k = w$ of visibility edges connecting v_0 and w exists (see Figure 3.3). 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 3.4):

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 2.1. 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 $|v_{i-1}a_i| + |a_iv_i|$.

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 $|v_{i-1}a_i|+|a_iv_i|$.

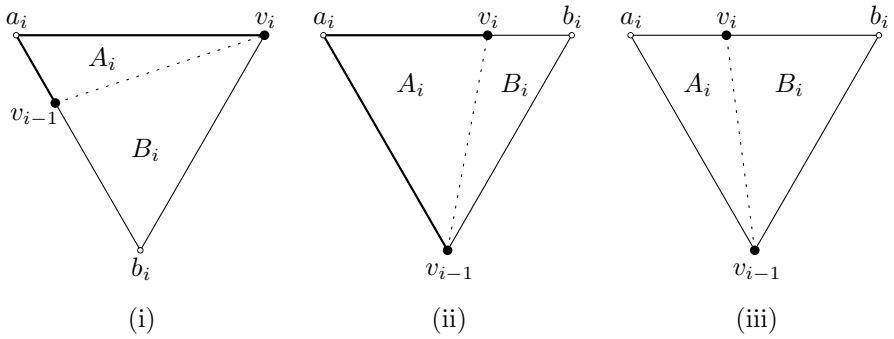

Figure 3.4: Charging the three types of configurations

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 3.1). 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 3.5):

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

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

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 2.1), all triangles B_i are also empty. Hence using the

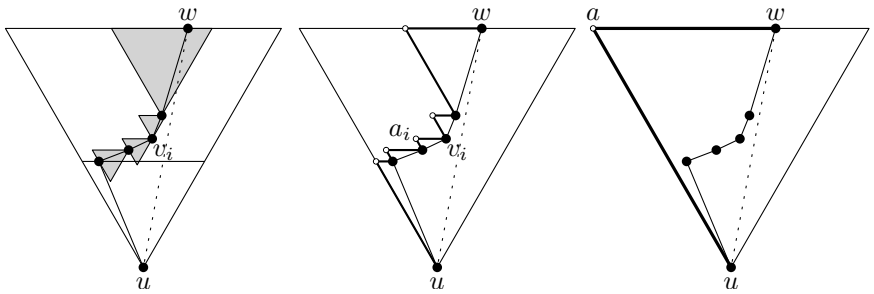

Figure 3.5: Visualization of the paths (thick lines) in the inequalities of case (a)

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 3.6). Note that since $\angle awu > \pi/2$ and v_j lies to the left of uw, $|a'v_j|$ is smaller than $|b'v_j|$.

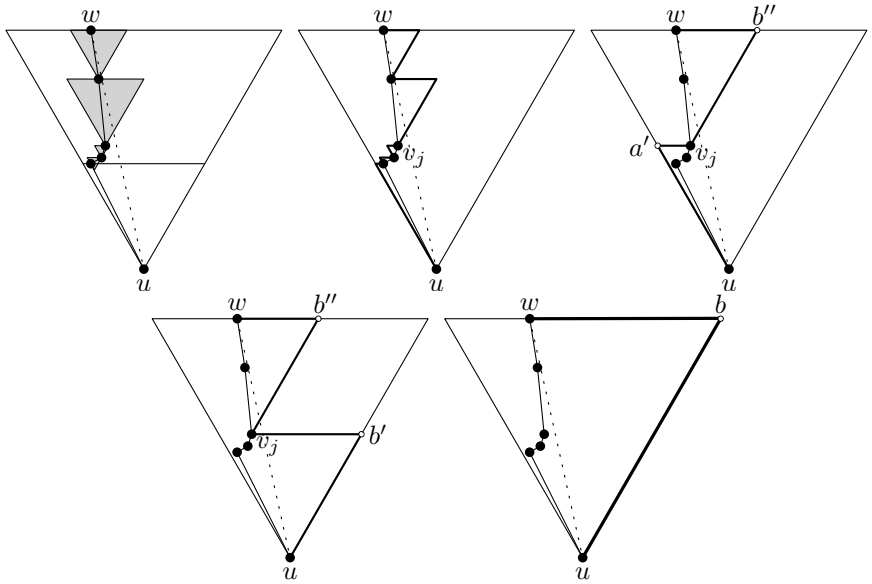

Figure 3.6: Visualization of the paths (thick lines) in the inequalities of case (c)

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

□

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 3.1 *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 3.1 *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}. □

Lemma 3.2 *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 3.1 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}.

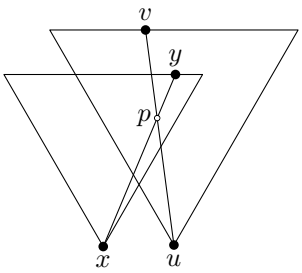

Figure 3.7: Edges uv and xy intersect at point p

Assume without loss of generality that $v \in C^u_{0,j}$ and $y \in T_{uv}$ (see Figure 3.7). If $y \in C^u_{0,j}$, we look at triangle upy. Since both u and y can see p, we get by Lemma 2.1 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. □

3.2 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 3.8). 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.* [7].

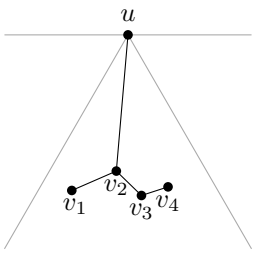

Figure 3.8: The edges that are added to G_9 for a negative subcone of a vertex u with canonical sequence v_1, v_2, v_3 and v_4

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.* [7]. We proceed to prove that G_9 is a spanning subgraph of the constrained half-θ_6-graph with spanning ratio 3.

Lemma 3.3 *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 2.1 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 P empty of vertices and constraints.

Since P 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 P 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 $C_0^{x_i}$ 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 P is empty, y must therefore lie strictly outside of P, 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_2^v (the case that v lies in C_1^w is similar). We need to show that w is the closest visible vertex in subcone $C_{2,j}^v$. We prove this by contradiction, so assume another vertex x in $C_{2,j}^v$ is the closest. Vertex x lies in T_{vw}, which is partitioned into a part inside P, a part to the right of wu, and a part below vw (see Figure 3.9). 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 P is empty, x must lie below vw (see Figure 3.9).

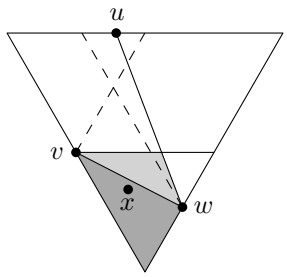

Figure 3.9: T_{vw} is partitioned into a part inside P (light gray), a part to the right of wu (white), and a part below vw (dark gray)

Applying Lemma 2.1 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 Q. Vertex x_1 cannot lie in C_0^x, as this would contradict that x is the closest visible vertex to v in $C_{2,j}^v$. Hence, since P and Q 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_0^x that contains

u. Neither v nor the convex chain from x to w lie in C_0^x. As P and Q 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 3.3 the following two properties were shown.

Corollary 3.2 *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 3.3 *Let u, v, and w be three vertices such that v and w are neighbors along a canonical path of u in $\overline{C}_i^u$. Triangle uvw is empty and does not contain any constraints.*

Theorem 3.2 *G_9 is a 3-spanner of the constrained half-θ_6-graph.*

Proof. We prove the lemma by showing that for each edge uw of the constrained half-θ_6-graph H that is not part of G_9, $d_{G_9}(u, w) \leq 3 \cdot d_H(u, w)$.

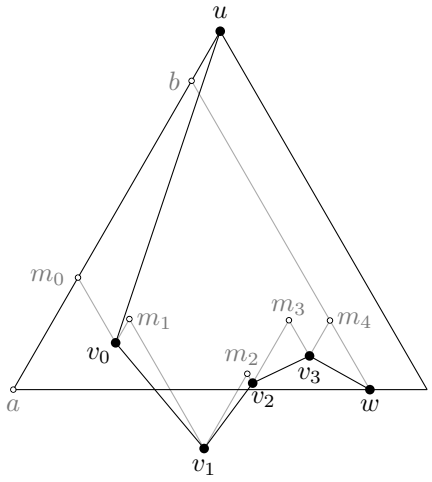

Figure 3.10: Bounding the length of the canonical path

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 3.10). 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 3.2 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. □

It follows from Theorems 3.1 and 3.2 that G_9 is a 6-spanner of the visibility graph.

Corollary 3.4 *G_9 is a 6-spanner of the visibility graph.*

To bound the degree of a vertex, we use a charging scheme that charges the edges added during the construction to the cones of that vertex. We look at the canonical path in $\overline{C}^u_{i,j}$, created by a vertex u. We use v to indicate an arbitrary vertex along the canonical path. Let v' be the closest vertex to u along the canonical path. The edges of G_9 are charged as follows:

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

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 3.11).

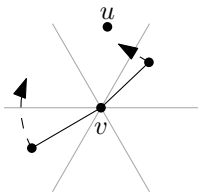

Figure 3.11: Two edges of a canonical path and the associated charges

Note that each edge is charged once to each of its endpoints and therefore the charge on a vertex is an upper bound on its degree. Also note that, since charges to a vertex v come from a canonical path it is part of, we can compute the total charge of a vertex by summing the charges of the canonical paths it is part of. By Corollary 3.2 no edge on the canonical path can lie in C^v_i or $\overline{C}^v_i$, so the charging scheme above is exhaustive.

Lemma 3.4 *Let v be a vertex that is incident to at least two constraints in the same positive cone C^v_i. Let $C^v_{i,j}$ be a subcone between two constraints and let u be the closest visible vertex in this subcone. Let $\overline{C}^u_{i,k}$ be the subcone of u that contains v and (when uv is a constraint) intersects $C^v_{i,j}$. Then v is the only vertex on the canonical path of $\overline{C}^u_{i,k}$.*

Proof. Let vw_1 and vw_2 be the two constraints between which subcone $C^v_{i,j}$ lies. By applying Lemma 2.1 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 $P \subset C^v_{i,j}$ empty of vertices and constraints. Since u is the closest vertex visible to v inside $C^v_{i,j}$, 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 P is empty, there can be no vertex other than v in $\overline{C}^u_{i,k}$ from which u is visible. Hence, v is the only vertex on the canonical path of $\overline{C}^u_{i,k}$. □

Lemma 3.5 *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^v_i.*

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^v_i, then u must lie in C^v_i. Assume u lies in C^v_j, $j \neq i$. Since all charges of this canonical path are shifted one cone towards C^v_j, a charge to C^v_i would have to come from $\overline{C}^v_j$. However, by Corollary 3.2, no edge on the canonical path of a vertex in C^v_j can lie in $\overline{C}^v_j$.

Next, we observe that there can be only one such vertex u for each subcone of C^v_i. 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^v_i is a single subcone and v is not the closest vertex to u on its canonical path, C^v_i 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^v_i 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 3.4. 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 3.2 that y cannot lie in $\overline{C}^v_i$. Since y must also lie in $\overline{C}^u_i$ to be on this canonical path, vertex u is not be visible from y due to the constraint xv. Hence, no such vertex 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 $\max\{2, c_i(v)+1\}$. □

Corollary 3.5 *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 3.6 *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 3.12).

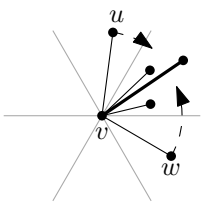

Figure 3.12: If vw is present, the negative cone does not contain edges having v as endpoint

We first show that vertices that do not lie in the positive subcones directly adjacent to $\overline{C}_i^v$ 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 $\overline{C}_i^v$ and let vx be the constraint closest to $\overline{C}_i^v$ 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 $\overline{C}_i^v$, 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.

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 $\overline{C}_i^v$ from the canonical path created by the closest visible vertices in the two positive subcones directly adjacent to $\overline{C}_i^v$. 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 3.3, 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 3.3, this triangle is empty. Furthermore, since uw is an edge of the constrained half-θ_6-graph and, by Lemma 3.2, 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 3.3 *Every vertex v in G_9 has degree at most $c(v) + 9$.*

Proof. From Lemmas 3.5 and 3.6, 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 $c(v) + 9$. □

3.3 Bounding the Maximum Degree Further

Using Lemmas 3.5 and 3.6, if we can avoid the situation where a positive cone gets a charge of $c_i(v) + 2$, with $c_i(v) = 0$, every positive cone is charged at most $c_i(v) + 1$, giving a total charge of $c(v) + 6$. By Corollary 3.5, a positive cone gets a charge of $c_i(v) + 2$, with $c_i(v) = 0$, when it does not contain any constraints having v as an endpoint and is charged for two edges in the adjacent negative cones (see Figure 3.13).

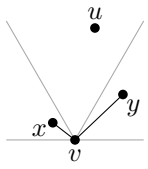

Figure 3.13: A positive cone having charge 2

For brevity, 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.

If x is the closest canonical vertex to v 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 $\overline{C}_1^v$ adjacent to C_0^v, 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 can assume that neither x nor y are the closest canonical vertex in the subcones that contain them. In that case, we perform the following transformation on G_9. 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 canonical path induced by u. If this sequence includes x, we remove vy. Otherwise we remove vx. Note that by Corollary 3.3, 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 note that by Corollary 3.3, triangle uvx is empty. Hence, since the constrained half-θ_6-graph is planar and u does not lie in $\overline{C}_2^v$, x is the last vertex along the canonical path of v in $\overline{C}_2^v$ and thus w is uniquely defined. 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.

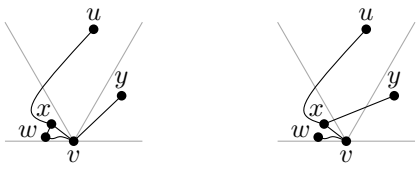

Figure 3.14: Transforming G_9 (a) into G_6 (b)

The situation before the transformation is shown in Figure 3.14 (a) and the situation after the transformation is shown in Figure 3.14 (b), where a curved line segment denotes the part of the canonical path that lies between its endpoints. To construct G_6, we apply this transformation on each positive cone with a charge of 2. 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 3.7 *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 contradiction. Given that x is not the closest canonical vertex to v in the subcones of $\overline{C}_i^v$ that contain it, 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 3.2. By Corollary 3.3 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 3.15).

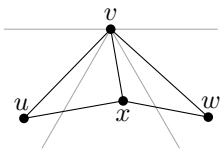

Figure 3.15: If edge vx is part of two canonical paths, x is the only vertex in the negative cone of v

Thus uvx and wvx form two disjoint triangles in the constrained half-θ_6-graph and, by Corollary 3.3, 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 $\overline{C}_{i,j}^v$, which is a contradiction. □

Lemma 3.8 *G_6 is a 3-spanner of the half-θ_6-graph.*

Proof. In the proof of Theorem 3.2 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 the canonical path and the edge between u and the closest canonical vertex of the subcone of u that contains

w. We now show that G_6 also contains a spanning path between u and w of length at most $3 \cdot |uw|$.

Lemma 3.7 tells us that any edge vx, where x lies in a negative cone of v, is part of a single canonical path, provided that x is not a closest canonical vertex of v. Since the construction of G_6 removes only edges meeting these requirements, for each removed edge 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 3.14, and Type 2, represented by xw in Figure 3.14. We first show that neither of these types removes edge vx. A Type 1 removal that has v as the middle vertex in the configuration, as shown in Figure 3.14, is called *centered at* v. A Type 1 removal of vy affects the canonical path containing x, v, and y (see Figure 3.14). 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 3.14). As argued during the construction of G_6, x is the last vertex along a canonical path of v and it is removed if w lies in a negative cone of x and x is not a closest canonical vertex to w. 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 present in 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 3.14, 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. 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.

Finally, we need to show that for every edge uw in G_9 that was removed during the transformation, there still exists a spanning path between u and w. Since every edge that is removed during the transformation is not an edge to a closest canonical vertex, Theorem 3.2 implies that G_9 contained a spanning path between u and w, which in turn implies that G_6 contains a transformed spanning path of length at most $3 \cdot |uw|$. Hence, every edge of the constrained half-θ_6-graph is approximated by a spanning path of length at most $3 \cdot |uw|$. □

Lemma 3.9 *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 3.5, 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 3.13). 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. We look at vertex w, the neighbor of x on the canonical path of v in $\overline{C}_{i-1}^v$. Since x is not the closest canonical vertex to v in a subcone of $\overline{C}_{i-1}^v$, the canonical path and vertex w exist. Since vertices uvx form a triangle in the constrained half-θ_6-graph, x cannot be incident to any constraints in C_{i-1}^x. Furthermore, since x is the last vertex along the canonical path of v, C_{i-1}^x has charge at most 1. 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 replace the charge of C_{i-1}^x by xy and the charge of C_{i-1}^x remains 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 is removed. Since this edge was charged to C_{i-1}^x, we can charge xy to C_{i-1}^x and the charge of C_{i-1}^x remains 1. □

Lemma 3.10 *G_6 is plane.*

Proof. Since the constrained half-θ_6-graph is plane, we focus on the added edges. Let xy be an edge added because a positive cone of vertex v had charge 2 on a canonical path of vertex u. This edge can potentially intersect edges of the half-θ_6-graph, edges that were added, and constraints.

Since xy was added, x, y, and v are part of the canonical path of some vertex u. Thus in the constrained half-θ_6-graph uvx and uvy form two triangles. 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. By construction, however, uv is not part of G_6. Hence xy cannot intersect any edge of the constrained half-θ_6-graph.

Next, we look at the added edges. Since by Corollary 3.3 $uxvy$ does not contain any vertices and xv and vy are edges of the constrained half-θ_6-graph, which we showed cannot be intersected by added edges, v must be an endpoint of the added edge that intersects xy. However, v lies in a positive cone of both x and y, so the edges are charged to the adjacent negative cones. Thus xy cannot intersect an added edge.

Finally, we look at constraints. By Corollary 3.3, 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. □

3.4 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.* [7] 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 we show in Chapter 6 that constrained θ-graphs with at least 6 cones are spanners, a logical next question is to see if the method shown in this chapter 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 plane spanner of $Vis(P, S)$, especially in light of recent results in the unconstrained setting [9] that show how to build a plane spanner of maximum degree 4.

Chapter 4

Routing on Constrained θ_6-Graphs

In this chapter, we consider how to route on the constrained θ_6-graph. Recall that an algorithm is 1-local when the routing algorithm must decide which vertex to forward a message to based solely on knowledge of the source and destination vertex, the current vertex and all neighbors the current vertex. We first show that no deterministic 1-local routing algorithm is $o(\sqrt{n})$-competitive on all pairs of non-visible vertices of the constrained θ_6-graph. After that, we show how to route between any two visible vertices of the constrained θ_6-graph using only 1-local information. Our routing algorithm guarantees that the returned path has length at most 2 times the Euclidean distance between the source and destination. To the best of our knowledge, this is the first local routing algorithm in the constrained setting with guarantees on the length of the returned path.

The results presented in this chapter are joint work with Prosenjit Bose, Rolf Fagerberg and Sander Verdonschot and have been submitted to the 26th ACM-SIAM Symposium on Discrete Algorithms (SODA 2015) [20].

4.1 Lower Bound on Local Routing Between Non-Visible Vertices

In this section we modify and generalize the proof by Bose *et al.* [10] that shows that no deterministic routing algorithm is $o(\sqrt{n})$-competitive for all triangulations to show that no deterministic 1-local routing algorithm is $o(\sqrt{n})$-competitive on all pairs of non-visible vertices in the constrained θ_6-graph, regardless of the amount of memory it is allowed to use.

Recall that when working on lower bounds, we use the notion of competitiveness

with respect to the shortest path: A routing algorithm is *c-competitive with respect to the shortest path* provided that the total distance traveled by the message is not more than c times the length of the shortest path between source and destination. The *routing ratio* of an algorithm *with respect to the shortest path* is the smallest c for which it is c-competitive with respect to the shortest path.

Lemma 4.1 *No deterministic 1-local routing algorithm is $o(\sqrt{n})$-competitive on all pairs of non-visible vertices of the θ_6-graph, regardless of the amount of memory it is allowed to use.*

Proof. Consider an $n \times n$ grid and shift every second row to the right by half a unit. We stretch the grid in horizontal direction, such that each edge has length n (see Figure 4.1). Next, we replace each horizontal edge by a constraint to prevent vertical visibility edges. Finally, we add two additional vertices, origin s and destination t, centered horizontally at one unit below the bottom row and one unit above the top row.

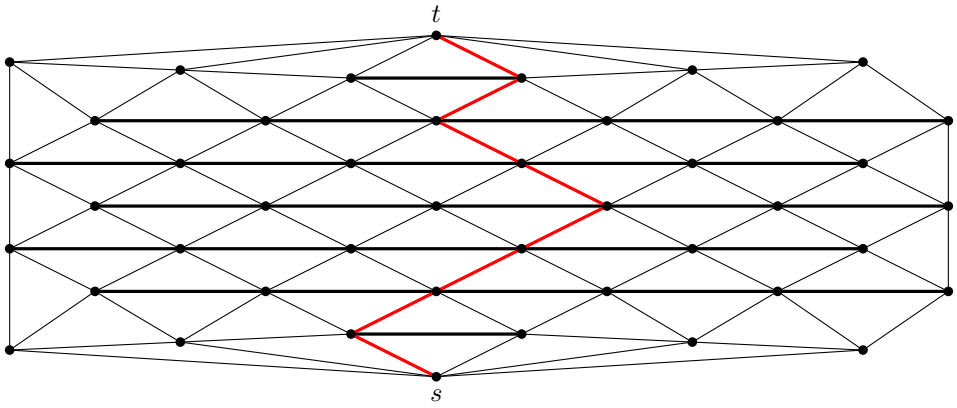

Figure 4.1: The constrained θ_6-graph starting from a hexagonal grid, using horizontal constraints to block vertical edges, and the red path of the routing algorithm

To conform to our general position assumption, we move all vertices by at most some arbitrarily small amount ϵ, such that no two points define a line parallel to one of the rays that define the cones and no three points are collinear. In particular, we ensure that all vertices on the bottom row have s as the closest vertex in one of their subcones and all vertices on the top row have t as the closest vertex in one of their subcones. This can be achieved by placing the bottom row on the upper hull of an

ellipse and placing the top row on the lower hull of an ellipse. Using this point set and these constraints, we build the constrained θ_6-graph G.

Consider any deterministic 1-local ∞-memory routing algorithm and let π be the path this algorithm takes when routing from s to t. We note that by construction, π consists of at least $n+1$ steps. If π consists of at least $n\sqrt{n}$ steps, the total length of the path is at least $n\sqrt{n} \cdot n = n^2\sqrt{n}$, since we stretched the grid such that each horizontal edge has length n. However, G contains a path of length $O(n^2)$ between s and t. Hence, in this case, the local routing algorithm is not $o(\sqrt{n})$-competitive.

Now, assume that π consists of $f(n)$, less than $n\sqrt{n}$, steps. Consider the $2\sqrt{f(n)}$ vertical lines through the neighbors of s at horizontal distance at most $n\sqrt{f(n)}$ from s. Since $f(n)$ is less than $n\sqrt{n}$ and G is a modified $n \times n$ grid, these $2\sqrt{f(n)}$ vertical lines exist. Now, let π' be the routing path π minus vertices s and t. We say that a vertex of π' *touches* a vertical line if it has a neighbor on that line. Hence, any vertex along π' touches at most 2 vertical lines. This implies that the total number of lines touched by the vertices along π' is at most $2f(n)$. Hence, on average, a line is touched $2f(n)/2\sqrt{f(n)} \leq \sqrt{f(n)}$ times. This implies that there exists a vertical line that is touched at most $\sqrt{f(n)}$ times. Let u be the neighbor of s on this vertical line.

We now create a new constrained θ_6-graph G' such that the deterministic 1-local routing algorithm follows the same path, but G' contains a short 'almost vertical' path via u. We start with s, t, and all vertices of π. Next, we add all vertices and constraints connected to these vertices in G. Using this point set and these constraints, we build the constrained θ_6-graph G' (see Figure 4.2).

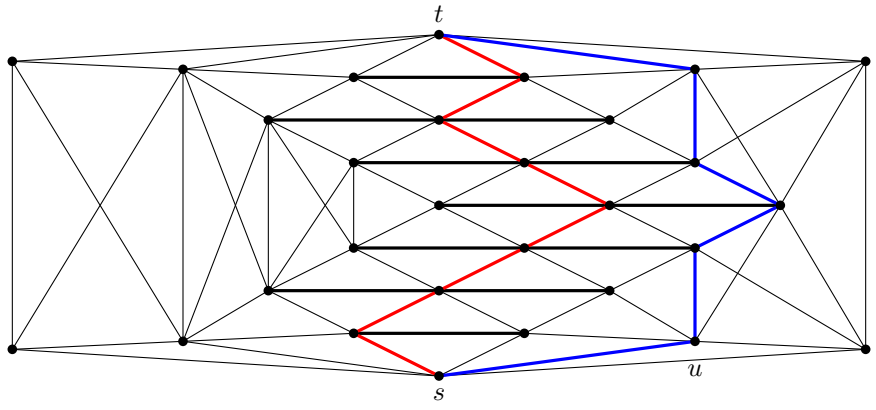

Figure 4.2: The constrained θ_6-graph that looks the same from the red path of the routing algorithm, but has an almost vertical blue path

Since the horizontal distance between vertices is far larger than their vertical distance, an 'almost vertical' path from u to the top row of G' is formed. This almost vertical path is a path that is vertical whenever possible and uses detours to avoid path π (see Figure 4.3): If π arrives at a vertex v that has a neighbor on the vertical line through u, we avoid π by following one edge away from π, followed by an edge back to the vertical line through u (see Figure 4.3a). If π arrives at a vertex on the vertical line through u, we avoid the vertex before and after v on π as before, and meet π at v (see Figure 4.3b). Since the vertical line through u is touched by at most $\sqrt{f(n)}$ vertices of π, only $O\left(\sqrt{f(n)}\right)$ of these detour edges are required. Hence, the length of the path from u to the top row of G' has length $O\left(n\sqrt{f(n)}\right)$. Thus, G' contains a path from s to t of length $O\left(n\sqrt{f(n)}\right)$: Follow the edge from s to u, follow the almost vertical path from u to the top row of G', and follow the edge to t.

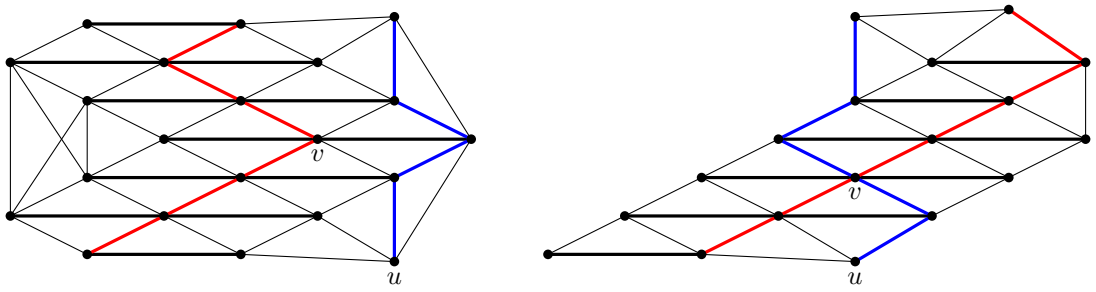

(a) (b)

Figure 4.3: The two types of detour: (a) when π does not visit the vertical line through u, (b) when π visits the vertical line through u

Since the 1-local routing algorithm is deterministic and the 1-local information of the vertices of π in G' is the same as in G, the algorithm stores the same data in its memory and follows the same path. The total length of this path is $n \cdot f(n)$, since it does not follow any of the shortcut edges. Hence, the deterministic 1-local routing algorithm is not $o\left(\sqrt{f(n)}\right)$-competitive on G'. Thus, since $f(n)$ is at least $n+1$, it is not $o(\sqrt{n})$-competitive. Furthermore, since G' can be constructed for any deterministic 1-local routing algorithm, we have shown that no deterministic 1-local routing algorithm is $o(\sqrt{n})$-competitive on all pairs of non-visible vertices. □

4.2 Routing on the Constrained θ_6-Graph

In this section, we provide a 1-local routing algorithm on the constrained θ_6-graph for any pair of visible vertices. Since the constrained θ_6-graph is the union of two constrained half-θ_6-graphs, we describe a routing algorithm for the constrained half-θ_6-graph for the case where the destination t lies in a positive subcone of the origin s. After describing this algorithm and proving that it is 2-competitive, we describe how to use it to route 1-locally on the constrained θ_6-graph.

Recall that when working on upper bounds, we use the notion of competitiveness with respect to the Euclidean distance: A routing algorithm is *c-competitive with respect to the Euclidean distance* provided that the total distance traveled by the message is not more than c times the Euclidean distance between source and destination. The *routing ratio* of an algorithm *with respect to the Euclidean distance* is the smallest c for which it is c-competitive with respect to the Euclidean distance.

4.2.1 Positive Routing on the Constrained Half-θ_6-Graph

Before describing how to route on the constrained half-θ_6-graph when t lies in a positive subcone of s, we first recall that there exists a path in canonical triangle T_{st}.

Lemma 4.2 *Given two vertices u and w such that u and w are visible to each other and w lies in a positive subcone $C^u_{i,j}$, there exists a path between u and w in T_{uw}.*

Proof. This follows from the fact that the spanning path between two visible vertices u and w, as constructed in Theorem 3.1, lies inside the canonical triangle T_{uw}. □

Positive Routing Algorithm for the Constrained Half-θ_6-Graph
Next, we describe how to route from s to t, when s can see t and t lies in a positive subcone $C^s_{i,j}$ (see Figure 4.4): When we are at s, we follow the edge to the closest vertex in the subcone that contains t. When we are at any other vertex u, we look at all edges in the subcones of C^u_i and all edges in the subcones of the adjacent negative cone $\overline{C}^u$ that is intersected by st. An edge in a subcone of $\overline{C}^u$ is considered only if it does not cross st. For example, in Figure 4.4, we do not consider the edge to v_1 since it lies in $\overline{C}^u$ and crosses st. It follows that we can cross st only when we follow an edge in C^u_i.

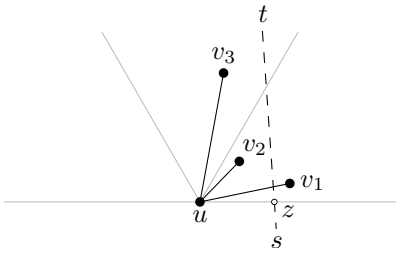

Figure 4.4: An example of routing from s to $t \in C_0^s$. The dashed line represents the visibility line between s and t

Let z be the intersection of st and the boundary of $\overline{C}^u$ that is not a boundary of C_i^u. We follow the edge uv that minimizes the unsigned angle $\angle zuv$. For example, in Figure 4.4, when we are at vertex u we follow the edge to v_2 since, out of the two remaining edges uv_2 and uv_3, $\angle zuv_2$ is smaller than $\angle zuv_3$. We note that since the algorithm uses only information about the location of s and t and the neighbors of the current vertex, it is a 1-local routing algorithm.

We proceed by proving that the above routing algorithm can always perform a step, i.e. at every vertex reached by the algorithm there exists an edge that is considered by the algorithm.

Lemma 4.3 *The routing algorithm can always perform a step.*

Proof. Given two vertices s and t such that s and t are visible to each other, we assume without loss of generality that $t \in C_0^s$. We maintain the following invariant (see Figure 4.5):

> **Invariant** Let x be the last intersection of an edge of the routing path with st (initially x is s), let $v_0, ..., v_k$ denote the endpoints of the edges following x as selected by the algorithm, and let x' be the intersection of st and the horizontal line through v_k. The simple polygon defined by $x, v_0, ..., v_k, x'$ is empty and does not contain any constraints.

When the routing algorithm starts at s, it looks at the subcone that contains t. Since t is visible from s, this subcone contains at least one visible vertex. Hence, it also contains a closest visible vertex v_0 and by construction, s has an edge to v_0. Therefore, when the routing algorithm starts at s, it can follow an edge.

To see that the invariant is satisfied, we need to show that triangle sv_0x' is empty and does not contain any constraints in its interior. By construction s cannot be

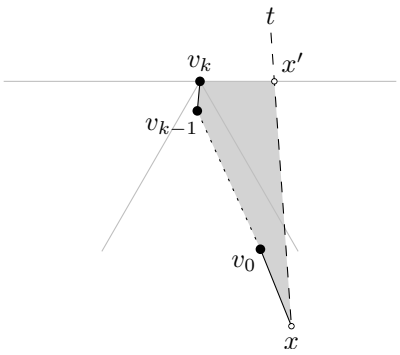

Figure 4.5: By the invariant, the gray region is empty and does not contain any constraints

the endpoint of any constraints in the interior of sv_0x', hence since sx' and sv_0 are visibility edges, any constraint has at least one endpoint in sv_0x'. Thus, it suffices to show that sv_0x' is empty. We prove this by contradiction, so assume that it is not empty. Since sv_0 and sx' are visibility edges and by construction s is not the endpoint of a constraint intersecting the interior of sv_0x', Lemma 2.1 gives us a convex chain of visibility edges between v_0 and x'. Since the region bounded by sv_0, sx', and this chain is empty and does not contain any constraints, the vertex along this chain that is closest to s is visible to s. However since every point in sv_0x' is closer to s than v_0, this contradicts the fact that v_0 is the closest visible vertex to s. Hence, triangle sv_0x' must be empty and the invariant is satisfied.

When the routing algorithm is at vertex u ($u \neq s$), we assume without loss of generality that u lies to the left of st. Let h be the halfplane below the horizontal line through t and let h' be the halfplane to the left of st. We need to show that u has at least one edge in the union of $C_0^u \cap h$ and $\overline{C}_1^u \cap h \cap h'$. We first show that there exists a vertex that is visible to u in the union of $C_0^u \cap h \cap h'$ and $\overline{C}_1^u \cap h \cap h'$. Since t lies in this region, we know that it is not empty. Consider all vertices in this region and let v be the vertex in this region that minimizes $\angle x'uv$. Note that we did not require there to be an edge between u and v. Since v minimizes $\angle x'uv$ and no constraint can cross st or ux', it is visible from u. We consider two cases: v lies in a subcone of C_0^u and v lies in a subcone of $\overline{C}_1^u$.

If v lies in $C_0^u \cap h$, it follows from Lemma 4.2 and the fact that v is visible from u that there exists a path between u and v that lies inside T_{uv}. Since T_{uv} is contained in $C_0^u \cap h$, there exists an edge in $C_0^u \cap h$ and the routing algorithm can perform a step.

If v lies in $\overline{C}_1^u \cap h$, it follows from Lemma 4.2 and the fact that v is visible from u that there exists a path between u and v that lies inside T_{vu}. Canonical triangle T_{vu} intersects three cones of u (see Figure 4.6): C_0^u, $\overline{C}_1^u$, and C_2^u. Since the routing algorithm follows edges in C_0^u or $\overline{C}_1^u$, the edge to u reaches u from either $\overline{C}_0^u$ or C_1^u. This implies that $T_{vu} \cap C_2^u$ is contained in the region of the invariant and is therefore empty. Hence, the first edge on the path from u to v lies in either $C_0^u \cap h$ or $\overline{C}_1^u \cap h$ and the algorithm can perform a step.

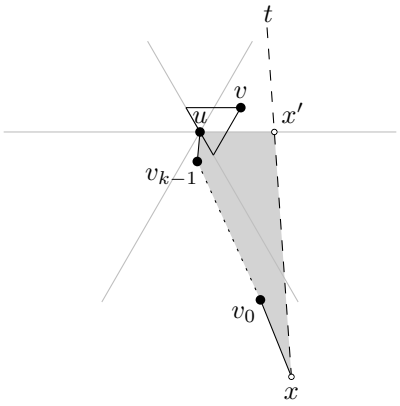

Figure 4.6: By the invariant, the gray region is empty, so the path between u and v lies inside $T_{vu} \cap (C_0^u \cup \overline{C}_1^u)$

It remains to show that after the algorithm takes a step, the invariant is satisfied at the new vertex v. Let uv be the edge that the algorithm followed and let x'' be the intersection of st and the horizontal line through v. We consider three cases (see Figure 4.7): (a) v lies in a subcone of $\overline{C}_1^u$, (b) v lies in a subcone of C_0^u and uv does not cross st, and (c) v lies in a subcone of C_0^u and uv crosses st.

Case (a): If v lies in a subcone of $\overline{C}_1^u$, we need to show that the quadrilateral $uvx''x'$ is empty and does not contain any constraints (see Figure 4.7a). We first show that u cannot be the endpoint of a constraint intersecting the interior of $uvx''x'$. We prove this by contradiction, so assume it is and let y be the other endpoint of the constraint. We first note that $\angle x'uy < \angle x'uv$. We look at $C_{1,j}^y$, the subcone of T_{yu} that lies below uy, and let z be the lowest vertex in this subcone. If u is the closest visible vertex in this subcone, uy would be an edge, which contradicts that v minimizes $\angle x'uv$. Otherwise, let z be the lowest vertex in this subcone. Since z is the lowest vertex in $C_{1,j}^y$, the visible region of T_{zu} is empty and uz is an edge. However, since $\angle x'uz < \angle x'uy < \angle x'uv$, we have a contradiction. Thus u cannot be

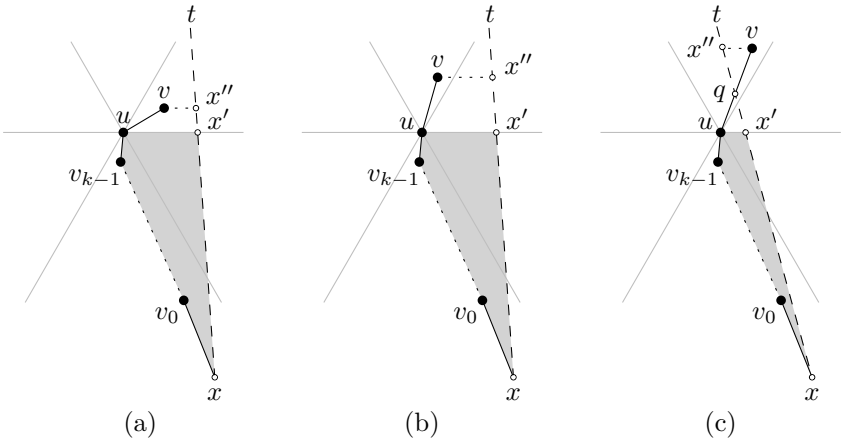

Figure 4.7: The three types of steps the algorithm can take: (a) v lies in a subcone of $\overline{C}_1^u$, (b) v lies in a subcone of C_0^u and uv does not cross st, and (c) v lies in a subcone of C_0^u and uv crosses st

the endpoint of a constraint intersecting the interior of $uvx''x'$.

Since u is not the endpoint of a constraint intersecting the interior of $uvx''x'$ and uv, ux', and $x'x''$ are visibility edges, any constraint intersecting the interior of $uvx''x'$ has at least one endpoint in $uvx''x'$. Thus it suffices to show that $uvx''x'$ is empty. We prove this by contradiction, so assume that $uvx''x'$ is not empty and let y be the lowest vertex in $uvx''x'$. Let $C_{1,j}^y$ be the subcone of C_1^y that contains u. Vertex u is visible to y, since any constraint crossing uy has an endpoint in $\overline{C}_1^u$ below y, contradicting that y is the lowest vertex, or in the region bounded by $x, v_0, ..., v_{k-1}, u, x'$ which contradicts the invariant. Hence y has an edge in $C_{1,j}^y$. This edge cannot be to u since $\angle x'uy < \angle x'uv$. Since y is the lowest vertex in $uvx''x'$, it cannot have an edge to a vertex in $uvx''x'$. Since by the invariant the region bounded by $x, v_0, ..., v_{k-1}, u, x'$ is empty, the edge of y in $C_{1,j}^y$ must cross uv. However, this contradicts the fact that the constrained half-θ_6-graph is plane.

Case (b): If v lies in a subcone of C_0^u and uv does not cross st, we need to show that the quadrilateral $uvx''x'$ is empty and does not contain any constraints (see Figure 4.7b). We first show that $uvx''x'$ is empty. We prove this by contradiction, so assume that $uvx''x'$ is not empty and let y be the lowest vertex in $uvx''x'$. We consider two cases: y lies in $\overline{C}_1^u$ and y lies in C_0^u.

If y lies in $\overline{C}_1^u$, let $C_{1,j}^y$ be the subcone of y that contains u. Vertex u is visible to y, since any constraint crossing uy has an endpoint in $\overline{C}_1^u$ below y, contradicting that y is the lowest vertex, or in the region bounded by $x, v_0, ..., v_{k-1}, u, x'$ which contradicts

the invariant. Hence y has an edge in $C^y_{1,j}$. This edge cannot be to u since $\angle x'uy$ is smaller than $\angle x'uv$. Furthermore, since y is the lowest vertex in $vvx''x'$, it cannot have an edge to a vertex in $uvx''x'$. Finally, since by the invariant the region bounded by $x, v_0, ..., v_{k-1}, u, x'$ is empty, the edge of y in $C^y_{1,j}$ must cross uv, which contradicts the fact that the constrained half-θ_6-graph is plane. In fact, this argument shows that $\overline{C}^u_1 \cap h'$ is empty.

If y lies in a subcone of C^u_0 and y is visible to u, uy would be an edge and $\angle x'uy < \angle x'uv$. So, assume that y is not visible from u. This means that there is a constraint that crosses uy. Since the line st and the edges of the region bounded by $x, v_0, ..., v_{k-1}, u, x'$ are visibility edges, the lower endpoint of this constraint must lie in $x, v_0, ..., v_{k-1}, u, v, x''$. By the invariant, it cannot lie in $x, v_0, ..., v_{k-1}, u, x'$, so it must lie in $uvx''x'$ and below y. However, this contradicts that y is the lowest vertex in $uvx''x'$. Since we arrived at a contradiction in both cases, we conclude that quadrilateral $uvx''x'$ is empty.

Next, we show that $uvx''x'$ does not contain any constraints. Since $uvx''x'$ is empty, the only way a constraint can intersect it, is when u is one of its endpoints. Hence, it remains to show that u cannot be the endpoint of a constraint intersecting the interior of $uvx''x'$. We prove this by contradiction, so assume it is and let y be the other endpoint of the constraint. Since $uvx''x'$ is empty, uy crosses vx''. Since st is a visibility edge, uy cannot cross it. Since $\overline{C}^u_1 \cap h'$ is empty, it suffices to consider the case where y lies in a subcone of C^u_0. Let $C^u_{0,j}$ be the subcone of C^u_0 to the right of uy. If y lies below t, $C^u_{0,j}$ contains a closest visible vertex whose angle with ux' is less than $\angle xuv$, contradicting that the routing algorithm routes to v.

If y lies above t, let z be the lowest vertex in the union of $C^u_{0,j}$ and $\overline{C}^u_1 \cap h'$. Since this region contains t, it is not empty and such a vertex z exists. If $z \in C^u_{0,j}$, it is the closest vertex in $C^u_{0,j}$. If $z \in \overline{C}^u_1$, u is the closest vertex to z. We also note that in both cases z is visible to u, since any constraint blocking it would have an endpoint below z. Hence, both cases result in an edge uz. However, since $\angle x'uz < \angle x'uv$, this contradicts that the routing algorithm routed to v. Therefore, u cannot be the endpoint of a constraint intersecting the interior of $uvx''x'$.

Case (c): If v lies in a subcone of C^u_0 and uv crosses st, let q be the intersection of uv and st. We need to show that the triangles uqx' and $qx''v$ are empty and do not contain any constraints (see Figure 4.7c). The proof that uqx' empty and does not contain any constraints is analogous to the previous case.

We prove that $qx''v$ is empty by contradiction, so assume that $qx''v$ is not empty. Since qx'' and qv are visibility edges, we can apply Lemma 2.1 and we obtain a vertex

y in $qx''v$ that is visible from q. If y is visible from u, v is not the closest vertex and edge uv would not exist. If y is not visible from u, we note that uq is visible and apply Lemma 2.1 on triangle uyq. This gives us a vertex z that is visible to u and closer to u than v, again contradicting the existence of edge uv. Hence, triangle $qx''v$ is empty.

Finally, we show that $qx''v$ does not contain any constraints. Since qx'' and qv are visibility edges and $qx''v$ is empty, any constraint intersecting the interior of $qx''v$ must have q as an endpoint. However, since q is not a vertex, it cannot be the endpoint of a constraint. □

Finally, we show that the path followed by the routing algorithm is 2-competitive.

Theorem 4.1 *The described routing algorithm is 2-competitive.*

Proof. Given two vertices s and t such that s and t are visible to each other, we assume without loss of generality that $t \in C_0^s$. The routing algorithm will thus only take steps in $C_0^{v_i}$, $\overline{C}_1^{v_i}$, and $\overline{C}_2^{v_i}$, where v_i is an arbitrary vertex along the routing path. Let a and b be the upper left and right corner of T_{st}. To bound the length of the routing path, we first bound the length of each edge. We consider three cases: (a) edges in subcones of $\overline{C}_1^{v_i}$ or $\overline{C}_2^{v_i}$, (b) edges in subcones of $C_0^{v_i}$ that do not cross st, (c) edges in subcones of $C_0^{v_i}$ that cross st. For ease of notation we use v_0 and v_k to denote s and t.

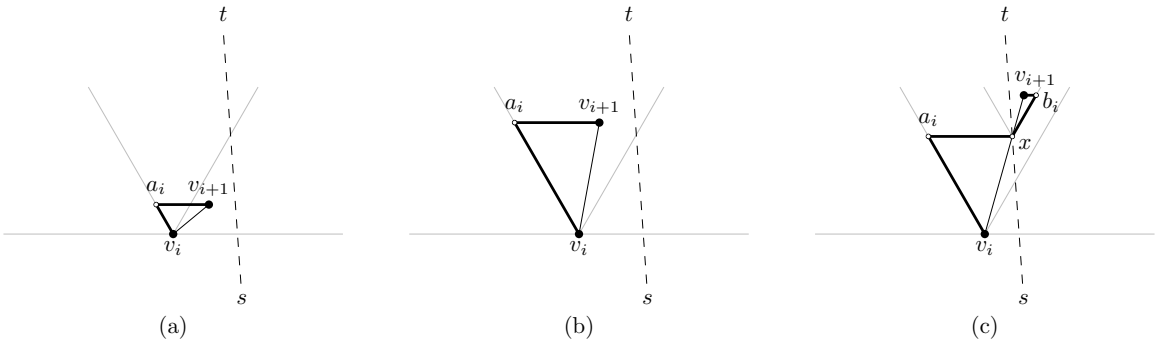

Figure 4.8: Bounding the edge lengths: (a) an edge in a subcone of $\overline{C}_1^u$, (b) an edge in a subcone of C_0^u that does not cross st, and (c) an edge in a subcone of C_0^u that crosses st

Case (a): If edge v_iv_{i+1} lies in a subcone of $\overline{C}_1^{v_i}$, let a_i be the upper left corner of $T_{v_{i+1}v_i}$ (see Figure 4.8a). By triangle inequality, we have that $|v_iv_{i+1}| \leq |v_ia_i|+|a_iv_{i+1}|$.

If edge v_iv_{i+1} lies in $\overline{C}_2^{v_i}$, let b_i be the upper right corner of $T_{v_{i+1}v_i}$. By triangle inequality, we have that $|v_iv_{i+1}| \leq |v_ib_i| + |b_iv_{i+1}|$.

Case (b): For edges v_iv_{i+1} in subcones of $C_0^{v_i}$ that do not cross st, let a_i and b_i be the upper left and right corner of $T_{v_iv_{i+1}}$ (see Figure 4.8b). If v_i lies to the left of st, we use that $|v_iv_{i+1}| \leq |v_ia_i| + |a_iv_{i+1}|$. If v_i lies to the right of st, we use that $|v_iv_{i+1}| \leq |v_ib_i| + |b_iv_{i+1}|$.

Case (c): Edges v_iv_{i+1} in subcones of $C_0^{v_i}$ that cross st are split into two parts, one for each side of st (see Figure 4.8c). Let x be the intersection of st and v_iv_{i+1}. If u lies to the left of st, let a_i be the upper left corner of T_{v_ix} and let b_i be the upper right corner of $T_{xv_{i+1}}$. By triangle inequality, we have that $|v_iv_{i+1}| \leq |v_ia_i| + |a_ix| + |xb_i| + |b_iv_{i+1}|$. If u lies to the right of st, let a_i be the upper left corner of $T_{xv_{i+1}}$ and let b_i be the upper right corner of T_{v_ix}. By triangle inequality, we have that $|v_iv_{i+1}| \leq |v_ib_i| + |b_ix| + |xa_i| + |a_iv_{i+1}|$.

Let x and x' be two consecutive crossings of the routing path and st and let v_i and v_i' be the last vertices before the path crosses st, i.e. v_iv_{i+1} crosses st at x and $v_i'v_{i+1}'$ crosses st at x'. Let a_x and b_x be the upper left and right corner of $T_{xx'}$. If the path between x and x' lies to the left of st, this part of the path is bounded by:

$$|xa_i| + \sum_{j=i}^{i'} |a_jv_{j+1}| + \sum_{j=i+1}^{i'} |v_ja_j| + |a_i'x'| = |xa_x| + |a_xx'|.$$

If the path between x and x' lies to the right of st, this part of the path is bounded by (see Figure 4.9a):

$$|xb_i| + \sum_{j=i}^{i'} |b_jv_{j+1}| + \sum_{j=i+1}^{i'} |v_jb_j| + |b_i'x'| = |xb_x| + |b_xx'|.$$

Next, we flip all unfolded bounds to the longer of the two sides at and bt: if $|at| \geq |bt|$, we replace all bounds of the form $|xb_x| + |b_xx'|$ by $|xa_x| + |a_xx'|$ and if $|at| < |bt|$, we replace all bounds of the form $|xa_x| + |a_xx'|$ by $|xb_x| + |b_xx'|$ (see Figure 4.9b). Note that this can only increase the length of the bounds. Finally, we sum these bounds and get that the sum is equal to $|sa| + \max\{|at|, |bt|\}$, which is at most $2 \cdot |st|$. □

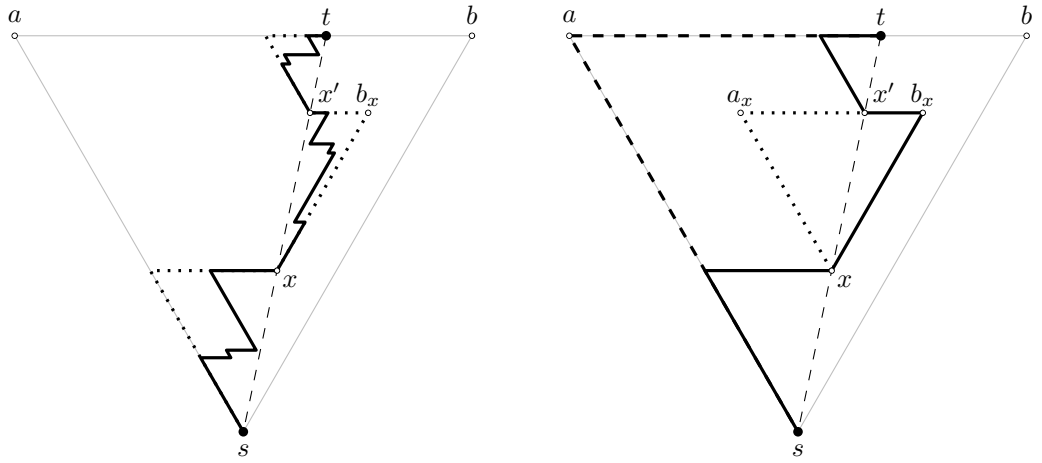

Figure 4.9: Bounding the total length: (a) the bounds (solid lines) are unfolded (dotted lines) and (b) the unfolded bounds (solid lines) are flipped to the longer of the two sides (dotted lines) and unfolded again (dashed lines)

4.2.2 Routing on the Constrained θ_6-Graph

To route on the constrained θ_6-graph, we split it into two constrained half-θ_6-graphs: the constrained half-θ_6-graph oriented as in Figure 2.5 and the constrained half-θ_6-graph where positive and negative cones are inverted. When we want to route from s to t, we pick the constrained half-θ_6-graph in which t lies in a positive subcone of s and apply the routing algorithm described in the previous section. Since this routing algorithm is 1-local and 2-competitive, we obtain a 1-local and 2-competitive routing algorithm for the constrained θ_6-graph, provided that we can determine locally, while routing, whether an edge is part of the constrained half-θ_6-graph, where t lies in a positive subcone of s. When at a vertex u, we consider the edges in order of increasing angle with the horizontal halfline through u that intersects st.

Lemma 4.4 *While executing the positive routing algorithm, we can determine locally at a vertex u for any edge uv in the constrained θ_6-graph whether it is part of the constrained half-θ_6-graph, where t lies in a positive subcone of s and s and t are visible.*

Proof. Suppose we color the edges of the constrained θ-graph red and blue such that red edges form the constrained half-θ_6-graph, where t lies in a positive subcone of s,

and blue edges form the constrained half-θ_6-graph, where t lies in a negative subcone of s. At a vertex u, we need to determine locally whether uv is red. Since an edge can be part of both constrained half-θ_6-graphs, it can be red and blue at the same time. This makes it harder to determine whether an edge is red, since determining that it is blue does not imply that it is not red.

If v lies in a positive subcone of u, we need to determine if it is the closest vertex in that subcone. Since by construction of the constrained half-θ_6-graph, u is connected to the closest vertex in this subcone, it suffices to check whether this vertex is v. Note that if uv is a constraint, v lies in two subcones of u and hence we need to check if it is the closest vertex in at least one of these subcones.

If v lies in a negative subcone of u, we know that if it is not the closest visible vertex in that subcone, it is red. Hence, it remains to determine for the edge to the closest vertex whether it is red: If it is the closest visible vertex, it is blue, but it may be red as well if u is also the closest visible vertex to v. Hence, we need to determine whether u is the closest vertex in $C^v_{i,j}$, a subcone of v that contains u. We consider two cases: (a) uv is a constraint, (b) uv is not a constraint.

Case (a): Since uv is a constraint, we know that it cannot cross st. Since we are considering uv, we also know that all edges that make a smaller angle with the horizontal halfline through u that intersects st are not red. Hence, uv is either part of the boundary of the routing path or the constraint is contained in the interior of the region bounded by the routing path and st. However, by the invariant of Lemma 4.3, the region bounded by the routing path and st does not contain any constraints in its interior. Therefore, uv is part of the boundary of the routing path and thus uv is red.

Case (b): If uv is not a constraint, let regions A and B be the intersection of C^v_i and the two subcones of u adjacent to $\overline{C}^u_i$ and let C be the intersection of $C^v_{i,j}$ and the negative subcone of u that contains v (see Figure 4.10). We first note that since uv lies in a negative subcone of u, the invariant of Lemma 4.3 implies that B is empty. Furthermore, since v is the closest visible vertex to u, C does not contain any vertices that are visible to u or v.

Since C does not contain any vertices that are visible to u or v, any constraint in $\overline{C}^u_i$ that has u as an endpoint and lies above uv, ensures that v cannot see A, i.e. it cannot block visibility of this region only partially. Hence, if such a constraint exists, u is the closest visible vertex to v in $C^v_{i,j}$, since neither B nor C contain any vertices visible to v. Therefore, uv is red.

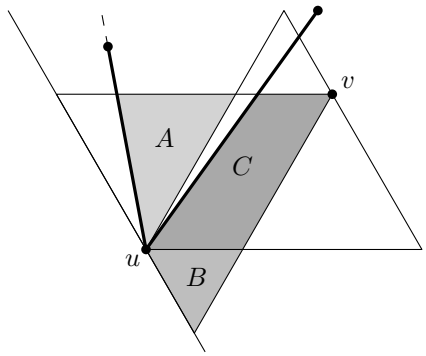

Figure 4.10: The three regions A, B, and C when determining whether an edge is part of the constrained half-θ_6-graph

If v can see A, we show that uv is red, if and only if the closest visible vertex in the subcone of u that contains A does not lie in A. We first show that if uv is red, then the closest visible vertex in the subcone of u that contains A does not lie in A. We prove the contrapositive of this statement: if the closest visible vertex x in the subcone of u that contains A lies in A, then uv is not red. Since A is visible to v, u is not the endpoint of a constraint in $\overline{C}_i^u$ above uv. Hence, we have two visibility edges uv and ux and u is not the endpoint of a constraint intersecting the interior of triangle uxv. Therefore, by Lemma 2.1, we have a convex chain of visibility vertices between x and v. Let y be the vertex adjacent to v along this chain. Since the polygon defined by ux, uv, and the convex chain is empty and does not contain any constraints, y lies in $C_{i,j}^v$. Thus, u is not the closest visible vertex in $C_{i,j}^v$ and uv is not red.

Next, we show that if the closest visible vertex x in the subcone of u that contains A does not lie in A, then uv is red. We prove this by contradiction, so assume that uv is not red. This implies that there exists a vertex $y \in C_{i,j}^v$ that is visible to v and closer than u. Since B is empty and C does not contain any vertices that are visible to v, y lies in A. Since uv and vy are visibility edges and v is not the endpoint of a constraint intersecting the interior of triangle uyv, by Lemma 2.1 there exists a convex chain of visibility edges between u and y. Furthermore, since C does not contain any vertices that are visible to u, the vertex adjacent to u along this chain lies in A. Since any vertex in A is closer to u than x, this leads to a contradiction, completing the proof. □

Routing Algorithm for the Constrained θ_6-Graph

Hence, to route on the constrained θ_6-graph, we apply the positive routing algorithm

on the constrained half-θ_6-graph, where t lies in a positive subcone of s, while determining whether an edge is part of this constrained half-θ_6-graph. The latter can be determined as follows: If v lies in a positive subcone, we need to check whether it is the closest vertex in that subcone. If v lies in a negative subcone and it is not the closest vertex, it is part of the constrained half-θ_6-graph. Finally, if v is the closest vertex in a negative subcone, it is part of the constrained half-θ_6-graph when it is a constraint or when the intersection of the cone of v that contains u and the subcone of C^u_{i-1} adjacent to $\overline{C}^u_i$ is empty.

4.2.3 Negative Routing on the Constrained Half-θ_6-Graph

In this section, we provide an $O(1)$-memory 1-local routing algorithm for the case where the destination t lies in a negative subcone of the origin s. For ease of exposition, we assume that s lies in a subcone of C^t_0. The $O(1)$-memory 1-local routing algorithm finds a path from s to t of length at most $2 \cdot |st|$ and travels a total distance of at most $18 \cdot |st|$ to do so.

We proceed to describe the edges considered by the negative routing algorithm. Given a vertex v and all neighbors of v whose projection along the bisector of C^t_0 is closer to t than the projection of v, we number the neighbors $u_0, ..., u_k$ of v the counterclockwise order, starting from the horizontal half-line to the left of v (see Figure 4.11). We create $k+1$ regions around v:

- We create $k-1$ triangular regions vu_iu_{i+1} for $0 \leq i \leq k-1$.
- We create one unbounded region using edge vu_0 and the two half-lines perpendicular to the bisector of C^t_0 starting at v and u_0 directed towards the left boundary C^t_0.
- We create one unbounded region using edge vu_k and the two half-lines perpendicular to the bisector of C^t_0 starting at v and u_k directed towards the right boundary C^t_0.

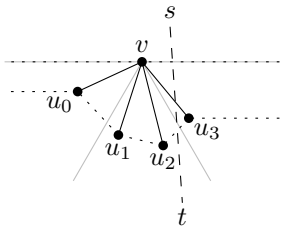

Figure 4.11: Triangle vu_2u_3 is the last region of v intersected by st

The *last region* of v intersected by st is defined as the last of these regions that is encountered when following st from s to t. In Figure 4.11, the region defined by v, u_2, and u_3 is the last region of v intersected by st.

We consider an edge of v for our routing algorithm when:

- Vertices u and v lie inside or on the boundary of T_{ts}.
- Edge uv is part of the last region of v that is intersected by st.
- Edge uv is the edge that the positive routing algorithm picks at u when routing from t to s. Note that for this condition, we do not require that u is part of the positive routing path, but only that should the positive routing path reach u, edge uv is the edge it would select for its next step.

Given s and t, the first two requirements can be checked using only 1-local information, i.e. the neighbors of the current vertex. The last requirement, on the other hand, may need 2-local information as it involves the neighbors of the neighbors of v. Hence, instead of using this last requirement, we ignore the edges that can never satisfy it and show that we can route competitively and 1-locally on the graph G formed by the edges that meet the first two requirements.

Since s lies in a subcone of C_0^t, the edges that define the last intersected region of a vertex v can lie in three cones: C_1^v, $\overline{C}_0^v$, and C_2^v. Since edges in the C_1^v and C_2^v of the negative routing algorithm correspond to edges in $\overline{C}_1^u$ and $\overline{C}_2^u$ of the positive routing algorithm, the positive routing algorithm never follows these edges if they intersect st. Hence, these edges need not be considered by the negative routing algorithm (see Figure 4.12a).

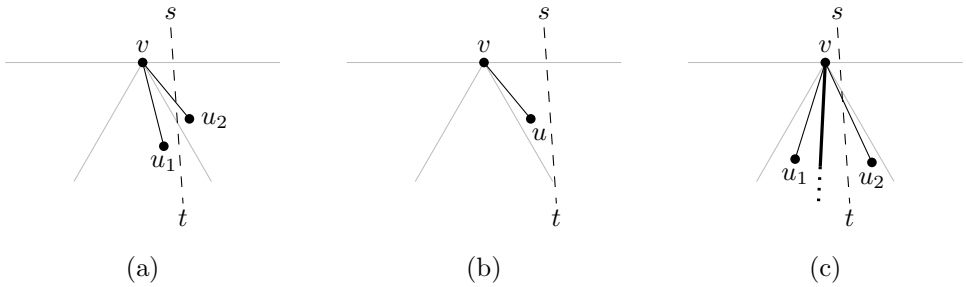

Figure 4.12: The edges ignored by the negative routing algorithm: (a) edge u_2v is ignored since it intersects st, (b) edge uv is ignored since C_2^v is intersected by st, (c) edge u_1v is ignored since it lies in a subcone that is not intersected by st and u_1vu_2 is intersected by a constraint that has v as an endpoint

We also do not need to consider edges in C_1^v and C_2^v when that cone is intersected by st (see Figure 4.12b): Assume C_1^v is intersected by st. Since we are considering edge uv, it cannot cross st. Hence, st intersects cone C_1^u, but more importantly st intersects $\overline{C}_2^u$. Hence, if the positive routing algorithm reaches u, it continues by following an edge in $\overline{C}_2^u$ or C_0^u. Since C_1^v corresponds to $\overline{C}_1^u$, no edge in this cone is followed by the positive routing algorithm, and we can ignore it.

Finally, we ignore edges that lie in a subcone that is not intersected by st when v is the endpoint of a constraint that intersects the interior of the last region of v that is intersected by st (see Figure 4.12c): If v is the endpoint of a constraint that intersects the interior of the last region of v that is intersected by st, we do not consider the edge that is not intersected by st. We can ignore this edge, since by the invariant, the region between the routing path and st does not contain any constraints.

Since these conditions can be checked using only s, t, v, the neighbors of v, and the constraints incident to v, we can determine 1-locally whether to consider an edge. Hence, the graph G on which we route is the graph formed by all edges uv such that:

- Vertices u and v lie inside or on the boundary of T_{ts}.
- Edge uv is part of the last region of v that is intersected by st.
- Edge uv does not meet any of the following three conditions:
  - Edge uv lies in C_1^v or C_2^v and crosses st.
  - Edge uv lies in C_1^v or C_2^v and this cone is intersected by st.
  - Edge uv lies in a subcone that is not intersected by st and v is the endpoint of a constraint that intersects the interior of the last region of v that is intersected by st.

For ease of exposition, we consider each edge of G to be oriented upward: Let u' and v' be the projections of u and v along the bisector of C_0^t. Edge uv is oriented from u to v if and only if $|tu'| \leq |tv'|$. Note that this does not imply that u lies in a negative cone of v. We proceed to prove that every vertex with two incoming edges is part of the positive routing path when routing from t to s.

Lemma 4.5 *Every vertex with in-degree 2 that is reached by the negative routing algorithm is part of the positive routing path from t to s.*

Proof. Let v be a vertex of in-degree 2 that is reached by the negative routing algorithm. Let u and w be the other endpoints of these edges to v, such that the projection of u along the bisector of T_{ts} is closer to t than the projection of w (see Figure 4.13). Since both uv and wv are part of the last intersected region of v, vertices u and w must lie on opposite sides of st. This implies that the positive routing algorithm reaches at least one of them when routing from t to s, since by the invariant the region between the routing path and st is empty. Thus it suffices to show that from both u and w the positive routing algorithm eventually reaches v.

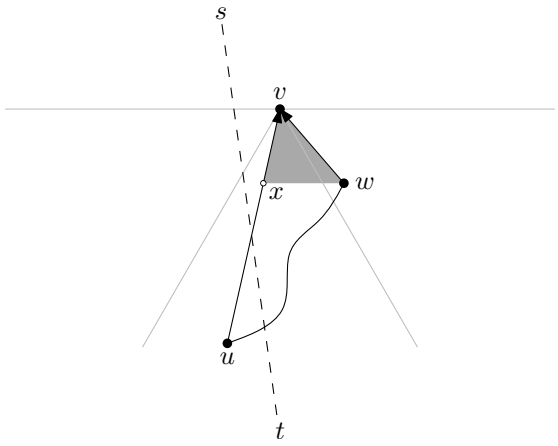

Figure 4.13: Vertex v has in-degree 2

If the positive routing algorithm reaches w, we show that it would follow the edge to v. Let x be the intersection of uv and the horizontal line through w (see Figure 4.13). First, we show that triangle vwx is empty. If w lies in a subcone of C_1^v or C_2^v, u lies in a subcone of $\overline{C}_0^v$, since otherwise one of the two edges would cross st and be ignored. Since vw and vx are visibility edges and v is not the endpoint of a constraint intersecting the interior of vwx, it follows from Lemma 2.1 that if vwx is not empty, there exists a convex chain of visibility edges between w and x and the region bounded by this chain, vw, and vx is empty. Let y be a vertex along this convex chain and note that y is visible to v. If y lies in the same cone of v as w, it also lies in the same subcone of v as w, since v is not the endpoint of a constraint intersecting the interior of vwx. However, this implies that w is not the closest visible vertex to v in this subcone, contradicting that vw is an edge. If y lies in $\overline{C}_0^v$, y has

an edge in its subcone that contains v, since v is a visible vertex in that subcone. This edge cannot cross vw and vu, since the constrained half-θ_6-graph is plane, and it cannot be connected to a vertex in the region bounded by the convex chain, vw, and vx, since it is empty. Hence, y would have an edge to v, contradicting that vu and vw are consecutive edges around v. We conclude that triangle vwx is empty.

Using an analogous argument, it can be shown that if u lies in a subcone of C_1^v or C_2^v, w lies in $\overline{C}_0^v$ and the existence of a vertex in vwx would contradict that uv is an edge or that u and w are consecutive edges around v. If both u and w lie in a subcone of $\overline{C}_0^v$, the argument reduces to the case where y lies in $\overline{C}_0^v$, again contradicting that u and w are consecutive edges around v. Hence, since vwx is empty, the positive routing algorithm routes to v when it reaches w, since it minimizes angle $\angle xwv$.

Next, we look at the case where the positive routing path reaches u. If it follows edge uv, we are done. If it does not follow edge uv, let z be the other endpoint of the edge the positive routing algorithm follows at u. By construction of the positive routing path, we know that the projection of z on the bisector of T_{ts} lies further from t than the projection of u. Since the constrained half-θ_6-graph is plane, the path from z to s cannot cross uv or wv, and since the positive routing path is monotone with respect to the bisector of T_{ts}, it cannot go down and around or through u. Furthermore, since the region enclosed by the positive routing path and st is empty, the path also cannot go around w without passing through w. Finally, since uv and wv are consecutive edges around v, the path from z to s cannot reach v by arriving from an edge between uv and wv. Hence, w must lie on the path from z to s. Thus, since we previously showed that when the positive routing algorithm reaches w, it routes to v, vertex v is also reached when the positive routing path reaches u. □

Negative Routing Algorithm for the Constrained Half-θ_6-Graph

Routing from s to t now comes down to searching for a vertex that has in-degree 2 on one of the two paths leaving s. When such a vertex v is found, we need to find the next vertex has in-degree 2 on one of the two paths leaving v. This process is repeated until we reach t. A single instance of this problem, i.e. finding the next vertex has in-degree 2 from another vertex can be viewed as searching for a specific point on a line. This problem has been studied extensively and a search strategy that is 9-competitive was presented by Baeza-Yates *et al.* [3]: We start by following the shorter of the two edges connected to s and call this distance 1. If we reached a vertex with in-degree 2, we are done. Otherwise, we go back to s and follow the other path up to distance 2 from s. Again, if we reached a vertex with in-degree 2, we

are done. Otherwise, we go back to s and follow the first path up to distance 4 from s. This process of backtracking and doubling the allowed travel distance is repeated until a vertex with in-degree 2 is reached. Since this strategy needs to keep track of the distance traveled, it uses $O(1)$-memory. Hence, we apply this search strategy and perform the following actions when we reach an unvisited vertex v:

- If v has in-degree 2, v is part of the positive routing path and we restart the searching strategy from v.
- If v has in-degree 1, we proceed to its neighbor u if we have enough budget left to traverse the edge. At u we check whether the positive routing algorithm would follow edge uv. If this is not the case, we know that v was a dead end and the path on the opposite side of st is part of the positive routing path. Hence, we backtrack and follow the path on the opposite side of st to the last visited vertex on that side.
- If v has in-degree 0, it is a dead end and we backtrack like in the previous case.

We conclude this section by showing that the above $O(1)$-memory 1-local routing algorithm has a routing ratio of at most 9 times the length of the positive routing path, which implies an 18-competitive 1-local routing algorithm for negative routing in the constrained half-θ_6-graph.

Theorem 4.2 *The above routing algorithm is an $O(1)$-memory 1-local 18-competitive routing algorithm for negative routing in the constrained half-θ_6-graph.*

Proof. Let p be the last vertex where the search strategy was restarted, initially p is s. We prove the theorem by showing that when we restart the search strategy at vertex q, we traveled at most 9 times the distance along the positive routing path between p and q. If we restart the search strategy because we reached a vertex of in-degree 2, this follows directly from the fact that the search strategy is 9-competitive, i.e. we found the vertex we are looking for and we spent at most 9 times the distance along the positive routing path between p and q.

If we reach a vertex v with in-degree 0 or we traverse an edge vu and the positive routing algorithm would not have routed from u to v, we backtrack to p and traverse the path on the opposite side of st. We follow this path until we reach w, the last vertex traversed on this side of st. Unfortunately, w is too close to p to prove that the total length traveled is at most 9 times the distance along the positive routing path between p and w. However, w must have in-degree 1: Since w is part of the

positive routing path, it cannot have in-degree 0, and since we did not restart the search strategy when we reached w the previous time, it cannot have in-degree 2. Hence, it has in-degree 1 and it follows that the vertex q to which w is connected is also part of the positive routing path. Since the distance along the positive routing path between p and v is at most 2 times the distance along the positive routing path between p and q, an argument analogous to the one used by Baeza-Yates *et al.* [3] shows that we traversed at most 9 times the distance along the positive routing path between p and q to reach q. □

4.2.4 Lower Bound on the Negative Routing Algorithm

In this section we show that the negative routing algorithm described in the previous section cannot be guaranteed to reach t while traveling less than $2\sqrt{39} \cdot |st| \approx 12.48 \cdot |st|$. This situation is shown in Figure 4.14: We place a vertex r_1 almost horizontally to the right of s at distance 1, followed by a vertex l_1 almost horizontally to the left of s at distance 2, followed a vertex r_2 almost horizontally to the right of s at distance 4. Once we reach the corners of T_{ts} at l_2 and r_3, we proceed down along the boundary of T_{ts} and place vertices l_3 and r_4 such that the distance between s and l_3 via l_2 is 8 and the distance between s and r_4 via r_3 is 16. Finally, we place vertices l_4 and r_5 arbitrarily close to t. The positive routing path from t to s would route to r_5, r_4, r_3, r_2, r_1, and finally s.

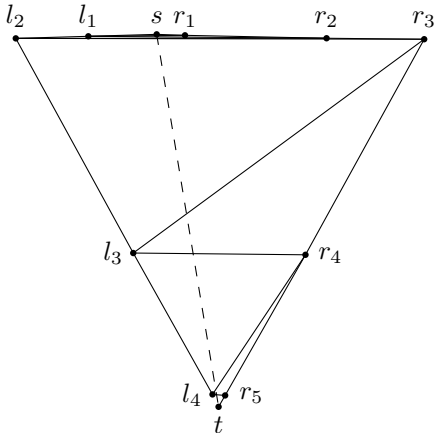

Figure 4.14: The situation where the negative routing algorithm uses $2\sqrt{39} \cdot |st|$ to reach t

The negative routing algorithm on the other hand would try both sides, going back through s each time it switches sides: go to r_1, go to l_1, go to r_2, go to l_3 (via

l_2), go to r_5 (via r_3), go to l_4 (via l_2), and finally go to t (via s and r_5). We can pick the edge lengths between the vertices in such a way that each time the next vertex along one of the two sides is reached (other than l_4), the negative routing algorithm runs out of budget and needs to backtrack to try the other side. The total length traveled this way is the sum of:

- $2 \cdot \delta(s, r_5)$, for going back and forth from s until the step before r_5 is reached for the first time,
- $2 \cdot \delta(s, r_5)$, for going to r_5 and back to s when the negative routing algorithm almost reaches t,
- $2 \cdot \delta(s, l_4)$, for going down the wrong path (and back up) after reaching r_5,
- $\delta(s, t)$, for finally reaching t,

where $\delta(x, y)$ is the distance along the negative routing path between x and y. Since r_5 can be arbitrarily close to t, this sums up to $5 \cdot \delta(s, t) + 2 \cdot \delta(s, l_4)$.

Let α be the angle between the bisector of T_{ts} and ts. Using the law of sinex, we can express $\delta(s, t)$ and $\delta(s, l_4)$ as follows:

$$\begin{aligned}
\delta(s,t) &= |sr_3| + |r_3t| \\
&= \left(\frac{\sin\left(\frac{\pi}{6} + \alpha\right)}{\sin\left(\frac{\pi}{3}\right)} + \frac{\sin\left(\frac{\pi}{2} - \alpha\right)}{\sin\left(\frac{\pi}{3}\right)} \right) \cdot |st| \\
&= \left(\sqrt{3} \cdot \cos\alpha + \sin\alpha \right) \cdot |st|
\end{aligned}$$

$$\begin{aligned}
\delta(s,l_4) &= |sl_2| + |l_2l_4| \\
&= \left(\frac{\sin\left(\frac{\pi}{6} - \alpha\right)}{\sin\left(\frac{\pi}{3}\right)} + \frac{\sin\left(\frac{\pi}{2} - \alpha\right)}{\sin\left(\frac{\pi}{3}\right)} \right) \cdot |st| \\
&= \left(\sqrt{3} \cdot \cos\alpha - \sin\alpha \right) \cdot |st|
\end{aligned}$$

Thus, the total distance traveled by the negative routing algorithm becomes:

$$\begin{aligned}
& 5 \cdot \delta(s,t) + 2 \cdot \delta(s,l_4) \\
= \; & 5 \cdot \left(\sqrt{3} \cdot \cos\alpha + \sin\alpha \right) \cdot |st| + 2 \cdot \left(\sqrt{3} \cdot \cos\alpha - \sin\alpha \right) \cdot |st| \\
= \; & \left(7\sqrt{3} \cdot \cos\alpha + 3\sin\alpha \right) \cdot |st|
\end{aligned}$$

When maximizing this function over α, with $0 \leq \alpha \leq \pi/6$, we find the maximum at $\alpha \approx 0.2425$, where the function has value $2\sqrt{39} \cdot |st| \approx 12.48 \cdot |st|$.

4.3 Conclusion

In this chapter, we showed that no deterministic 1-local routing algorithm is $o(\sqrt{n})$-competitive on all pairs of non-visible vertices of the constrained θ_6-graph, regardless of the amount of memory it is allowed to use. Following this negative result, we showed how to route between any two visible vertices of the constrained θ_6-graph using only 1-local information by routing on one of the two constrained half-θ_6-graphs. This routing algorithm guarantees that the returned path has length at most 2 times the Euclidean distance between the source and destination. To the best of our knowledge, this is the first 1-local routing algorithm in the constrained setting with guarantees on the length of the returned path.

There remain a number of open problems in the area of local competitive routing in the constrained setting. For example, though we showed that no deterministic 1-local routing algorithm is $o(\sqrt{n})$-competitive on all pairs of non-visible vertices of the θ_6-graph, it would still be interesting to construct a routing algorithm that reaches any non-visible vertex.

Furthermore, we showed how to route on a specific constrained θ-graph. It would be a very nice if there exists a routing local algorithm that is competitive on all constrained θ-graphs. In the unconstrained setting, the θ-routing algorithm is such an algorithm, provided that at least 7 cones are being used. In the constrained setting, however, this particular algorithm need not reach the destination, since even if the source can see the destination, this does not necessarily hold for every vertex along the path. Because of this, there need not be any edge in the cone that contains the destination, meaning that the θ-routing algorithm can get stuck.

Finally, constrained θ-graphs are not the only graphs that are known to be spanners in the constrained setting. For example, constrained Yao-graphs and constrained (generalized) Delaunay graphs have also been shown to be spanners. Like in the case of constrained θ-graphs, currently no routing algorithms are known for to exist for those graphs.

Chapter 5

Theta-Graphs

In this chapter, we present improved upper and lower bounds on the spanning ratio of θ-graphs with at least six cones. We show that for any integer $k \geq 1$, θ-graphs with $4k + 2$ cones have a spanning ratio of $1 + 2\sin(\theta/2)$ and we provide a matching lower bound, showing that this spanning ratio is tight.

Next, we show that for any integer $k \geq 1$, θ-graphs with $4k+4$ cones have spanning ratio at most $1 + 2\sin(\theta/2)/(\cos(\theta/2) - \sin(\theta/2))$. We also show that θ-graphs with $4k + 3$ and $4k + 5$ cones have spanning ratio at most $\cos(\theta/4)/(\cos(\theta/2) - \sin(3\theta/4))$. This is a significant improvement on all families of θ-graphs for which exact bounds are not known. For example, the spanning ratio of the θ-graph with 7 cones is decreased from at most 7.5625 to at most 3.5132. These spanning proofs also imply improved upper bounds on the competitiveness of the θ-routing algorithm. In particular, we show that the θ-routing algorithm is $(1+2\sin(\theta/2)/(\cos(\theta/2)-\sin(\theta/2)))$-competitive on θ-graphs with $4k + 4$ cones and that this ratio is tight.

Finally, we present improved lower bounds on the spanning ratio of these graphs. Using these bounds, we provide a partial order on these families of θ-graphs. In particular, we show that θ-graphs with $4k + 4$ cones have spanning ratio at least $1 + 2\tan(\theta/2) + 2\tan^2(\theta/2)$, where θ is $2\pi/(4k + 4)$. This is somewhat surprising since, for equal values of k, the spanning ratio of θ-graphs with $4k+4$ cones is greater than that of θ-graphs with $4k + 2$ cones, showing that increasing the number of cones can make the spanning ratio worse.

The results presented in this chapter are joint work with Prosenjit Bose, Jean-Lou De Carufel, Pat Morin and Sander Verdonschot and appear in the Proceedings of the 24th Canadian Conference on Computational Geometry (CCCG 2012) [15] and the Proceedings of the 13th Workshop on Algorithms and Data Structures (WADS

2013) [32], and have been submitted to Theoretical Computer Science (TCS) [16].

5.1 Upper Bounds

In this section, we provide improved upper bounds for the four families of θ-graphs: the $\theta_{(4k+2)}$-graph, the $\theta_{(4k+3)}$-graph, the $\theta_{(4k+4)}$-graph, and the $\theta_{(4k+5)}$-graph. We first prove that the $\theta_{(4k+2)}$-graph has a tight spanning ratio of $1 + 2\sin(\theta/2)$. Next, we provide a generic framework for the spanning proof for the three other families of θ-graphs. After providing this framework, we complete the proofs for the individual families.

5.1.1 Optimal Bounds on the $\theta_{(4k+2)}$-Graph

We start by showing that the $\theta_{(4k+2)}$-graph has a spanning ratio of $1 + 2\sin(\theta/2)$. At the end of this section, we also provide a matching lower bound, proving that this spanning ratio is tight.

Theorem 5.1 *Let u and w be two vertices in the plane. Let m be the midpoint of the side of T_{uw} opposite u and let α be the unsigned angle between uw and um. There exists a path connecting u and w in the $\theta_{(4k+2)}$-graph ($k \geq 1$) of length at most*

$$\left(\left(\frac{1+\sin\left(\frac{\theta}{2}\right)}{\cos\left(\frac{\theta}{2}\right)}\right)\cdot\cos\alpha+\sin\alpha\right)\cdot|uw|.$$

Proof. We assume without loss of generality that $w \in C_0^u$. We prove the theorem by induction on the area of T_{uw} (formally, induction on the rank, when ordered by area, of the canonical triangles for all pairs of vertices). Let a and b be the upper left and right corners of T_{uw} and let y and z be the left and right intersections of the left and right boundaries of T_{uw} and the boundaries of C_{2k+1}^w, the cone of w that contains u (see Figure 5.1). Our inductive hypothesis is the following, where $\delta(u, w)$ denotes the length of the shortest path from u to w in the $\theta_{(4k+2)}$-graph:

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

Note that if both ayw and bzw are empty, the induction hypothesis implies that $\delta(u,w) \leq \min\{|ua| + |aw|, |ub| + |bw|\}$.

We first show that this induction hypothesis implies the theorem. Basic trigonometry gives us the following equalities: $|um| = |uw| \cdot \cos\alpha$, $|mw| = |uw| \cdot \sin\alpha$, $|am| = |bm| = |uw| \cdot \cos\alpha \tan(\theta/2)$, and $|ua| = |ub| = |uw| \cdot \cos\alpha / \cos(\theta/2)$. Thus, the induction hypothesis gives us that

$$\delta(u,w) \leq |ua| + |am| + |mw| = \left(\left(\frac{1 + \sin\left(\frac{\theta}{2}\right)}{\cos\left(\frac{\theta}{2}\right)} \right) \cdot \cos\alpha + \sin\alpha \right) \cdot |uw|.$$

Base case: T_{uw} has rank 1. Since the triangle is a smallest triangle, w is the closest vertex to u in that cone. Hence, the edge uw is part of the $\theta_{(4k+2)}$-graph and $\delta(u,w) = |uw|$. From the triangle inequality, we have $|uw| \leq \min\{|ua| + |aw|, |ub| + |bw|\}$, so the induction hypothesis holds.

Induction step: We assume that the induction hypothesis holds for all pairs of vertices with canonical triangles of rank up to j. Let T_{uw} be a canonical triangle of rank $j+1$.

If uw is an edge in the $\theta_{(4k+2)}$-graph, the induction hypothesis follows from the same argument as in the base case. If there is no edge between u and w, let v be the vertex closest to u in C_0^u, and let a' and b' be the upper left and right corners of T_{uv} (see Figure 5.1). By definition, $\delta(u,w) \leq |uv| + \delta(v,w)$, and by the triangle inequality, $|uv| \leq \min\{|ua'| + |a'v|, |ub'| + |b'v|\}$.

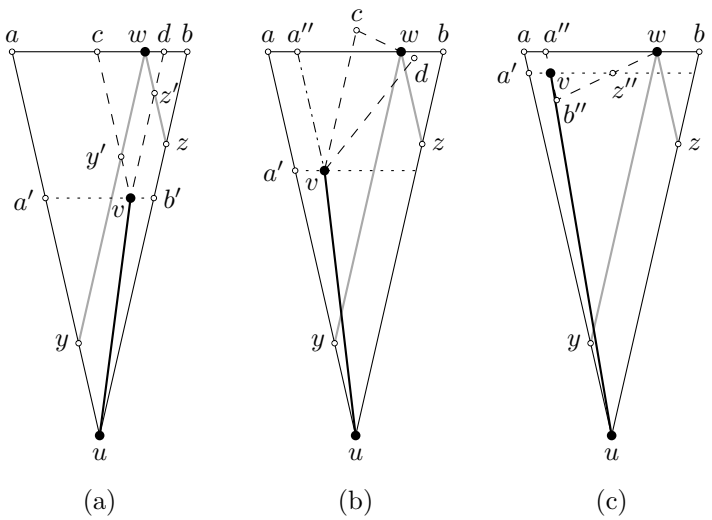

Figure 5.1: The three cases of the induction step based on the cone of v that contains w, in this case for the θ_{14}-graph

Without loss of generality, we assume that v lies to the left of w. We perform a case analysis based on the cone of v that contains w: (a) $w \in C_0^v$, (b) $w \in C_i^v$ where $1 \leq i \leq k-1$, (c) $w \in C_k^v$.

Case (a): Vertex w lies in C_0^v (see Figure 5.1a). Let c and d be the upper left and right corners of T_{vw}, and let y' and z' be the left and right intersections of T_{vw} and the boundaries of C_{2k+1}^w. Since T_{vw} has smaller area than T_{uw}, we apply the inductive hypothesis to T_{vw}. We need to prove all three statements of the inductive hypothesis for T_{uw}.

1. If ayw is empty, then $cy'w$ is also empty, so by induction $\delta(v,w) \leq |vd| + |dw|$. Since v, d, b, and b' form a parallelogram, we have:

$$\begin{aligned} \delta(u,w) &\leq |uv| + \delta(v,w) \\ &\leq |ub'| + |b'v| + |vd| + |dw| \\ &= |ub| + |bw|, \end{aligned}$$

   which proves the first statement of the induction hypothesis.

2. If bzw is empty, an analogous argument proves the second statement of the induction hypothesis.

3. If neither ayw nor bzw is empty, by induction we have $\delta(v,w) \leq \max\{|vc| + |cw|, |vd| + |dw|\}$. Assume, without loss of generality, that the maximum of the right hand side is attained by its second argument $|vd| + |dw|$ (the other case is similar). Since vertices v, d, b, and b' form a parallelogram, we have that:

$$\begin{aligned} \delta(u,w) &\leq |uv| + \delta(v,w) \\ &\leq |ub'| + |b'v| + |vd| + |dw| \\ &\leq |ub| + |bw| \\ &\leq \max\{|ua| + |aw|, |ub| + |bw|\}, \end{aligned}$$

   which proves the third statement of the induction hypothesis.

Case (b): Vertex w lies in C_i^v where $1 \leq i \leq k-1$ (see Figure 5.1b). In this case, v lies in ayw. Therefore, the first statement of the induction hypothesis for T_{uw} is vacuously true. It remains to prove the second and third statement of the induction hypothesis. Let a'' be the intersection of the side of T_{uw} opposite u and the left boundary of C_0^v. Since T_{vw} is smaller than T_{uw}, by induction we have

$\delta(v,w) \leq \max\{|vc|+|cw|, |vd|+|dw|\}$. Since $w \in C_i^v$ where $1 \leq i \leq k-1$, we can apply Lemma 2.4. Note that point a in Lemma 2.4 corresponds to point a'' in this proof. Hence, we get that $\max\{|vc|+|cw|, |vd|+|dw|\} \leq |va''|+|a''w|$. Since $|uv| \leq |ua'|+|a'v|$ and v, a'', a, and a' form a parallelogram, we have that $\delta(u,w) \leq |ua|+|aw|$, proving the induction hypothesis for T_{uw}.

Case (c): Vertex w lies in C_k^v (see Figure 5.1c). Since v lies in ayw, the first statement of the induction hypothesis for T_{uw} is vacuously true. It remains to prove the second and third statement of the induction hypothesis. Let a'' and b'' be the upper and lower left corners of T_{wv}, and let z'' be the intersection of T_{wv} and the lower boundary of C_k^v, i.e. the cone of v that contains w. Note that z'' is also the right intersection of T_{uv} and T_{wv}. Since v is the closest vertex to u, T_{uv} is empty. Hence, $b''z''v$ is empty. Since T_{wv} is smaller than T_{uw}, we can apply induction on it. As $b''z''v$ is empty, the induction hypothesis for T_{wv} gives $\delta(v,w) \leq |va''|+|a''w|$. Since $|uv| \leq |ua'|+|a'v|$ and v, a'', a, and a' form a parallelogram, we have that $\delta(u,w) \leq |ua|+|aw|$, proving the second and third statement of the induction hypothesis for T_{uw}. □

Since $((1+\sin(\theta/2))/\cos(\theta/2)) \cdot \cos\alpha + \sin\alpha$ is increasing for $\alpha \in [0, \theta/2]$, for $\theta \leq \pi/3$, it is maximized when $\alpha = \theta/2$, and we obtain the following corollary:

Corollary 5.1 *The $\theta_{(4k+2)}$-graph ($k \geq 1$) is a $(1+2\sin(\theta/2))$-spanner.*

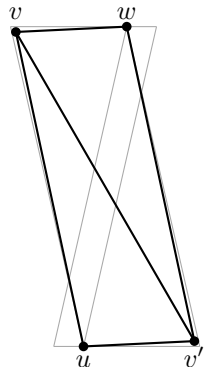

Figure 5.2: The lower bound for the $\theta_{(4k+2)}$-graph

The upper bounds given in Theorem 5.1 and Corollary 5.1 are tight, as shown in Figure 5.2: we place a vertex v arbitrarily close to the upper corner of T_{uw} that is furthest from w. Likewise, we place a vertex v' arbitrarily close to the lower corner of T_{wu} that is furthest from u. Both shortest paths between u and w visit either v or v',

so the path length is arbitrarily close to $(((1+\sin(\theta/2))/\cos(\theta/2))\cdot\cos\alpha+\sin\alpha)\cdot|uw|$, showing that the upper bounds are tight.

5.1.2 Generic Framework for the Spanning Proof

In this section, we provide a generic framework for the spanning proof for the three other families of θ-graphs: the $\theta_{(4k+3)}$-graph, the $\theta_{(4k+4)}$-graph, and the $\theta_{(4k+5)}$-graph. This framework contains those parts of the spanning proof that are identical for all three families. In the subsequent sections, we handle the single case that depends on each specific family and determines their respective spanning ratios.

Theorem 5.2 *Let u and w be two vertices in the plane. Let m be the midpoint of the side of T_{uw} opposite u and let α be the unsigned angle between uw and um. There exists a path connecting u and w in the $\theta_{(4k+x)}$-graph ($k \geq 1$) of length at most*

$$\left(\frac{\cos\alpha}{\cos\left(\frac{\theta}{2}\right)} + \boldsymbol{c}\cdot\left(\cos\alpha\tan\left(\frac{\theta}{2}\right)+\sin\alpha\right)\right)\cdot|uw|,$$

where $\boldsymbol{c} \geq 1$ is a function that depends on $x \in \{3,4,5\}$ and θ. For the $\theta_{(4k+4)}$-graph, $\boldsymbol{c}$ equals $1/(\cos(\theta/2)-\sin(\theta/2))$ and for the $\theta_{(4k+3)}$-graph and $\theta_{(4k+5)}$-graph, $\boldsymbol{c}$ equals $\cos(\theta/4)/\ (\cos(\theta/2)-\sin(3\theta/4))$.

Proof. We assume without loss of generality that $w \in C_0^u$. We prove the theorem by induction on the area of T_{uw} (formally, induction on the rank, when ordered by area, of the canonical triangles for all pairs of vertices). Let a and b be the upper left and right corners of T_{uw}. Our inductive hypothesis is the following, where $\delta(u,w)$ denotes the length of the shortest path from u to w in the $\theta_{(4k+x)}$-graph: $\delta(u,w) \leq \max\{|ua|+\boldsymbol{c}\cdot|aw|, |ub|+\boldsymbol{c}\cdot|bw|\}$.

We first show that this induction hypothesis implies the theorem. Basic trigonometry gives us the following equalities: $|um| = |uw|\cdot\cos\alpha$, $|mw| = |uw|\cdot\sin\alpha$, $|am| = |bm| = |uw|\cdot\cos\alpha\tan(\theta/2)$, and $|ua| = |ub| = |uw|\cdot\cos\alpha/\cos(\theta/2)$. Thus the induction hypothesis gives that

$$\delta(u,w) \leq |ua| + \boldsymbol{c}\cdot(|am|+|mw|) = \left(\frac{\cos\alpha}{\cos\left(\frac{\theta}{2}\right)} + \boldsymbol{c}\cdot\left(\cos\alpha\tan\left(\frac{\theta}{2}\right)+\sin\alpha\right)\right)\cdot|uw|.$$

Base case: T_{uw} has rank 1. Since the triangle is a smallest triangle, w is the closest vertex to u in that cone. Hence, the edge uw is part of the $\theta_{(4k+x)}$-graph

and $\delta(u, w) = |uw|$. From the triangle inequality and the fact that $\boldsymbol{c} \geq 1$, we have $|uw| \leq \max\{|ua| + \boldsymbol{c} \cdot |aw|, |ub| + \boldsymbol{c} \cdot |bw|\}$, so the induction hypothesis holds.

Induction step: We assume that the induction hypothesis holds for all pairs of vertices with canonical triangles of rank up to j. Let T_{uw} be a canonical triangle of rank $j + 1$.

If uw is an edge in the $\theta_{(4k+x)}$-graph, the induction hypothesis follows from the same argument as in the base case. If there is no edge between u and w, let v be the vertex closest to u in T_{uw}, and let a' and b' be the upper left and right corners of T_{uv} (see Figure 5.3). By definition, $\delta(u, w) \leq |uv| + \delta(v, w)$, and by the triangle inequality, $|uv| \leq \min\{|ua'| + |a'v|, |ub'| + |b'v|\}$.

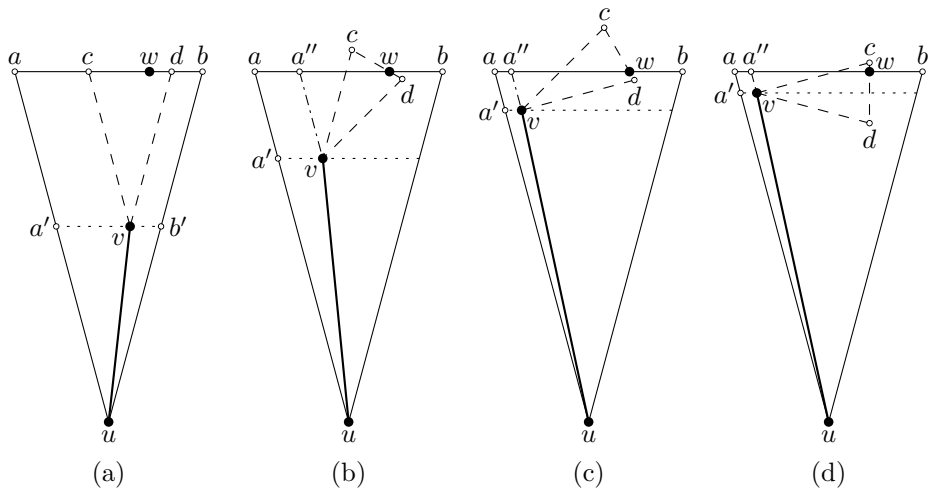

Figure 5.3: The four cases of the induction step based on the cone of v that contains w, in this case for the θ_{12}-graph

Without loss of generality, we assume that v lies to the left of w. We perform a case analysis based on the cone of v that contains w, where c and d are the left and right corners of T_{vw}, opposite to v: (a) $w \in C_0^v$, (b) $w \in C_i^v$ where $1 \leq i \leq k-1$, or $i = k$ and $|cw| \leq |dw|$, (c) $w \in C_k^v$ and $|cw| > |dw|$, (d) $w \in C_{k+1}^v$.

Case (a): Vertex w lies in C_0^v (see Figure 5.3a). Since T_{vw} has smaller area than T_{uw}, we apply the inductive hypothesis to T_{vw}. Hence we have $\delta(v, w) \leq \max\{|vc| + \boldsymbol{c} \cdot |cw|, |vd| + \boldsymbol{c} \cdot |dw|\}$. Since v lies to the left of w, the maximum of the right hand side is attained by its first argument, $|vc| + \boldsymbol{c} \cdot |cw|$. Since vertices v, c, a, and a' form

a parallelogram, and $\boldsymbol{c} \geq 1$, we have that

$$\begin{aligned}\delta(u,w) &\leq |uv| + \delta(v,w)\\ &\leq |ua'| + |a'v| + |vc| + \boldsymbol{c}\cdot|cw|\\ &\leq |ua| + \boldsymbol{c}\cdot|aw|\\ &\leq \max\{|ua| + \boldsymbol{c}\cdot|aw|, |ub| + \boldsymbol{c}\cdot|bw|\},\end{aligned}$$

which proves the induction hypothesis.

Case (b): Vertex w lies in C_i^v, where $1 \leq i \leq k-1$, or $i = k$ and $|cw| \leq |dw|$ (see Figure 5.3b). Let a'' be the intersection of the side of T_{uw} opposite u and the left boundary of C_0^v. Since T_{vw} is smaller than T_{uw}, by induction we have $\delta(v,w) \leq \max\{|vc| + \boldsymbol{c}\cdot|cw|, |vd| + \boldsymbol{c}\cdot|dw|\}$. Since $w \in C_i^v$ where $1 \leq i \leq k-1$, or $i = k$ and $|cw| \leq |dw|$, we can apply Lemma 2.4. Note that point a in Lemma 2.4 corresponds to point a'' in this proof. Hence, we get that $\max\{|vc| + |cw|, |vd| + |dw|\} \leq |va''| + |a''w|$ and $\max\{|cw|, |dw|\} \leq |a''w|$. Since $\boldsymbol{c} \geq 1$, this implies that $\max\{|vc| + \boldsymbol{c}\cdot|cw|, |vd| + \boldsymbol{c}\cdot|dw|\} \leq |va''| + \boldsymbol{c}\cdot|a''w|$. Since $|uv| \leq |ua'| + |a'v|$ and v, a'', a, and a' form a parallelogram, we have that $\delta(u,w) \leq |ua| + \boldsymbol{c}\cdot|aw|$, proving the induction hypothesis for T_{uw}.

Case (c) and (d) Vertex w lies in C_k^v and $|cw| > |dw|$, or w lies in C_{k+1}^v (see Figures 5.3c and d). Let a'' be the intersection of the side of T_{uw} opposite u and the left boundary of C_0^v. Since T_{vw} is smaller than T_{uw}, we can apply induction on it. The actual application of the induction hypothesis varies for the three families of θ-graphs and, using Lemma 2.5, determines the value of $\boldsymbol{c}$. Hence, these cases are discussed in the spanning proofs of the three families. □

5.1.3 Upper Bound on the $\theta_{(4k+4)}$-Graph

In this section, we improve the upper bounds on the spanning ratio of the $\theta_{(4k+4)}$-graph, for any integer $k \geq 1$.

Theorem 5.3 *Let u and w be two vertices in the plane. Let m be the midpoint of the side of T_{uw} opposite u and let α be the unsigned angle between uw and um. There exists a path connecting u and w in the $\theta_{(4k+4)}$-graph ($k \geq 1$) of length at most*

$$\left(\frac{\cos\alpha}{\cos\left(\frac{\theta}{2}\right)} + \frac{\cos\alpha\tan\left(\frac{\theta}{2}\right) + \sin\alpha}{\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{\theta}{2}\right)}\right)\cdot|uw|.$$

Proof. We apply Theorem 5.2 using $\boldsymbol{c} = 1/(\cos(\theta/2) - \sin(\theta/2))$. It remains to handle Case (c), where $w \in C_k^v$ and $|cw| > |dw|$, and Case (d), where $w \in C_{k+1}^v$.

Recall that c and d are the left and right corners of T_{vw}, opposite to v, and a'' is the intersection of the side of T_{uw} opposite u and the left boundary of C_0^v. Let β be $\angle a''wv$ and let γ be the angle between vw and the bisector of T_{vw}. Since T_{vw} is smaller than T_{uw}, the induction hypothesis gives an upper bound on $\delta(v, w)$. Since $|uv| \leq |ua'| + |a'v|$ and v, a'', a, and a' form a parallelogram, we need to show that $\delta(v, w) \leq |va''| + \boldsymbol{c} \cdot |a''w|$ for both cases in order to complete the proof.

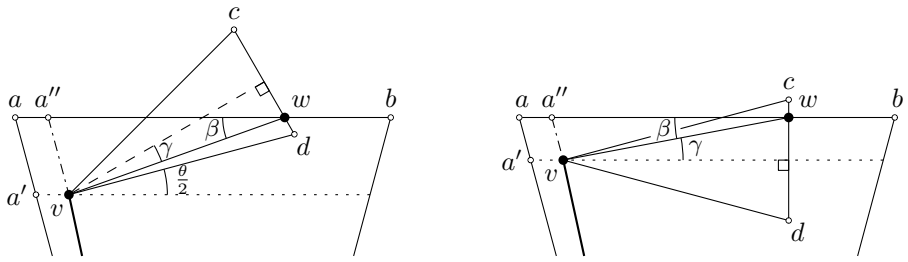

Figure 5.4: The remaining cases of the induction step for the $\theta_{(4k+4)}$-graph: (a) w lies in C_k^v and $|cw| > |dw|$, (b) w lies in C_{k+1}^v

Case (c): When w lies in C_k^v and $|cw| > |dw|$, the induction hypothesis for T_{vw} gives $\delta(v, w) \leq |vc| + \boldsymbol{c} \cdot |cw|$ (see Figure 5.4a). We note that $\gamma = \theta - \beta$. Hence, the inequality follows from Lemma 2.5 when $\boldsymbol{c} \geq (\cos(\theta - \beta) - \sin\beta)/(\cos(\theta/2 - \beta) - \sin(3\theta/2 - \beta))$. Since this function is decreasing in β for $\theta/2 \leq \beta \leq \theta$, it is maximized when β equals $\theta/2$. Hence, $\boldsymbol{c}$ needs to be at least $(\cos(\theta/2) - \sin(\theta/2))/(1 - \sin\theta)$, which can be rewritten to $1/(\cos(\theta/2) - \sin(\theta/2))$.

Case (d): When w lies in C_{k+1}^v, w lies above the bisector of T_{vw} (see Figure 5.4b) and the induction hypothesis for T_{vw} gives $\delta(v, w) \leq |wd| + \boldsymbol{c} \cdot |dv|$. We note that $\gamma = \beta$. Hence, the inequality follows from Lemma 2.5 when $\boldsymbol{c} \geq (\cos\beta - \sin\beta)/(\cos(\theta/2 - \beta) - \sin(\theta/2 + \beta))$, which is equal to $1/(\cos(\theta/2) - \sin(\theta/2))$. □

Since $\cos\alpha/\cos(\theta/2) + (\cos\alpha\tan(\theta/2) + \sin\alpha)/(\cos(\theta/2) - \sin(\theta/2))$ is increasing for $\alpha \in [0, \theta/2]$, for $\theta \leq \pi/4$, it is maximized when $\alpha = \theta/2$, and we obtain the following corollary:

Corollary 5.2 *The $\theta_{(4k+4)}$-graph ($k \geq 1$) is a $\left(1 + \frac{2\sin\left(\frac{\theta}{2}\right)}{\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{\theta}{2}\right)}\right)$-spanner.*

Furthermore, we observe that the proof of Theorem 5.3 follows the same path as the θ-routing algorithm follows: if the direct edge to the destination is part of the graph, it follows this edge, and if it is not, it follows the edge to the closest vertex in the cone that contains the destination.

Corollary 5.3 *The θ-routing algorithm is $\left(1+\frac{2\sin\left(\frac{\theta}{2}\right)}{\cos\left(\frac{\theta}{2}\right)-\sin\left(\frac{\theta}{2}\right)}\right)$-competitive on the $\theta_{(4k+4)}$-graph ($k \geq 1$).*

5.1.4 Upper Bound on the $\theta_{(4k+3)}$-Graph and $\theta_{(4k+5)}$-Graph

In this section, we improve the upper bounds on the spanning ratio of the $\theta_{(4k+3)}$-graph and the $\theta_{(4k+5)}$-graph, for any integer $k \geq 1$.

Theorem 5.4 *Let u and w be two vertices in the plane. Let m be the midpoint of the side of T_{uw} opposite u and let α be the unsigned angle between uw and um. There exists a path connecting u and w in the $\theta_{(4k+3)}$-graph ($k \geq 1$) of length at most*

$$\left(\frac{\cos\alpha}{\cos\left(\frac{\theta}{2}\right)}+\frac{\left(\cos\alpha\tan\left(\frac{\theta}{2}\right)+\sin\alpha\right)\cdot\cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right)-\sin\left(\frac{3\theta}{4}\right)}\right)\cdot|uw|.$$

Proof. We apply Theorem 5.2 using $\boldsymbol{c}=\cos(\theta/4)/(\cos(\theta/2)-\sin(3\theta/4))$. It remains to handle Case (c), where $w \in C_k^v$ and $|cw| > |dw|$, and Case (d), where $w \in C_{k+1}^v$.

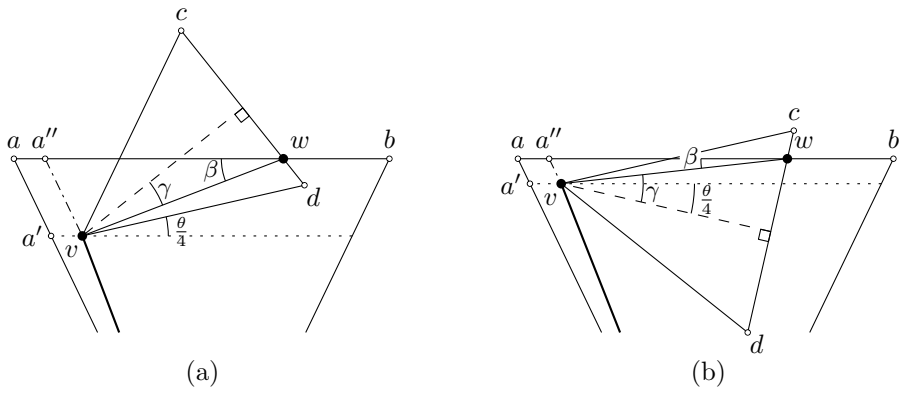

Figure 5.5: The remaining cases of the induction step for the $\theta_{(4k+3)}$-graph: (a) w lies in C_k^v and $|cw| > |dw|$, (b) w lies in C_{k+1}^v

Recall that c and d are the left and right corners of T_{vw}, opposite to v, and a'' is the intersection of the side of T_{uw} opposite u and the left boundary of C_0^v. Let β

be $\angle a''wv$ and let γ be the angle between vw and the bisector of T_{vw}. Since T_{vw} is smaller than T_{uw}, the induction hypothesis gives an upper bound on $\delta(v,w)$. Since $|uv| \leq |ua'| + |a'v|$ and v, a'', a, and a' form a parallelogram, we need to show that $\delta(v,w) \leq |va''| + \boldsymbol{c} \cdot |a''w|$ for both cases in order to complete the proof.

Case (c): When w lies in C_k^v and $|cw| > |dw|$, the induction hypothesis for T_{vw} gives $\delta(v,w) \leq |vc| + \boldsymbol{c} \cdot |cw|$ (see Figure 5.5a). We note that $\gamma = 3\theta/4 - \beta$. Hence, the inequality follows from Lemma 2.5 when $\boldsymbol{c} \geq (\cos(3\theta/4 - \beta) - \sin\beta)/(\cos(\theta/2 - \beta) - \sin(5\theta/4 - \beta))$. Since this function is decreasing in β for $\theta/4 \leq \beta \leq 3\theta/4$, it is maximized when β equals $\theta/4$. Hence, $\boldsymbol{c}$ needs to be at least $(\cos(\theta/2) - \sin(\theta/4))/(\cos(\theta/4) - \sin\theta)$, which is equal to $\cos(\theta/4)/(\cos(\theta/2) - \sin(3\theta/4))$.

Case (d): When w lies in C_{k+1}^v, w lies above the bisector of T_{vw} (see Figure 5.5b) and the induction hypothesis for T_{vw} gives $\delta(v,w) \leq |wd| + \boldsymbol{c} \cdot |dv|$. We note that $\gamma = \theta/4 + \beta$. Hence, the inequality follows from Lemma 2.5 when $\boldsymbol{c} \geq (\cos(\theta/4 + \beta) - \sin\beta)/(\cos(\theta/2 - \beta) - \sin(3\theta/4 + \beta))$, which is equal to $\cos(\theta/4)/(\cos(\theta/2) - \sin(3\theta/4))$. □

Theorem 5.5 *Let u and w be two vertices in the plane. Let m be the midpoint of the side of T_{uw} opposite u and let α be the unsigned angle between uw and um. There exists a path connecting u and w in the $\theta_{(4k+5)}$-graph ($k \geq 1$) of length at most*

$$\left(\frac{\cos\alpha}{\cos\left(\frac{\theta}{2}\right)} + \frac{\left(\cos\alpha \tan\left(\frac{\theta}{2}\right) + \sin\alpha\right) \cdot \cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{3\theta}{4}\right)} \right) \cdot |uw|.$$

Proof. We apply Theorem 5.2 using $\boldsymbol{c} = \cos(\theta/4)/(\cos(\theta/2) - \sin(3\theta/4))$. It remains to handle Case (c), where $w \in C_k^v$ and $|cw| > |dw|$, and Case (d), where $w \in C_{k+1}^v$.

Recall that c and d are the left and right corners of T_{vw}, opposite to v, and a'' is the intersection of the side of T_{uw} opposite u and the left boundary of C_0^v. Let β be $\angle a''wv$ and let γ be the angle between vw and the bisector of T_{vw}. Since T_{vw} is smaller than T_{uw}, the induction hypothesis gives an upper bound on $\delta(v,w)$. Since $|uv| \leq |ua'| + |a'v|$ and v, a'', a, and a' form a parallelogram, we need to show that $\delta(v,w) \leq |va''| + \boldsymbol{c} \cdot |a''w|$ for both cases in order to complete the proof.

Case (c): When w lies in C_k^v and $|cw| > |dw|$, the induction hypothesis for T_{vw} gives $\delta(v,w) \leq |vc| + \boldsymbol{c} \cdot |cw|$ (see Figure 5.6a). We note that $\gamma = 5\theta/4 - \beta$. Hence, the inequality follows from Lemma 2.5 when $\boldsymbol{c} \geq (\cos(5\theta/4 - \beta) - \sin\beta)/(\cos(\theta/2 - \beta) - \sin(7\theta/4 - \beta))$. Since this function is decreasing in β for $3\theta/4 \leq \beta \leq 5\theta/4$, it is maximized when β equals $3\theta/4$. Hence, $\boldsymbol{c}$ needs to be at least $(\cos(\theta/2) - \sin(3\theta/4))/(\cos(\theta/4) - \sin\theta)$, which is less than $\cos(\theta/4)/(\cos(\theta/2) - \sin(3\theta/4))$.

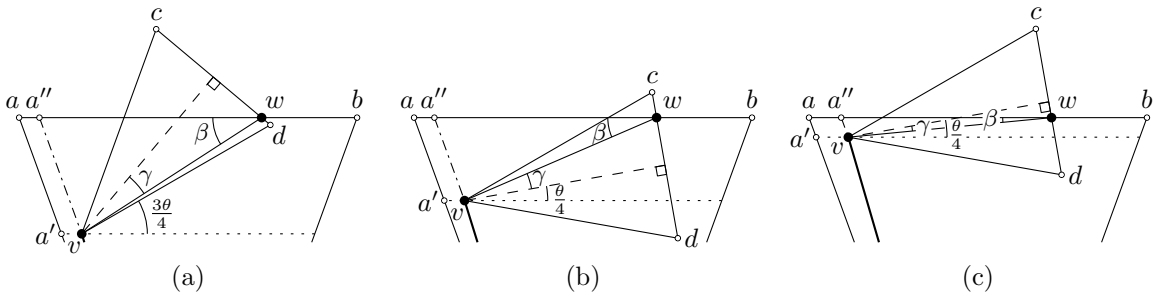

Figure 5.6: The remaining cases of the induction step for the $\theta_{(4k+5)}$-graph: (a) w lies in C_k^v and $|cw| > |dw|$, (b) w lies in C_{k+1}^v and $|cw| < |dw|$, (c) w lies in C_{k+1}^v and $|cw| \geq |dw|$

Case (d): When w lies in C_{k+1}^v, the induction hypothesis for T_{vw} gives $\delta(v,w) \leq \max\{|vc| + \boldsymbol{c} \cdot |cw|, |vd| + \boldsymbol{c} \cdot |dw|\}$. If $|cw| < |dw|$ (see Figure 5.6b), the induction hypothesis for T_{vw} gives $\delta(v,w) \leq |vd| + \boldsymbol{c} \cdot |dw|$. We note that $\gamma = \beta - \theta/4$. Hence, the inequality follows from Lemma 2.5 when $\boldsymbol{c} \geq (\cos(\beta - \theta/4) - \sin\beta)/(\cos(\theta/2 - \beta) - \sin(\theta/4 + \beta))$, which is equal to $\cos(\theta/4)/(\cos(\theta/2) - \sin(3\theta/4))$.

If $|cw| \geq |dw|$, the induction hypothesis for T_{vw} gives $\delta(v,w) \leq |vc| + \boldsymbol{c} \cdot |cw|$ (see Figure 5.6c). We note that $\gamma = \theta/4 - \beta$. Hence, the inequality follows from Lemma 2.5 when $\boldsymbol{c} \geq (\cos(\theta/4 - \beta) - \sin\beta)/(\cos(\theta/2 - \beta) - \sin(3\theta/4 - \beta))$. Since this function is decreasing in β for $0 \leq \beta \leq \theta/4$, it is maximized when β equals 0. Hence, $\boldsymbol{c}$ needs to be at least $\cos(\theta/4)/(\cos(\theta/2) - \sin(3\theta/4))$. □

By looking at two vertices u and w in the $\theta_{(4k+3)}$-graph and the $\theta_{(4k+5)}$-graph, we can see that when the angle between uw and the bisector of T_{uw} is α, the angle between wu and the bisector of T_{wu} is $\theta/2 - \alpha$. Hence the worst case spanning ratio corresponds to the minimum of the spanning ratio when looking at T_{uw} and the spanning ratio when looking at T_{wu}.

Theorem 5.6 *The $\theta_{(4k+3)}$-graph and $\theta_{(4k+5)}$-graph ($k \geq 1$) are $\frac{\cos(\frac{\theta}{4})}{\cos(\frac{\theta}{2}) - \sin(\frac{3\theta}{4})}$-spanners.*

Proof. The spanning ratio of the $\theta_{(4k+3)}$-graph and the $\theta_{(4k+5)}$-graph is at most

$$\min \left\{ \begin{array}{l} \frac{\cos\alpha}{\cos(\frac{\theta}{2})} + \frac{(\cos\alpha \tan(\frac{\theta}{2}) + \sin\alpha) \cdot \cos(\frac{\theta}{4})}{\cos(\frac{\theta}{2}) - \sin(\frac{3\theta}{4})}, \\ \frac{\cos(\frac{\theta}{2} - \alpha)}{\cos(\frac{\theta}{2})} + \frac{(\cos(\frac{\theta}{2} - \alpha) \tan(\frac{\theta}{2}) + \sin(\frac{\theta}{2} - \alpha)) \cdot \cos(\frac{\theta}{4})}{\cos(\frac{\theta}{2}) - \sin(\frac{3\theta}{4})} \end{array} \right\}.$$

Since $\cos\alpha/\cos(\theta/2) + \boldsymbol{c}\cdot(\cos\alpha\tan(\theta/2) + \sin\alpha)$ is increasing for $\alpha \in [0, \theta/2]$, for $\theta \leq 2\pi/7$, the minimum of these two functions is maximized when the two functions are equal, i.e. when $\alpha = \theta/4$. Thus the $\theta_{(4k+3)}$-graph and the $\theta_{(4k+5)}$-graph have spanning ratio at most

$$\frac{\cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right)} + \frac{\left(\cos\left(\frac{\theta}{4}\right)\tan\left(\frac{\theta}{2}\right) + \sin\left(\frac{\theta}{4}\right)\right)\cdot\cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{3\theta}{4}\right)} = \frac{\cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{3\theta}{4}\right)}. \qquad \square$$

Furthermore, we observe that the proofs of Theorem 5.4 and Theorem 5.5 follow the same path as the θ-routing algorithm follows. Since in the case of routing, we are forced to consider the canonical triangle with the source as apex, the arguments that decreased the spanning ratio cannot be applied. Hence, we obtain the following corollary.

Corollary 5.4 *The θ-routing algorithm is $\left(1 + \frac{2\sin\left(\frac{\theta}{2}\right)\cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right)-\sin\left(\frac{3\theta}{4}\right)}\right)$-competitive on the $\theta_{(4k+3)}$-graph and the $\theta_{(4k+5)}$-graph ($k \geq 1$).*

5.2 Lower Bounds

In this section, we provide lower bounds for the $\theta_{(4k+3)}$-graph, the $\theta_{(4k+4)}$-graph, and the $\theta_{(4k+5)}$-graph. For each of the families, we construct a lower bound example by extending the shortest path between two vertices u and w. For brevity, we describe only how to extend one of the shortest paths between these vertices. To extend all shortest paths between u and w, the same transformation is applied to all equivalent paths or canonical triangles.

For example, when constructing the lower bound for the $\theta_{(4k+3)}$-graph, our first step is to ensure that there is no edge between u and w. To this end, the proof of Theorem 5.7 states that we place a vertex v_1 in the corner of T_{uw} that is furthest from w. Placing only this single vertex, however, does not prevent the edge uw from being present, as u is still the closest vertex in T_{wu}. Hence, we also place a vertex in the corner of T_{wu} that is furthest from u. Since these two modifications are essentially the same, but applied to different canonical triangles, we describe only the placement of one of these vertices. The full result of each step is shown in the accompanying figures.

5.2.1 Lower Bound on the $\theta_{(4k+3)}$-Graph

In this section, we construct a lower bound on the spanning ratio of the $\theta_{(4k+3)}$-graph, for any integer $k \geq 1$.

Theorem 5.7 *The worst case spanning ratio of the $\theta_{(4k+3)}$-graph ($k \geq 1$) is at least*

$$\frac{3\cos\left(\frac{\theta}{4}\right) + \cos\left(\frac{3\theta}{4}\right) + \sin\left(\frac{\theta}{2}\right) + \sin\theta + \sin\left(\frac{3\theta}{2}\right)}{3\cos\left(\frac{\theta}{2}\right) + \cos\left(\frac{3\theta}{2}\right)}.$$

Proof. We construct the lower bound example by extending the shortest path between two vertices u and w in three steps. We describe only how to extend one of the shortest paths between these vertices. To extend all shortest paths, the same modification is performed in each of the analogous cases, as shown in Figure 5.7.

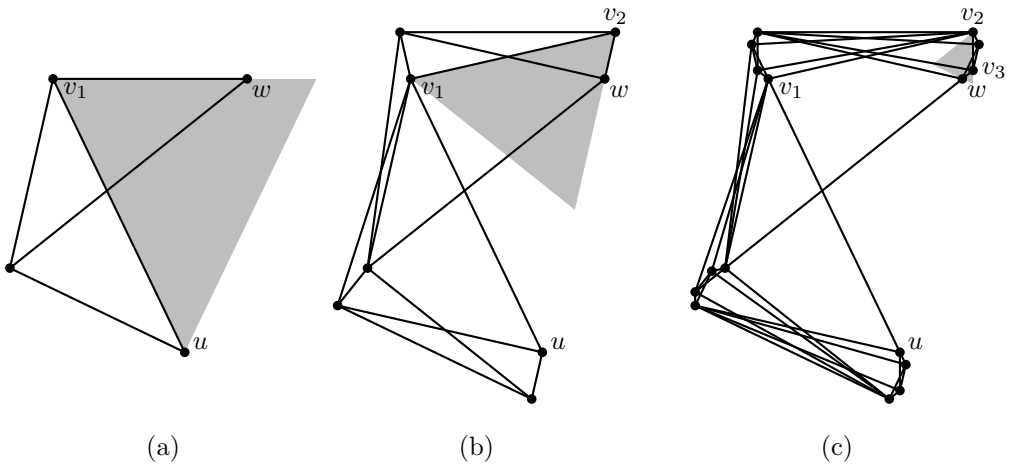

Figure 5.7: The construction of the lower bound for the $\theta_{(4k+3)}$-graph

First, we place w such that the angle between uw and the bisector of the cone of u that contains w is $\theta/4$. Next, we ensure that there is no edge between u and w by placing a vertex v_1 in the upper corner of T_{uw} that is furthest from w (see Figure 5.7a). Next, we place a vertex v_2 in the corner of T_{v_1w} that lies outside T_{uw} (see Figure 5.7b). Finally, to ensure that there is no edge between v_2 and w, we place a vertex v_3 in T_{v_2w} such that T_{v_2w} and T_{v_3w} have the same orientation (see Figure 5.7c). Note that we cannot place v_3 in the lower right corner of T_{v_2w} since this would cause an edge between u and v_3 to be added, creating a shortcut to w.

One of the shortest paths in the resulting graph visits u, v_1, v_2, v_3, and w. Thus, to obtain a lower bound for the $\theta_{(4k+3)}$-graph, we compute the length of this path.

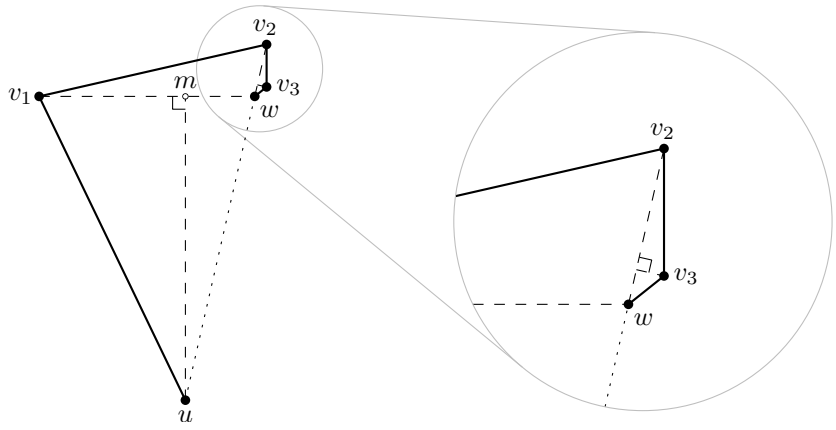

Figure 5.8: The lower bound for the $\theta_{(4k+3)}$-graph

Let m be the midpoint of the side of T_{uw} opposite u. By construction, we have that $\angle v_1um = \theta/2$, $\angle wum = \angle v_2v_1w = \angle v_3v_2w = \theta/4$, $\angle v_3wv_2 = 3\theta/4$, $\angle uv_1w = \angle v_1v_2w = \pi/2 - \theta/2$, and $\angle v_2v_3w = \pi - \theta$ (see Figure 5.8). We can express the various line segments as follows:

$$
\begin{aligned}
|uv_1| &= \frac{\cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right)} \cdot |uw| \\
|v_1w| &= \frac{\sin\left(\frac{3\theta}{4}\right)}{\sin\left(\frac{\pi}{2} - \frac{\theta}{2}\right)} \cdot |uw| = \frac{\sin\left(\frac{3\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right)} \cdot |uw| \\
|v_1v_2| &= \frac{\cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right)} \cdot |v_1w| \\
|v_2w| &= \frac{\sin\left(\frac{\theta}{4}\right)}{\sin\left(\frac{\pi}{2} - \frac{\theta}{2}\right)} \cdot |v_1w| = \frac{\sin\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right)} \cdot |v_1w| \\
|v_2v_3| &= \frac{\sin\left(\frac{3\theta}{4}\right)}{\sin(\pi - \theta)} \cdot |v_2w| = \frac{\sin\left(\frac{3\theta}{4}\right)}{\sin(\theta)} \cdot |v_2w| \\
|v_3w| &= \frac{\sin\left(\frac{\theta}{4}\right)}{\sin(\pi - \theta)} \cdot |v_2w| = \frac{\sin\left(\frac{\theta}{4}\right)}{\sin(\theta)} \cdot |v_2w|
\end{aligned}
$$

Hence, the total length of the shortest path is $|uv_1| + |v_1v_2| + |v_2v_3| + |v_3w|$, which can be rewritten to

$$\frac{3\cos\left(\frac{\theta}{4}\right) + \cos\left(\frac{3\theta}{4}\right) + \sin\left(\frac{\theta}{2}\right) + \sin\theta + \sin\left(\frac{3\theta}{2}\right)}{3\cos\left(\frac{\theta}{2}\right) + \cos\left(\frac{3\theta}{2}\right)} \cdot |uw|,$$

proving the theorem. □

5.2.2 Lower Bound on the $\theta_{(4k+4)}$-Graph

The $\theta_{(4k+2)}$-graph has the nice property that any line perpendicular to the bisector of a cone is parallel to the boundary of a cone (Lemma 2.2). As a result of this, if u, v, and w are vertices with v in one of the upper corners of T_{uw}, then T_{wv} is completely contained in T_{uw}. The $\theta_{(4k+4)}$-graph does not have this property. In this section, we show how to exploit this to construct a lower bound for the $\theta_{(4k+4)}$-graph whose spanning ratio exceeds the worst case spanning ratio of the $\theta_{(4k+2)}$-graph.

Theorem 5.8 *The worst case spanning ratio of the $\theta_{(4k+4)}$-graph ($k \geq 1$) is at least*

$$1 + 2\tan\left(\frac{\theta}{2}\right) + 2\tan^2\left(\frac{\theta}{2}\right).$$

Proof. We construct the lower bound example by extending the shortest path between two vertices u and w in three steps. We describe only how to extend one of the shortest paths between these vertices. To extend all shortest paths, the same modification is performed in each of the analogous cases, as shown in Figure 5.9.

First, we place w such that the angle between uw and the bisector of the cone of u that contains w is $\theta/2$. Next, we ensure that there is no edge between u and w by placing a vertex v_1 in the upper corner of T_{uw} that is furthest from w (see Figure 5.9a). Next, we place a vertex v_2 in the corner of T_{v_1w} that lies in the same cone of u as w and v_1 (see Figure 5.9b). Finally, we place a vertex v_3 in the intersection of the left boundary of T_{v_2w} and the right boundary of T_{wv_2} to ensure that there is no edge between v_2 and w (see Figure 5.9c). Note that we cannot place v_3 in the lower right corner of T_{v_2w} since this would cause an edge between u and v_3 to be added, creating a shortcut to w.

One of the shortest paths in the resulting graph visits u, v_1, v_2, v_3, and w. Thus, to obtain a lower bound for the $\theta_{(4k+4)}$-graph, we compute the length of this path.

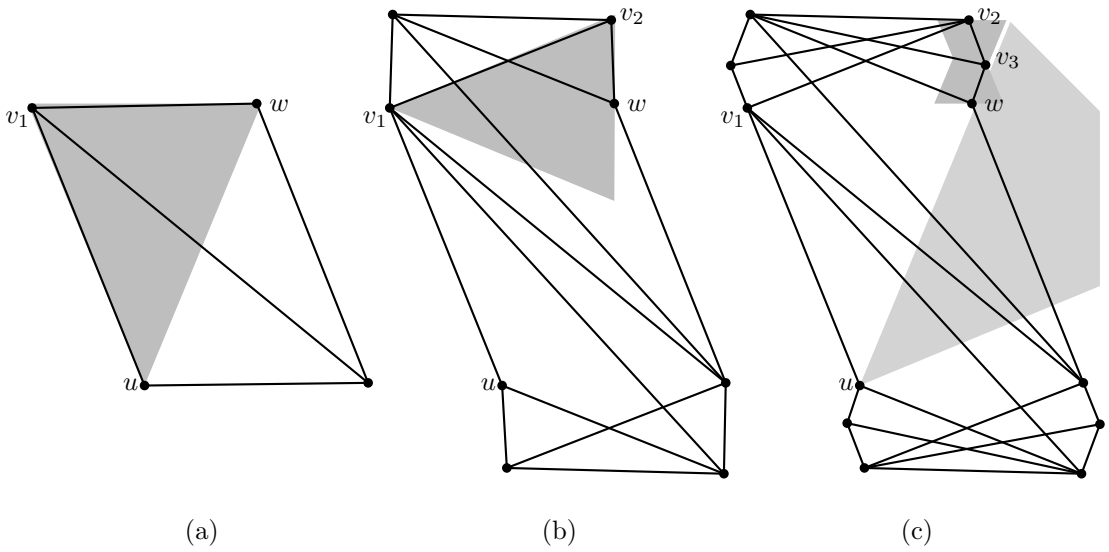

Figure 5.9: The construction of the lower bound for the $\theta_{(4k+4)}$-graph

Let m be the midpoint of the side of T_{uw} opposite u. By construction, we have that $\angle v_1um = \angle wum = \angle v_2v_1w = \angle v_3v_2w = \angle v_3wv_2 = \theta/2$ (see Figure 5.10). We can express the various line segments as follows:

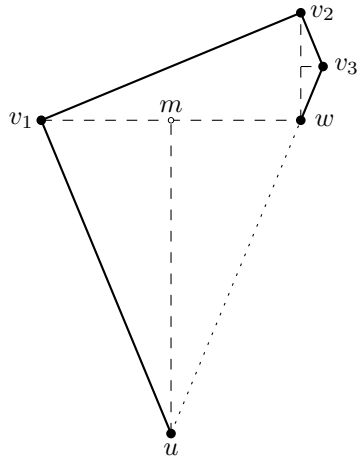

Figure 5.10: The lower bound for the $\theta_{(4k+4)}$-graph

$$
\begin{aligned}
|uv_1| &= |uw| \\
|v_1w| &= 2\sin\left(\frac{\theta}{2}\right)\cdot|uw| \\
|v_1v_2| &= \frac{|v_1w|}{\cos\left(\frac{\theta}{2}\right)} = 2\tan\left(\frac{\theta}{2}\right)\cdot|uw| \\
|v_2w| &= \tan\left(\frac{\theta}{2}\right)\cdot|v_1w| = 2\sin\left(\frac{\theta}{2}\right)\tan\left(\frac{\theta}{2}\right)\cdot|uw| \\
|v_2v_3| &= |v_3w| = \frac{\frac{1}{2}|v_1w|}{\cos\left(\frac{\theta}{2}\right)} = \tan^2\left(\frac{\theta}{2}\right)\cdot|uw|
\end{aligned}
$$

Hence, the total length of the shortest path is $|uv_1|+|v_1v_2|+|v_2v_3|+|v_3w|$, which can be rewritten to

$$\left(1+2\tan\left(\frac{\theta}{2}\right)+2\tan^2\left(\frac{\theta}{2}\right)\right)\cdot|uw|. \qquad \square$$

5.2.3 Lower Bound on the $\theta_{(4k+5)}$-Graph

In this section, we give a lower bound on the spanning ratio of the $\theta_{(4k+5)}$-graph, for any integer $k \geq 1$.

Theorem 5.9 *The worst case spanning ratio of the $\theta_{(4k+5)}$-graph ($k \geq 1$) is at least*

$$\frac{1}{2}\sqrt{4\sec\left(\frac{\theta}{2}\right)+7\sec^2\left(\frac{\theta}{2}\right)+4\sec^3\left(\frac{\theta}{2}\right)+\sec^4\left(\frac{\theta}{2}\right)-8\cos\left(\frac{\theta}{2}\right)-4}$$
$$+\tan\left(\frac{\theta}{2}\right)+\frac{1}{2}\sec\left(\frac{\theta}{2}\right)\tan\left(\frac{\theta}{2}\right).$$

Proof. We construct the lower bound example by extending the shortest path between two vertices u and w in two steps. We describe only how to extend one of the shortest paths between these vertices. To extend all shortest paths, the same modification is performed in each of the analogous cases, as shown in Figure 5.11.

First, we place w such that the angle between uw and the bisector of the cone of u that contains w is $\theta/4$. Next, we ensure that there is no edge between u and w by placing a vertex v_1 in the upper corner of T_{uw} that is furthest from w (see

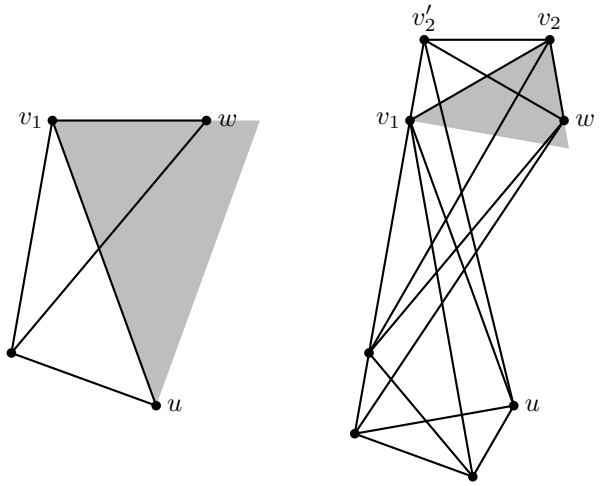

Figure 5.11: The construction of the lower bound for the $\theta_{(4k+5)}$-graph

Figure 5.11a). Finally, we place a vertex v_2 in the corner of T_{v_1w} that lies outside T_{uw}. We also place a vertex v_2' in the corner of T_{wv_1} that lies in the same cone of u as w and v_1 (see Figure 5.11b). Note that placing v_2' creates a shortcut between u and v_2', as u is the closest vertex in one of the cones of v_2'.

One of the shortest paths in the resulting graph visits u, v_2', and w. Thus, to obtain a lower bound for the $\theta_{(4k+5)}$-graph, we compute the length of this path.

Let m be the midpoint of the side of T_{uw} opposite u. By construction, we have that $\angle v_1um = \theta/2$, $\angle wum = \theta/4$, $\angle v_1wv_2' = 3\theta/4$, and $\angle uv_1v_2' = \angle uv_1w + \angle wv_1v_2' = (\pi - \theta)/2 + (\pi - (\pi - \theta)/2 - 3\theta/4) = \pi - 3\theta/4$ (see Figure 5.12). We can express the various line segments as follows:

$$|uv_1| = \frac{\cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right)} \cdot |uw|$$

$$|v_2'w| = \frac{\cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right)} \cdot \left(\sin\left(\frac{\theta}{4}\right) + \cos\left(\frac{\theta}{4}\right)\tan\left(\frac{\theta}{2}\right)\right) \cdot |uw|$$

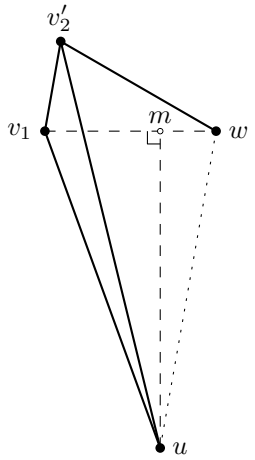

Figure 5.12: The lower bound for the $\theta_{(4k+5)}$-graph

$$|v_1v_2'| = \left(\sin\left(\frac{\theta}{4}\right) + \cos\left(\frac{\theta}{4}\right)\tan\left(\frac{\theta}{2}\right)\right)^2 \cdot |uw|$$

$$|uv_2'| = \sqrt{|uv_1|^2 + |v_1v_2'|^2 - 2\cdot|uv_1|\cdot|v_1v_2'|\cdot\cos\left(\pi - \frac{3\theta}{4}\right)}$$

Hence, the total length of the shortest path is $|uv_2'|+|v_2'w|$, which can be rewritten to

$$\frac{1}{2}\sqrt{4\sec\left(\frac{\theta}{2}\right) + 7\sec^2\left(\frac{\theta}{2}\right) + 4\sec^3\left(\frac{\theta}{2}\right) + \sec^4\left(\frac{\theta}{2}\right) - 8\cos\left(\frac{\theta}{2}\right) - 4} + \tan\left(\frac{\theta}{2}\right) + \frac{1}{2}\sec\left(\frac{\theta}{2}\right)\tan\left(\frac{\theta}{2}\right)$$

times the length of uw. □

5.3 Comparison

In this section we prove that the upper and lower bounds of the four families of θ-graphs admit a partial ordering. We need the following lemma that can be proved by elementary calculus.

Lemma 5.1 *Let $x \in \left[0, \frac{\pi}{4}\right]$ be a real number. Then the following inequalities hold:*

1. $\sin(x) \leq x$ *with equality if and only if* $x = 0$.
2. $\cos(x) \geq 1 - \frac{x^2}{2}$ *with equality if and only if* $x = 0$.
3. $\sin(x) \geq x - \frac{x^3}{6}$ *with equality if and only if* $x = 0$.
4. $\cos(x) \leq 1 - \frac{x^2}{2} + \frac{x^4}{24}$ *with equality if and only if* $x = 0$.
5. $\tan(x) \geq x$ *with equality if and only if* $x = 0$.
6. $\tan^2(x) \geq x^2$ *with equality if and only if* $x = 0$.

Using the above properties, we proceed to prove a number of relations between the four families of θ-graphs.

Lemma 5.2 *Let $ub(m)$ and $lb(m)$ denote the upper and lower bound on the θ_m-graph:*

$$ub(m) = \begin{cases} 1 + 2\sin\left(\frac{\pi}{4k+2}\right) & \text{if } m = 4k+2 \quad (k \geq 1) \\ \frac{\cos\left(\frac{\pi}{2(4k+3)}\right)}{\cos\left(\frac{\pi}{4k+3}\right) - \sin\left(\frac{3\pi}{2(4k+3)}\right)} & \text{if } m = 4k+3 \quad (k \geq 1) \\ 1 + 2\frac{\sin\left(\frac{\pi}{4k+4}\right)}{\cos\left(\frac{\pi}{4k+4}\right) - \sin\left(\frac{\pi}{4k+4}\right)} & \text{if } m = 4k+4 \quad (k \geq 1) \\ \frac{\cos\left(\frac{\pi}{2(4k+5)}\right)}{\cos\left(\frac{\pi}{4k+5}\right) - \sin\left(\frac{3\pi}{2(4k+5)}\right)} & \text{if } m = 4k+5 \quad (k \geq 1) \end{cases}$$

$$lb(m) = \begin{cases} 1 + 2\sin\left(\frac{\pi}{4k+2}\right) & \text{if } m = 4k+2 \quad (k \geq 1) \\ \frac{3\cos\left(\frac{\pi}{2(4k+3)}\right) + \cos\left(\frac{3\pi}{2(4k+3)}\right) + \sin\left(\frac{\pi}{4k+3}\right) + \sin\left(\frac{2\pi}{4k+3}\right) + \sin\left(\frac{3\pi}{4k+3}\right)}{3\cos\left(\frac{\pi}{4k+3}\right) + \cos\left(\frac{3\pi}{4k+3}\right)} & \text{if } m = 4k+3 \quad (k \geq 1) \\ 1 + 2\tan\left(\frac{\pi}{4k+4}\right) + 2\tan^2\left(\frac{\pi}{4k+4}\right) & \text{if } m = 4k+4 \quad (k \geq 1) \\ \frac{\sqrt{4\sec\left(\frac{\pi}{4k+5}\right) + 7\sec^2\left(\frac{\pi}{4k+5}\right) + 4\sec^3\left(\frac{\pi}{4k+5}\right) + \sec^4\left(\frac{\pi}{4k+5}\right) - 8\cos\left(\frac{\pi}{4k+5}\right) - 4}}{2} & \\ + \tan\left(\frac{\pi}{4k+5}\right) + \frac{1}{2}\sec\left(\frac{\pi}{4k+5}\right)\tan\left(\frac{\pi}{4k+5}\right) & \text{if } m = 4k+5 \quad (k \geq 1) \end{cases}$$

Then the following inequalities hold where k is an integer.

$$ub(4(k+1)+2) < lb(4k+2) \qquad (k \geq 1) \tag{a}$$
$$ub(4(k+1)+3) < lb(4k+3) \qquad (k \geq 1) \tag{b}$$
$$ub(4(k+1)+4) < lb(4k+4) \qquad (k \geq 1) \tag{c}$$
$$ub(4(k+1)+5) < lb(4k+5) \qquad (k \geq 1) \tag{d}$$
$$ub(4k+2) < lb(4k+4) \qquad (k \geq 1) \tag{e}$$
$$ub(4(k+1)+4) < lb(4k+2) \qquad (k \geq 1) \tag{f}$$
$$ub(4(k+1)+5) < lb(4k+3) \qquad (k \geq 1) \tag{g}$$
$$ub(4(k+1)+3) < lb(4k+5) \qquad (k \geq 1) \tag{h}$$
$$ub(4k+5) < lb(4k+2) \qquad (k \geq 2) \tag{i}$$

Proof. We use the same strategy for each inequality. We use the definitions of ub and lb in combination with Lemma 5.1. Notice that the restriction on k in each of these inequalities ensures that we can apply Lemma 5.1. We are then left with an algebraic inequality that can be translated into a polynomial inequality, which is easy to verify.

(a)
$$\begin{aligned}
& ub(4(k+1)+2) \\
&= 1 + 2\sin\left(\frac{\pi}{4(k+1)+2}\right) && \text{by the definition of } ub, \\
&< 1 + 2\left(\frac{\pi}{4(k+1)+2}\right) && \text{by Lemma 5.1-1,} \\
&< 1 + 2\left(\left(\frac{\pi}{4k+2}\right) - \frac{1}{6}\left(\frac{\pi}{4k+2}\right)^3\right) && \text{see below,} \qquad (3) \\
&< 1 + 2\sin\left(\frac{\pi}{4k+2}\right) && \text{by Lemma 5.1-3,} \\
&= lb(4k+2) && \text{by the definition of } lb.
\end{aligned}$$

We now explain why (3) holds. The inequality

$$1 + 2\left(\frac{\pi}{4(k+1)+2}\right) < 1 + 2\left(\left(\frac{\pi}{4k+2}\right) - \frac{1}{6}\left(\frac{\pi}{4k+2}\right)^3\right)$$

can be simplified to

$$192k^2 + \left(192 - 2\pi^2\right)k + \left(48 - 3\pi^2\right) > 0. \tag{4}$$

The largest real root of the polynomial involved in (4) is negative. Moreover, (3) holds for $k = 1$. Therefore, (3) holds for any $k \geq 1$.

(b) The proof is analogous to the one of (a).

(c) The proof is analogous to the one of (a).

(d) We let

$$f(k) = \frac{\cos\left(\frac{\pi}{2(4(k+1)+5)}\right)}{\cos\left(\frac{\pi}{4(k+1)+5}\right) - \sin\left(\frac{3\pi}{2(4(k+1)+5)}\right)},$$

$$r(k) = 4\sec\left(\frac{\pi}{4k+5}\right) + 7\sec^2\left(\frac{\pi}{4k+5}\right) + 4\sec^3\left(\frac{\pi}{4k+5}\right) + \sec^4\left(\frac{\pi}{4k+5}\right) - 8\cos\left(\frac{\pi}{4k+5}\right) - 4,$$

$$g(k) = 2\tan\left(\frac{\pi}{4k+5}\right) + \sec\left(\frac{\pi}{4k+5}\right)\tan\left(\frac{\pi}{4k+5}\right),$$

so that

$$ub(4(k+1)+5) = f(k),$$

$$lb(4k+5) = \frac{\sqrt{r(k)} + g(k)}{2}.$$

Using a proof similar to the one of (a), we can prove that

$$(2f(k) - g(k))^2 < r(k).$$

Using a proof similar to the one of (a), we can prove that $2f(k) - g(k) > 0$, for $k \geq 1$, thus we can proceed as follows

$$2f(k) - g(k) < \sqrt{r(k)}$$

$$f(k) < \frac{\sqrt{r(k)} + g(k)}{2}$$

$$ub(4(k+1)+5) < lb(4k+5).$$

(e) The proof is analogous to the one of (a).

(f) The proof is analogous to the one of (a).

(g) The proof is analogous to the one of (d).

(h) The proof is analogous to the one of (d).

(i) The proof is analogous to the one of (d). □

We note that inequalities (a), (b), (c), and (d) imply that the spanning ratio is monotonic within each of the four families. We also note that increasing the number of cones of a θ-graph by 2 from $4k+2$ to $4k+4$ increases the worst case spanning ratio, thus showing that adding cones can make the spanning ratio worse instead of better. Therefore, the spanning ratio is non-monotonic between families. Figure 5.13 illustrates the partial ordering on the spanning ratios of the four families implied by Lemma 5.2.

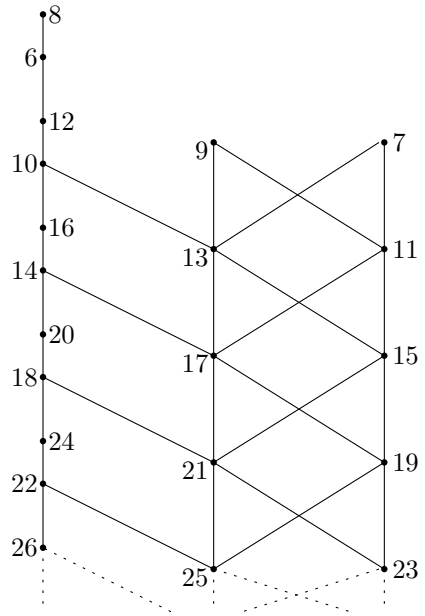

Figure 5.13: Partial order on the spanning ratios of the four families

Finally, to give a better idea of the separation between the upper and lower bounds of the $\theta_{(4k+3)}$-graph, the $\theta_{(4k+4)}$-graph, and the $\theta_{(4k+5)}$-graph, we provide the first few terms of the Taylor series in the following table.

	Upper Bound	Lower Bound
$\theta_{(4k+3)}$-graph	$\mathbf{1 + \frac{3\theta}{4}} + \frac{21\theta^2}{32} + \frac{33\theta^3}{64} + O(\theta^4)$	$\mathbf{1 + \frac{3\theta}{4}} + \frac{9\theta^2}{32} + \frac{3\theta^3}{32} + \Omega(\theta^4)$
$\theta_{(4k+4)}$-graph	$\mathbf{1 + \theta + \frac{\theta^2}{2}} + \frac{\theta^3}{3} + O(\theta^4)$	$\mathbf{1 + \theta + \frac{\theta^2}{2}} + \frac{\theta^3}{12} + \Omega(\theta^4)$
$\theta_{(4k+5)}$-graph	$\mathbf{1 + \frac{3\theta}{4} + \frac{21\theta^2}{32}} + \frac{33\theta^3}{64} + O(\theta^4)$	$\mathbf{1 + \frac{3\theta}{4} + \frac{21\theta^2}{32}} + \frac{3\theta^3}{32} + \Omega(\theta^4)$

Table 5.1: The Taylor series of the upper and lower bounds on the spanning ratio of θ-graphs

5.4 Tight Routing Bounds

While improving the upper bounds on the spanning ratio of the $\theta_{(4k+4)}$-graph, we also improved the upper bound on the routing ratio of the θ-routing algorithm. In this section we show that this bound of $1 + 2\sin(\theta/2)/\left(\cos(\theta/2) - \sin(\theta/2)\right)$ and the current upper bound of $1/\left(1 - 2\sin(\theta/2)\right)$ on the θ_{10}-graph are tight, i.e. we provide matching lower bounds on the routing ratio of the θ-routing algorithm on these families of graphs.

5.4.1 Tight Routing Bounds for the $\theta_{(4k+4)}$-Graph

In this section we show that the upper bound of $1 + (2\sin(\theta/2))/(\cos(\theta/2) - \sin(\theta/2))$ on the routing ratio of the θ-routing algorithm for the $\theta_{(4k+4)}$-graph is a tight bound.

Theorem 5.10 *The θ-routing algorithm is $\left(1 + \frac{2\sin\left(\frac{\theta}{2}\right)}{\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{\theta}{2}\right)}\right)$-competitive on the $\theta_{(4k+4)}$-graph ($k \geq 1$) and this bound is tight.*

Proof. An upper bound of $1 + (2\sin(\theta/2))/(\cos(\theta/2) - \sin(\theta/2))$ on the routing ratio was shown in Corollary 5.3, hence it suffices to show that this is also a lower bound.

We construct the lower bound example on the competitiveness of the θ-routing algorithm on the $\theta_{(4k+4)}$-graph by repeatedly extending the routing path from source u to destination w. First, we place w in the right corner of T_{uw}. To ensure that the θ-routing algorithm does not follow the edge between u and w, we place a vertex v_1 in the left corner of T_{uw}. Next, to ensure that the θ-routing algorithm does not follow the edge between v_1 and w, we place a vertex v_1' in the left corner of T_{v_1w}. We repeat this step until we have created a cycle around w (see Figure 5.14a).

To extend the routing path further, we again place a vertex v_2 in the corner of the current canonical triangle. To ensure that the routing algorithm still routes to v_1 from u, we place v_2 slightly outside of T_{uv_1}. However, another problem arises: vertex v_1' is no longer the vertex closest to v_1 in T_{v_1w}, as v_2 is closer. To solve this problem, we also place a vertex x_1 in $T_{v_1v_2}$ such that v_1' lies in T_{x_1w} (see Figure 5.14b). By repeating this process four times, we create a second cycle around w.

To add more cycles around w, we repeat the same process as described above: place a vertex in the corner of the current canonical triangle and place an auxiliary vertex to ensure that the previous cycle stays intact. Note that when placing x_i, we also need to ensure that it does not lie in $T_{x_{i-1}w}$, to prevent shortcuts from being formed. A lower bound example consisting of two cycles is shown in Figure 5.15.

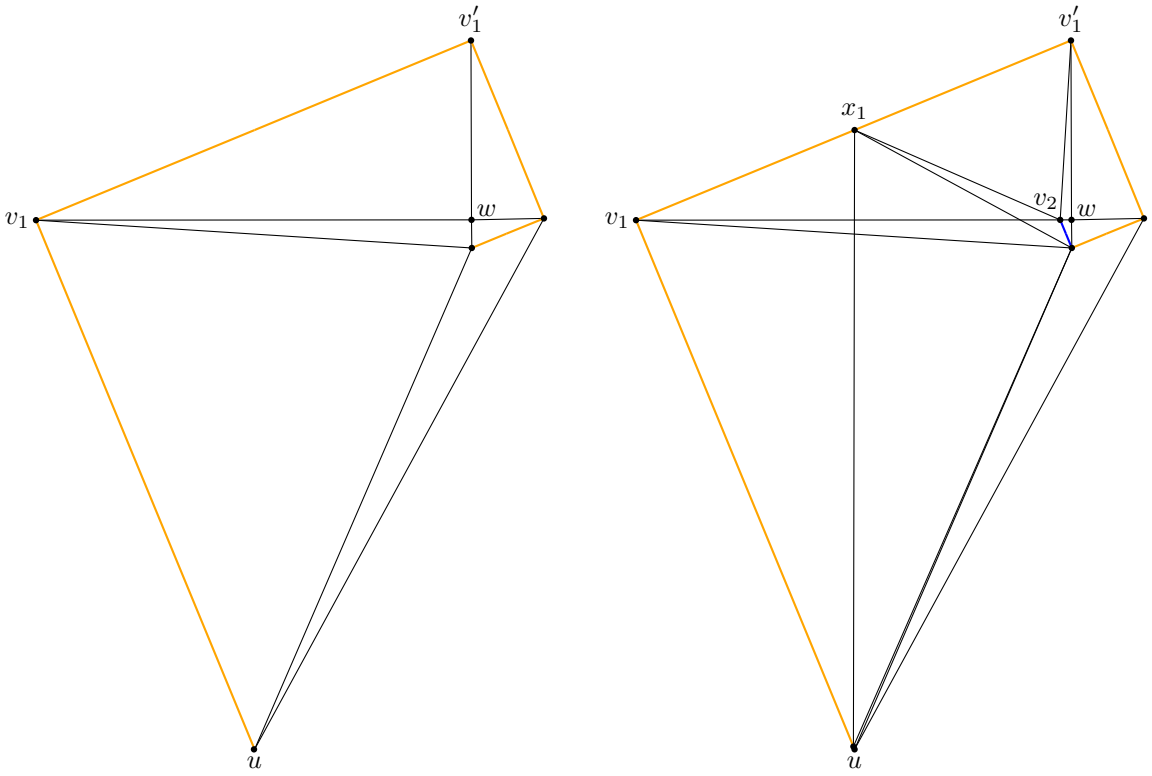

(a) (b)

Figure 5.14: Constructing a lower bound example for θ-routing on the $\theta_{(4k+4)}$-graph: (a) after constructing the first cycle, (b) after adding v_2, the first vertex of the second cycle, and x_1, the auxiliary vertex needed to maintain the first cycle

This way we need to add auxiliary vertices only to the $(k-1)$-th cycle, when adding the k-th cycle, hence we can add an additional cycle using only a constant number of vertices. Since we can place the vertices arbitrarily close to the corners of the canonical triangles, we ensure that $|uv_1| = |uw|$ and that the distance between consecutive vertices v_i and v_i' is always $1/\cos(\theta/2)$ times $|v_iw|$. Hence, when we take $|uw| = 1$ and let the number of vertices approach infinity, we get that the total length of the path is $1 + 2\sin(\theta/2) \cdot \sum_{i=0}^{\infty}(\tan^i(\theta/2)/\cos(\theta/2))$, which can be rewritten to $1 + (2\sin(\theta/2))/(\cos(\theta/2) - \sin(\theta/2))$. □

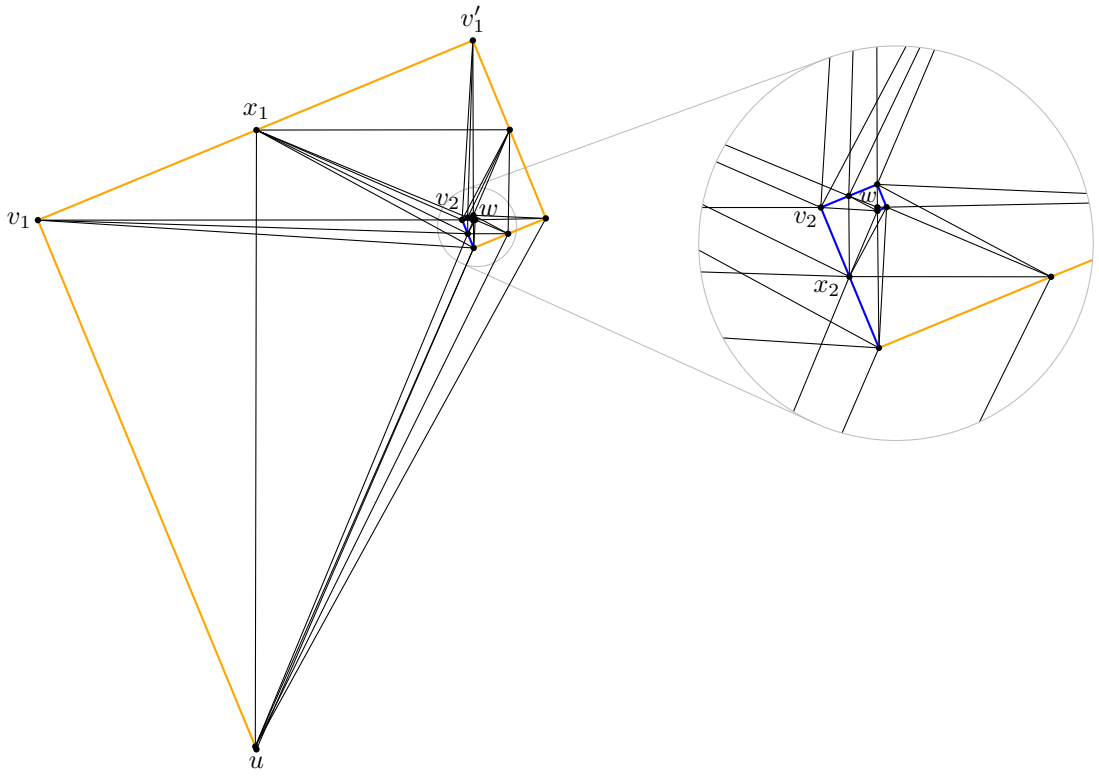

Figure 5.15: A lower bound example for θ-routing on the $\theta_{(4k+4)}$-graph, consisting of two cycles: the first cycle is coloured orange and the second cycle is coloured blue

5.4.2 Tight Routing Bounds for the θ_{10}-Graph

In this section we show that the upper bound of $1/(1-2\sin(\theta/2))$ on the routing ratio of the θ-routing algorithm for the θ_{10}-graph is a tight bound.

Theorem 5.11 *The θ-routing algorithm is $(1/(1-2\sin(\theta/2)))$-competitive on the θ_{10}-graph and this bound is tight.*

Proof. Ruppert and Seidel [55] showed that the routing ratio is at most $1/(1-2\sin(\theta/2))$, hence it suffices to show that this is also a lower bound.

We construct the lower bound example on the competitiveness of the θ-routing algorithm on the θ_{10}-graph by repeatedly extending the routing path from source u to destination w. First, we place w in the right corner of T_{uw}. To ensure that the θ-routing algorithm does not follow the edge between u and w, we place a vertex v_1 in the left corner of T_{uw}. Next, to ensure that the θ-routing algorithm does not follow

the edge between v_1 and w, we place a vertex v_1' in the left corner of T_{v_1w}. We repeat this step until we have created a cycle around w (see Figure 5.16).

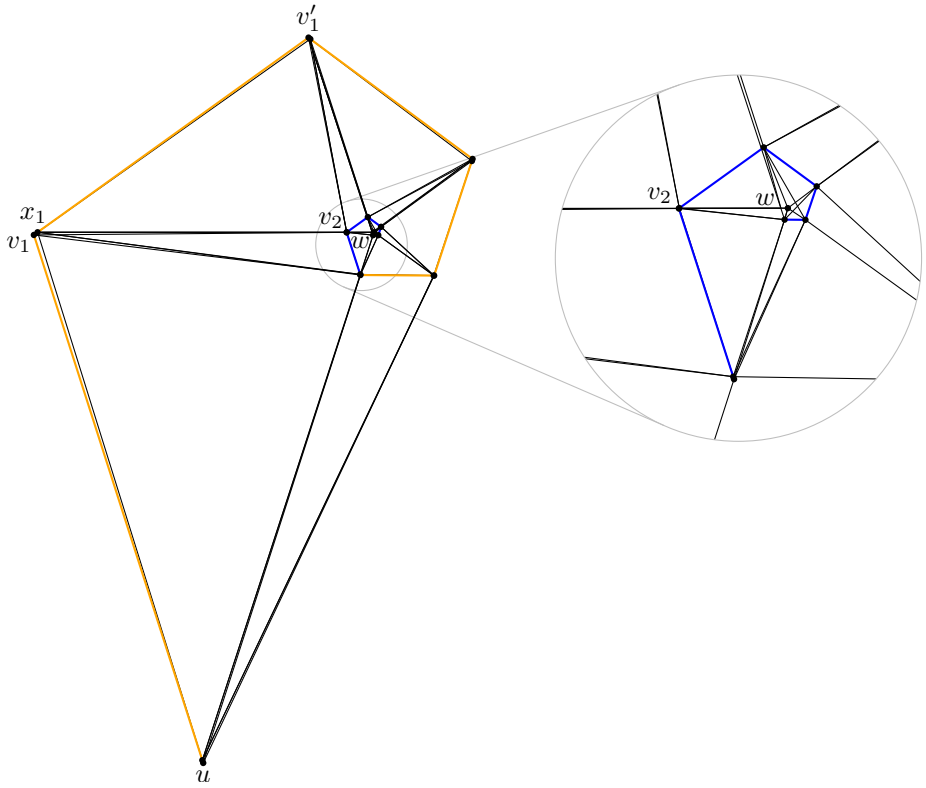

Figure 5.16: A lower bound example for θ-routing on the θ_{10}-graph, consisting of two cycles: the first cycle is coloured orange and the second cycle is coloured blue

To extend the routing path further, we again place a vertex v_2 in the corner of the current canonical triangle. To ensure that the routing algorithm still routes to v_1 from u, we place v_2 slightly outside of T_{uv_1}. However, another problem arises: vertex v_1' is no longer the vertex closest to v_1 in T_{v_1w}, as v_2 is closer. To solve this problem, we also place a vertex x_1 in $T_{v_1v_2}$ such that v_1' lies in T_{x_1w} (see Figure 5.17). By repeating this process four times, we create a second cycle around w.

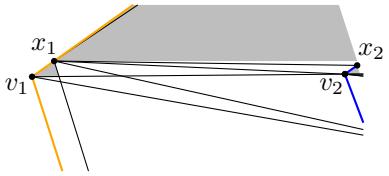

Figure 5.17: The placement of vertices such that previous cycles stay intact when adding a new cycle

To add more cycles around w, we repeat the same process as described above: place a vertex in the corner of the current canonical triangle and place an auxiliary vertex to ensure that the previous cycle stays intact. Note that when placing x_i, we also need to ensure that it does not lie in $T_{x_{i-1}w}$, to prevent shortcuts from being formed (see Figure 5.17). This means that in general x_i does not lie arbitrarily close to the corner of $T_{v_iv_{i+1}}$.

This way we need to add auxiliary vertices only to the $(k-1)$-th cycle, when adding the k-th cycle, hence we can add an additional cycle using only a constant number of vertices. Since we can place the vertices arbitrarily close to the corners of the canonical triangles, we ensure that the distance to w is always $2\sin(\theta/2)$ times the distance between w and the previous vertex along the path. Hence, when we take $|uw| = 1$ and let the number of vertices approach infinity, we get that the total length of the path is $\sum_{i=0}^{\infty} (2\sin(\theta/2))^i$, which can be rewritten to $1/(1-2\sin(\theta/2))$. □

5.5 Conclusion

We showed that the $\theta_{(4k+2)}$-graph has a tight spanning ratio of $1+2\sin(\theta/2)$. This is the first time tight spanning ratios have been found for a large family of θ-graphs. Previously, the only θ-graph for which tight bounds were known was the θ_6-graph. We also gave improved upper bounds on the spanning ratio of the $\theta_{(4k+3)}$-graph, the $\theta_{(4k+4)}$-graph, and the $\theta_{(4k+5)}$-graph.

We also constructed lower bounds for all four families of θ-graphs and provided a partial order on these families. In particular, we showed that the $\theta_{(4k+4)}$-graph has a spanning ratio of at least $1+2\tan(\theta/2)+2\tan^2(\theta/2)$. This result is somewhat surprising since, for equal values of k, the worst case spanning ratio of the $\theta_{(4k+4)}$-graph is greater than that of the $\theta_{(4k+2)}$-graph, showing that increasing the number of cones does not necessarily give a better spanner. In other words, it can happen that

the worst-case spanning ratio of the θ_m-graph is less than that of the θ_{m+1}-graph, due to the different orientation of the cones.

There remain a number of open problems, such as finding tight spanning ratios for the $\theta_{(4k+3)}$-graph, the $\theta_{(4k+4)}$-graph, and the $\theta_{(4k+5)}$-graph. Similarly, for the θ_4 and θ_5-graphs, though upper and lower bounds are known [5, 27], these are far from tight. It would also be nice if we could improve the routing algorithms for θ-graphs. At the moment, θ-routing is the standard routing algorithm for general θ-graphs, but it is unclear whether this is the best routing algorithm for general θ-graphs: though we showed that the current bounds on the competitiveness of the θ-routing algorithm are tight in case of the $\theta_{(4k+4)}$-graph, this does not imply that there exists no algorithm that can do better on these graphs. As a special case, we note that the θ-routing algorithm is not $o(n)$-competitive on the θ_6-graph, but a better (tight) algorithm is known to exist [18].

Chapter 6

Constrained Theta-Graphs

In this chapter, we generalize the results on unconstrained θ-graphs from the previous chapter to the constrained setting. There are two main obstacles that differentiate this work from previous results. First, the main difficulty with the constrained setting is that induction cannot be applied directly, as the destination need not be visible from the vertex closest to the source (see Figure 6.1, where w is not visible from v_0, the vertex closest to u). Second, when the graph does not have $4k+2$ cones, the cones do not line up as nicely as in Chapter 3, making it more difficult to apply induction.

We overcome these two difficulties and show that constrained θ-graphs with $4k+2$ cones have a spanning ratio of at most $1+2\sin(\theta/2)$, where θ is $2\pi/(4k+2)$. Since the lower bounds of the unconstrained θ-graphs carry over to the constrained setting, this shows that this spanning ratio is tight. We also show that constrained θ-graphs with $4k+4$ cones have a spanning ratio of at most $1+2\sin(\theta/2)/\left(\cos(\theta/2)-\sin(\theta/2)\right)$, where θ is $2\pi/(4k+4)$. Finally, we show that constrained θ-graphs with $4k+3$ or $4k+5$ cones have a spanning ratio of at most $\cos(\theta/4)/\left(\cos(\theta/2)-\sin(3\theta/4)\right)$, where θ is $2\pi/(4k+3)$ or $2\pi/(4k+5)$.

The results presented in this chapter are joint work with Prosenjit Bose and appear in the Proceedings of the 11th Latin American Symposium on Theoretical Informatics (LATIN 2014) [30] and in the ArXiv e-prints [31].

6.1 The Constrained $\theta_{(4k+2)}$-Graph

In this section we prove that the constrained $\theta_{(4k+2)}$-graph has spanning ratio at most $1+2\cdot\sin(\theta/2)$ by generalizing Theorem 5.1. Since this is also a lower bound (see Section 5.1.1), this proves that this spanning ratio is tight.

Theorem 6.1 *Let u and w be two vertices in the plane such that u can see w. Let m be the midpoint of the side of T_{uw} opposing u and let α be the unsigned angle between uw and um. There exists a path connecting u and w in the constrained $\theta_{(4k+2)}$-graph ($k \geq 1$) of length at most*

$$\left(\left(\frac{1+\sin\left(\frac{\theta}{2}\right)}{\cos\left(\frac{\theta}{2}\right)}\right)\cdot\cos\alpha+\sin\alpha\right)\cdot|uw|.$$

Proof. We assume without loss of generality that $w \in C_0^u$. 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 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 6.1).

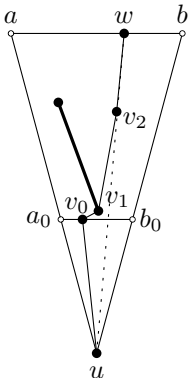

Figure 6.1: A convex chain from v_0 to w

Our inductive hypothesis is the following, where $\delta(u,w)$ denotes the length of the shortest path from u to w in the constrained $\theta_{(4k+2)}$-graph:

- 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 show that this induction hypothesis implies the theorem: $|um| = |uw| \cdot \cos\alpha$, $|mw| = |uw| \cdot \sin\alpha$, $|am| = |bm| = |uw| \cdot \cos\alpha \cdot \tan(\theta/2)$, and $|ua| = |ub| = |uw| \cdot \cos\alpha / \cos(\theta/2)$. Thus the induction hypothesis gives that

$$\delta(u,w) \leq |ua| + |am| + |mw| = \left(\left(\frac{1+\sin\left(\frac{\theta}{2}\right)}{\cos\left(\frac{\theta}{2}\right)}\right)\cdot\cos\alpha+\sin\alpha\right)\cdot|uw|.$$

Base case: T_{uw} has rank 1. Since the triangle is a smallest triangle such that u and w can see each other, w is the closest visible vertex to u in that cone. Hence the edge uw is part of the constrained $\theta_{(4k+2)}$-graph, and $\delta(u, w) = |uw|$. From the triangle inequality, we have $|uw| \leq \min\{|ua| + |aw|, |ub| + |bw|\}$, so the induction hypothesis holds.

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 $\theta_{(4k+2)}$-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 closest visible vertex to u in the subcone of u that contains w, and let a_0 and b_0 be the upper left and right corner of T_{uv_0} (see Figure 6.1). 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 2.1 to triangle v_0uw, a convex chain $v_0, ..., v_l = w$ of visibility edges connecting v_0 and w exists (see Figure 6.1). 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.

We now look at two consecutive vertices v_{j-1} and v_j along the convex chain. There are four types of configurations (see Figure 6.2): (i) $v_j \in C_k^{v_{j-1}}$, (ii) $v_j \in C_i^{v_{j-1}}$ where $1 \leq i < k$, (iii) $v_j \in C_0^{v_{j-1}}$ and v_j lies to the right of or has the same x-coordinate as v_{j-1}, (iv) $v_j \in C_0^{v_{j-1}}$ and v_j lies to the left of v_{j-1}. By convexity, the direction of $\overrightarrow{v_jv_{j+1}}$ is rotating counterclockwise for increasing j. Thus, these configurations occur in the order Type (i), Type (ii), Type (iii), Type (iv) along the convex chain from v_0 to w. We bound $\delta(v_{j-1}, v_j)$ as follows:

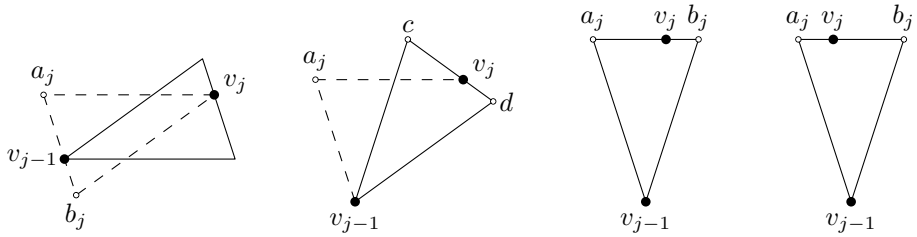

Figure 6.2: The four types of configurations

Type (i): If $v_j \in C_k^{v_{j-1}}$, let a_j and b_j be the upper and lower left corner of $T_{v_j v_{j-1}}$ and let $B_j = v_{j-1}b_j v_j$. Note that since $v_j \in C_k^{v_{j-1}}$, a_j is also the intersection of the left boundary of $C_0^{v_{j-1}}$ and the horizontal line through v_j. Triangle B_j lies between the convex chain and uw, so it must be empty. Since v_j can see v_{j-1} and $T_{v_j v_{j-1}}$ has smaller area than T_{uw}, the induction hypothesis gives that $\delta(v_{j-1}, v_j)$ is at most $|v_{j-1}a_j| + |a_j v_j|$.

Type (ii): If $v_j \in C_i^{v_{j-1}}$ where $1 \leq i < k$, let c and d be the upper and lower right corner of $T_{v_{j-1}v_j}$. Let a_j be the intersection of the left boundary of $C_0^{v_{j-1}}$ and the horizontal line through v_j. Since v_j can see v_{j-1} and $T_{v_{j-1}v_j}$ has smaller area than T_{uw}, the induction hypothesis gives that $\delta(v_{j-1}, v_j)$ is at most $\max\{|v_{j-1}c|+|cv_j|, |v_{j-1}d|+|dv_j|\}$. Since $v_j \in C_i^{v_{j-1}}$ where $1 \leq i < k$, we can apply Lemma 2.4 (where v, w, and a from Lemma 2.4 are v_{j-1}, v_j, and a_j), which gives us that $\max\{|v_{j-1}c|+|cv_j|, |v_{j-1}d|+|dv_j|\} \leq |v_{j-1}a_j| + |a_j v_j|$.

Type (iii): If $v_j \in C_0^{v_{j-1}}$ and v_j lies to the right of or has the same x-coordinate as v_{j-1}, let a_j and b_j be the left and right corner of $T_{v_{j-1}v_j}$ and let $A_j = v_{j-1}a_j v_j$ and $B_j = v_{j-1}b_j v_j$. Since v_j can see v_{j-1} and $T_{v_{j-1}v_j}$ has smaller area than T_{uw}, we can apply the induction hypothesis. Regardless of whether A_j and B_j are empty or not, $\delta(v_{j-1}, v_j)$ is at most $\max\{|v_{j-1}a_j|+|a_j v_j|, |v_{j-1}b_j|+|b_j v_j|\}$. Since v_j lies to the right of or has the same x-coordinate as v_{j-1}, we know that $|v_{j-1}a_j|+|a_j v_j| \geq |v_{j-1}b_j|+|b_j v_j|$, so $\delta(v_{j-1}, v_j)$ is at most $|v_{j-1}a_j| + |a_j v_j|$.

Type (iv): If $v_j \in C_0^{v_{j-1}}$ and v_j lies to the left of v_{j-1}, let a_j and b_j be the left and right corner of $T_{v_{j-1}v_j}$ and let $A_j = v_{j-1}a_j v_j$ and $B_j = v_{j-1}b_j v_j$. Since v_j can see v_{j-1} and $T_{v_{j-1}v_j}$ has smaller area than T_{uw}, we can apply the induction hypothesis. Thus, if B_j is empty, $\delta(v_{j-1}, v_j)$ is at most $|v_{j-1}a_j| + |a_j v_j|$ and if B_j is not empty, $\delta(v_{j-1}, v_j)$ is at most $|v_{j-1}b_j| + |b_j v_j|$.

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 (iv) configurations: for Type (iv) configurations to occur, v_j needs to lie to the left of v_{j-1}. However, by construction, v_j lies on or to the right of the line through v_{j-1} and w. Hence, since $\angle awv_{j-1} < \angle awu \leq \pi/2$, v_j lies to the right of or has the same

x-coordinate as v_{j-1}. We can now bound $\delta(u, w)$ by using these bounds:

$$\begin{aligned}\delta(u,w) &\leq |uv_0| + \sum_{j=1}^{l} \delta(v_{j-1}, v_j) \\ &\leq |ua_0| + |a_0v_0| + \sum_{j=1}^{l} (|v_{j-1}a_j| + |a_jv_j|) \\ &= |ua| + |aw|\end{aligned}$$

Case (b): If $\angle awu > \pi/2$ and B is empty, the convex chain can contain Type (iv) configurations. However, since B is empty and the area between the convex chain and uw is empty (by Lemma 2.1), all B_j are also empty. Using the computed bounds on the lengths of the paths between the points along the convex chain, we can bound $\delta(u, w)$ as in the previous case.

Case (c): If $\angle awu > \pi/2$ and B is not empty, the convex chain can contain Type (iv) configurations and since B is not empty, the triangles B_j 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 $T_{v_{j'}v_{j'+1}}$ be the first Type (iv) 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 $T_{v_{j'}w}$. We can bound $\delta(u, w)$ as follows (see Figure 6.3):

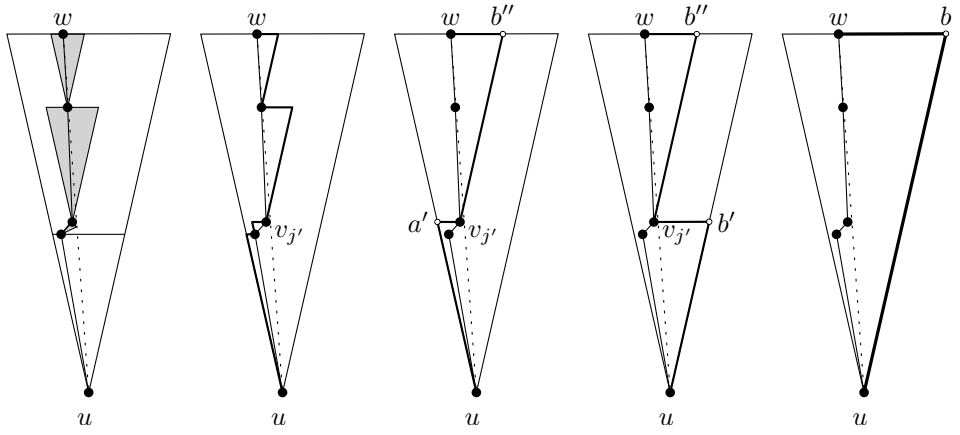

Figure 6.3: Visualization of the paths (thick lines) in the inequalities of case (c)

$$\begin{aligned}
\delta(u,w) &\leq |uv_0| + \sum_{j=1}^{l} \delta(v_{j-1}, v_j) \\
&\leq |ua_0| + |a_0v_0| + \sum_{j=1}^{j'} (|v_{j-1}a_j| + |a_jv_j|) + \sum_{j=j'+1}^{l} (|v_{j-1}b_j| + |b_jv_j|) \\
&= |ua'| + |a'v_{j'}| + |v_{j'}b''| + |b''w| \\
&\leq |ub'| + |b'v_{j'}| + |v_{j'}b''| + |b''w| \\
&= |ub| + |bw|
\end{aligned}$$

□

Since $((1+\sin(\theta/2))\,/\,\cos(\theta/2)) \cdot \cos\alpha + \sin\alpha$ is increasing for $\alpha \in [0, \theta/2]$, for $\theta \leq \pi/3$, it is maximized when $\alpha = \theta/2$, and we obtain the following corollary:

Corollary 6.1 *The constrained $\theta_{(4k+2)}$-graph ($k \geq 1$) is a $\left(1 + 2 \cdot \sin\left(\frac{\theta}{2}\right)\right)$-spanner of $Vis(P, S)$.*

6.2 Generic Framework for the Spanning Proof

Next, we modify the spanning proof from the previous section and provide a generic framework for the spanning proof for the other three families of θ-graphs. After providing this framework, we complete the proofs for the individual families. These proofs are generalizations of Theorems 5.2, 5.3, 5.4, and 5.5.

Theorem 6.2 *Let u and w be two vertices in the plane such that u can see w. Let m be the midpoint of the side of T_{uw} opposing u and let α be the unsigned angle between uw and um. There exists a path connecting u and w in the constrained $\theta_{(4k+x)}$-graph ($k \geq 1$) of length at most*

$$\left(\frac{\cos\alpha}{\cos\left(\frac{\theta}{2}\right)} + \boldsymbol{c} \cdot \left(\cos\alpha \cdot \tan\left(\frac{\theta}{2}\right) + \sin\alpha\right)\right) \cdot |uw|,$$

where $\boldsymbol{c} \geq 1$ is a function that depends on $x \in \{3, 4, 5\}$ and θ. For the $\theta_{(4k+4)}$-graph, $\boldsymbol{c}$ equals $1/(\cos(\theta/2) - \sin(\theta/2))$ and for the $\theta_{(4k+3)}$-graph and $\theta_{(4k+5)}$-graph, $\boldsymbol{c}$ equals $\cos(\theta/4)/\ (\cos(\theta/2) - \sin(3\theta/4))$.

Proof. 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. We assume without loss of generality that $w \in C_0^u$. Let a and b be the upper left and right corner of T_{uw} (see Figure 6.1).

Our inductive hypothesis is the following, where $\delta(u, w)$ denotes the length of the shortest path from u to w in the constrained $\theta_{(4k+x)}$-graph: $\delta(u, w) \leq \max\{|ua| + \boldsymbol{c} \cdot |aw|, |ub| + \boldsymbol{c} \cdot |bw|\}$.

We first show that this induction hypothesis implies the theorem. Basic trigonometry gives us the following equalities: $|um| = |uw| \cdot \cos\alpha$, $|mw| = |uw| \cdot \sin\alpha$, $|am| = |bm| = |uw| \cdot \cos\alpha \cdot \tan(\theta/2)$, and $|ua| = |ub| = |uw| \cdot \cos\alpha / \cos(\theta/2)$. Thus the induction hypothesis gives that

$$\delta(u, w) \leq |ua| + \boldsymbol{c} \cdot (|am| + |mw|) = \left(\frac{\cos\alpha}{\cos\left(\frac{\theta}{2}\right)} + \boldsymbol{c} \cdot \left(\cos\alpha \tan\left(\frac{\theta}{2}\right) + \sin\alpha \right) \right) \cdot |uw|.$$

Base case: T_{uw} has rank 1. Since the triangle is a smallest triangle such that u and w can see each other, w is the closest visible vertex to u in that cone. Hence the edge uw is part of the constrained $\theta_{(4k+x)}$-graph, and $\delta(u, w) = |uw|$. From the triangle inequality and the fact that $\boldsymbol{c} \geq 1$, we have $|uw| \leq \min\{|ua| + \boldsymbol{c} \cdot |aw|, |ub| + \boldsymbol{c} \cdot |bw|\}$, so the induction hypothesis holds.

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 $\theta_{(4k+x)}$-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 closest visible vertex to u in the subcone of u that contains w, and let a_0 and b_0 be the upper left and right corner of T_{uv_0} (see Figure 6.1). 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.

Since uw and uv_0 are visibility edges, by applying Lemma 2.1 to triangle v_0uw, a convex chain $v_0, ..., v_l = w$ of visibility edges connecting v_0 and w exists (see Figure 6.1). 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.

We now look at two consecutive vertices v_{j-1} and v_j along the convex chain. When $v_j \notin C_0^{v_{j-1}}$, let c and d be the left and right corners of $T_{v_{j-1}v_j}$. We distinguish four types of configurations: (i) $v_j \in C_i^{v_{j-1}}$ where $i > k$, or $i = k$ and $|cw| > |dw|$, (ii) $v_j \in C_i^{v_{j-1}}$ where $1 \leq i \leq k-1$, or $i = k$ and $|cw| \leq |dw|$, (iii) $v_j \in C_0^{v_{j-1}}$ and v_j

lies to the right of or has the same x-coordinate as v_{j-1}, (iv) $v_j \in C_0^{v_{j-1}}$ and v_j lies to the left of v_{j-1}. By convexity, the direction of $\overrightarrow{v_j v_{j+1}}$ is rotating counterclockwise for increasing j. Thus, these configurations occur in the order Type (i), Type (ii), Type (iii), Type (iv) along the convex chain from v_0 to w. We bound $\delta(v_{j-1}, v_j)$ as follows:

Type (i): $v_j \in C_i^{v_{j-1}}$ where $i > k$, or $i = k$ and $|cw| > |dw|$. Since v_j can see v_{j-1} and $T_{v_j v_{j-1}}$ has smaller area than T_{uw}, the induction hypothesis gives that $\delta(v_{j-1}, v_j)$ is at most $\max\{|v_{j-1}c| + \boldsymbol{c} \cdot |cv_j|, |v_{j-1}d| + \boldsymbol{c} \cdot |dv_j|\}$.

Let a_j is the intersection of the horizontal line through v_j and the left boundary of $C_0^{v_{j-1}}$. We aim to show that $\max\{|v_{j-1}c|+\boldsymbol{c}\cdot|cv_j|, |v_{j-1}d|+\boldsymbol{c}\cdot|dv_j|\} \leq |v_{j-1}a_j|+\boldsymbol{c}\cdot|a_j v_j|$. We use Lemma 2.5 to do this. However, since the precise application of this lemma depends on the family of θ-graphs and determines the value of $\boldsymbol{c}$, this case is discussed in the spanning proofs of the three families.

Type (ii): $v_j \in C_i^{v_{j-1}}$ where $1 \leq i \leq k-1$, or $i = k$ and $|cw| \leq |dw|$. Since v_j can see v_{j-1} and $T_{v_j v_{j-1}}$ has smaller area than T_{uw}, the induction hypothesis gives that $\delta(v_{j-1}, v_j)$ is at most $\max\{|v_{j-1}c| + \boldsymbol{c} \cdot |cv_j|, |v_{j-1}d| + \boldsymbol{c} \cdot |dv_j|\}$.

Let a_j be the intersection of the left boundary of $C_0^{v_{j-1}}$ and the horizontal line through v_j. Since $v_j \in C_i^{v_{j-1}}$ where $1 \leq i \leq k-1$, or $i = k$ and $|cw| \leq |dw|$, we can apply Lemma 2.4 in this case (where v, w, and a from Lemma 2.4 are v_{j-1}, v_j, and a_j) and we get that $\max\{|v_{j-1}c| + |cv_j|, |v_{j-1}d| + |dv_j|\} \leq |v_{j-1}a_j| + |a_j v_j|$ and $\max\{|cv_j|, |dv_j|\} \leq |a_j v_j|$. Since $\boldsymbol{c} \geq 1$, this implies that $\max\{|v_{j-1}c|+\boldsymbol{c}\cdot|cv_j|, |v_{j-1}d|+ \boldsymbol{c} \cdot |dv_j|\} \leq |v_{j-1}a_j| + \boldsymbol{c} \cdot |a_j v_j|$.

Type (iii): If $v_j \in C_0^{v_{j-1}}$ and v_j lies to the right of or has the same x-coordinate as v_{j-1}, let a_j and b_j be the left and right corner of $T_{v_{j-1}v_j}$. Since v_j can see v_{j-1} and $T_{v_{j-1}v_j}$ has smaller area than T_{uw}, we can apply the induction hypothesis. Thus, since v_j lies to the right of or has the same x-coordinate as v_{j-1}, $\delta(v_{j-1}, v_j)$ is at most $|v_{j-1}a_j| + \boldsymbol{c} \cdot |a_j v_j|$.

Type (iv): If $v_j \in C_0^{v_{j-1}}$ and v_j lies to the left of v_{j-1}, let a_j and b_j be the left and right corner of $T_{v_{j-1}v_j}$. Since v_j can see v_{j-1} and $T_{v_{j-1}v_j}$ has smaller area than T_{uw}, we can apply the induction hypothesis. Thus, since v_j lies to the left of v_{j-1}, $\delta(v_{j-1}, v_j)$ is at most $|v_{j-1}b_j| + \boldsymbol{c} \cdot |b_j v_j|$.

To complete the proof, we consider two cases: (a) $\angle awu \leq \pi/2$, (b) $\angle awu > \pi/2$.

Case (a): We need to prove that $\delta(u, w) \leq \max\{|ua| + |aw|, |ub| + |bw|\} = |ua| + |aw|$. We first show that the convex chain cannot contain any Type (iv) configurations: for Type (iv) configurations to occur, v_j needs to lie to the left of v_{j-1}. However, by construction, v_j lies on or to the right of the line through v_{j-1} and w. Hence, since $\angle awv_{j-1} < \angle awu \leq \pi/2$, v_j lies to the right of v_{j-1}. We can now

bound $\delta(u,w)$ by using these bounds:

$$\begin{aligned}\delta(u,w) &\leq |uv_0| + \sum_{j=1}^{l} \delta(v_{j-1}, v_j) \\ &\leq |ua_0| + |a_0v_0| + \sum_{j=1}^{l} (|v_{j-1}a_j| + \boldsymbol{c} \cdot |a_jv_j|) \\ &\leq |ua| + \boldsymbol{c} \cdot |aw|\end{aligned}$$

Case (b): If $\angle awu > \pi/2$, the convex chain can contain Type (iv) configurations. We need to prove that $\delta(u,w) \leq \max\{|ua|+|aw|, |ub|+|bw|\} = |ub|+|bw|$. Let $T_{v_{j'}v_{j'+1}}$ be the first Type (iv) 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 $T_{v_{j'}w}$. We now bound $\delta(u,w)$ as follows (see Figure 6.3):

$$\begin{aligned}\delta(u,w) &\leq |uv_0| + \sum_{j=1}^{l} \delta(v_{j-1}, v_j) \\ &\leq |ua_0| + |a_0v_0| + \sum_{j=1}^{j'} (|v_{j-1}a_j| + \boldsymbol{c} \cdot |a_jv_j|) + \sum_{j=j'+1}^{l} (|v_{j-1}b_j| + \boldsymbol{c} \cdot |b_jv_j|) \\ &\leq |ua'| + \boldsymbol{c} \cdot |a'v_{j'}| + |v_{j'}b''| + \boldsymbol{c} \cdot |b''w| \\ &\leq |ub'| + \boldsymbol{c} \cdot |b'v_{j'}| + |v_{j'}b''| + \boldsymbol{c} \cdot |b''w| \\ &= |ub| + \boldsymbol{c} \cdot |bw|\end{aligned}$$

□

6.3 The Constrained $\theta_{(4k+4)}$-Graph

In this section we complete the proof of Theorem 6.2 for the constrained $\theta_{(4k+4)}$-graph.

Theorem 6.3 *Let u and w be two vertices in the plane such that u can see w. Let m be the midpoint of the side of T_{uw} opposite u and let α be the unsigned angle between uw and um. There exists a path connecting u and w in the constrained $\theta_{(4k+4)}$-graph ($k \geq 1$) of length at most*

$$\left(\frac{\cos\alpha}{\cos\left(\frac{\theta}{2}\right)} + \frac{\cos\alpha \cdot \tan\left(\frac{\theta}{2}\right) + \sin\alpha}{\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{\theta}{2}\right)} \right) \cdot |uw|.$$

Proof. We apply Theorem 6.2 using $\boldsymbol{c} = 1/\left(\cos(\theta/2) - \sin(\theta/2)\right)$. The assumptions made in Theorem 6.2 still apply. Recall that c and d are the left and right corners of $T_{v_{j-1}v_j}$, opposite to v, and a_j is the intersection of the horizontal line through v_j and the left boundary of $C_0^{v_{j-1}}$. It remains to show that for the Type (i) configurations, we have that $\max\{|v_{j-1}c| + \boldsymbol{c} \cdot |cv_j|, |v_{j-1}d| + \boldsymbol{c} \cdot |dv_j|\} \leq |v_{j-1}a_j| + \boldsymbol{c} \cdot |a_j v_j|$. Let β be $\angle a_j v_j v_{j-1}$ and let γ be the angle between $v_j v_{j-1}$ and the bisector of $T_{v_{j-1}v_j}$.

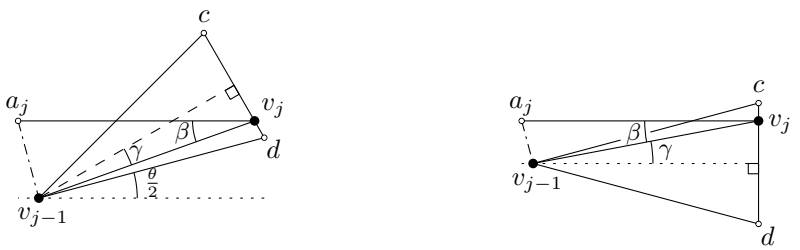

Figure 6.4: The remaining cases of the induction step for the $\theta_{(4k+4)}$-graph: (a) $v_j \in C_k^{v_{j-1}}$ and $|cw| > |dw|$, (b) $v_j \in C_{k+1}^{v_{j-1}}$

We distinguish two cases: (a) $v_j \in C_k^{v_{j-1}}$ and $|cw| > |dw|$, (b) $v_j \in C_{k+1}^{v_{j-1}}$.

Case (a): When $v_j \in C_k^{v_{j-1}}$ and $|cw| > |dw|$, the induction hypothesis for $T_{v_{j-1}v_j}$ gives $\delta(v_{j-1}, v_j) \leq |v_{j-1}c| + \boldsymbol{c} \cdot |cv_j|$ (see Figure 6.4a). We note that $\gamma = \theta - \beta$. Hence Lemma 2.5 gives that the inequality holds when $\boldsymbol{c} \geq (\cos(\theta - \beta) - \sin\beta)/(\cos(\theta/2 - \beta) - \sin(3\theta/2 - \beta))$. As this function is decreasing in β for $\theta/2 \leq \beta \leq \theta$, it is maximized when β equals $\theta/2$. Hence $\boldsymbol{c}$ needs to be at least $(\cos(\theta/2) - \sin(\theta/2))/(1 - \sin\theta)$, which can be rewritten to $1/(\cos(\theta/2) - \sin(\theta/2))$.

Case (b): When $v_j \in C_{k+1}^{v_{j-1}}$, v_j lies above the bisector of $T_{v_{j-1}v_j}$ and the induction hypothesis for $T_{v_{j-1}v_j}$ gives $\delta(v_{j-1}, v_j) \leq |v_j d| + \boldsymbol{c} \cdot |dv_{j-1}|$ (see Figure 6.4b). We note that $\gamma = \beta$. Hence Lemma 2.5 gives that the inequality holds when $\boldsymbol{c} \geq (\cos\beta - \sin\beta)/(\cos(\theta/2 - \beta) - \sin(\theta/2 + \beta))$, which is equal to $1/(\cos(\theta/2) - \sin(\theta/2))$. □

Since $\cos\alpha/\cos(\theta/2) + (\cos\alpha\tan(\theta/2) + \sin\alpha)/(\cos(\theta/2) - \sin(\theta/2))$ is increasing for $\alpha \in [0, \theta/2]$, for $\theta \leq \pi/4$, it is maximized when $\alpha = \theta/2$, and we obtain the following corollary:

Corollary 6.2 *The constrained $\theta_{(4k+4)}$-graph ($k \geq 1$) is a $\left(1 + \frac{2\cdot\sin\left(\frac{\theta}{2}\right)}{\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{\theta}{2}\right)}\right)$-spanner of $Vis(P, S)$.*

6.4 The Constrained $\theta_{(4k+3)}$-Graph and $\theta_{(4k+5)}$-Graph

In this section we complete the proof of Theorem 6.2 for the constrained $\theta_{(4k+3)}$-graph and $\theta_{(4k+5)}$-graph.

Theorem 6.4 *Let u and w be two vertices in the plane such that u can see w. Let m be the midpoint of the side of T_{uw} opposite u and let α be the unsigned angle between uw and um. There exists a path connecting u and w in the constrained $\theta_{(4k+3)}$-graph ($k \geq 1$) of length at most*

$$\left(\frac{\cos\alpha}{\cos\left(\frac{\theta}{2}\right)} + \frac{\left(\cos\alpha \cdot \tan\left(\frac{\theta}{2}\right) + \sin\alpha\right) \cdot \cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{3\theta}{4}\right)} \right) \cdot |uw|.$$

Proof. We apply Theorem 6.2 using $\boldsymbol{c} = \cos(\theta/4)/(\cos(\theta/2) - \sin(3\theta/4))$. The assumptions made in Theorem 6.2 still apply. Recall that c and d are the left and right corners of $T_{v_{j-1}v_j}$, opposite to v, and a_j is the intersection of the horizontal line through v_j and the left boundary of $C_0^{v_{j-1}}$. It remains to show that for the Type (i) configurations, we have that $\max\{|v_{j-1}c| + \boldsymbol{c} \cdot |cv_j|, |v_{j-1}d| + \boldsymbol{c} \cdot |dv_j|\} \leq |v_{j-1}a_j| + \boldsymbol{c} \cdot |a_jv_j|$. Let β be $\angle a_jv_jv_{j-1}$ and let γ be the angle between v_jv_{j-1} and the bisector of $T_{v_{j-1}v_j}$.

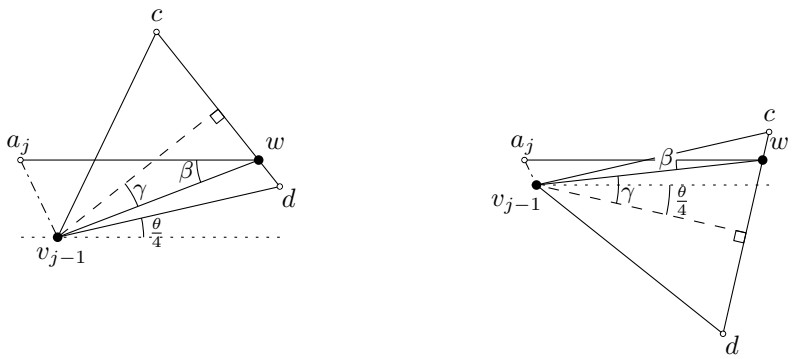

Figure 6.5: The remaining cases of the induction step for the $\theta_{(4k+3)}$-graph: (a) $v_j \in C_k^{v_{j-1}}$ and $|cw| > |dw|$, (b) $v_j \in C_{k+1}^{v_{j-1}}$

We distinguish two cases: (a) $v_j \in C_k^{v_{j-1}}$ and $|cw| > |dw|$, (b) $v_j \in C_{k+1}^{v_{j-1}}$.

Case (a): When $v_j \in C_k^{v_{j-1}}$ and $|cw| > |dw|$, the induction hypothesis for $T_{v_{j-1}v_j}$ gives $\delta(v_{j-1}, v_j) \leq |v_{j-1}c| + \boldsymbol{c} \cdot |cv_j|$ (see Figure 6.5a). We note that $\gamma = 3\theta/4 -$

β. Hence Lemma 2.5 gives that the inequality holds when $\boldsymbol{c} \geq (\cos(3\theta/4 - \beta) - \sin\beta)/(\cos(\theta/2 - \beta) - \sin(5\theta/4 - \beta))$. As this function is decreasing in β for $\theta/4 \leq \beta \leq 3\theta/4$, it is maximized when β equals $\theta/4$. Hence $\boldsymbol{c}$ needs to be at least $(\cos(\theta/2) - \sin(\theta/4))/(\cos(\theta/4) - \sin\theta)$, which is equal to $\cos(\theta/4)/(\cos(\theta/2) - \sin(3\theta/4))$.

Case (b): When $v_j \in C_{k+1}^{v_{j-1}}$, v_j lies above the bisector of $T_{v_{j-1}v_j}$ and the induction hypothesis for $T_{v_{j-1}v_j}$ gives $\delta(v_{j-1}, v_j) \leq |v_j d| + \boldsymbol{c} \cdot |dv_{j-1}|$ (see Figure 6.5b). We note that $\gamma = \theta/4 + \beta$. Hence Lemma 2.5 gives that the inequality holds when $\boldsymbol{c} \geq (\cos(\theta/4 + \beta) - \sin\beta)/(\cos(\theta/2 - \beta) - \sin(3\theta/4 + \beta))$, which is equal to $\cos(\theta/4)/(\cos(\theta/2) - \sin(3\theta/4))$. □

Theorem 6.5 *Let u and w be two vertices in the plane such that u can see w. Let m be the midpoint of the side of T_{uw} opposite u and let α be the unsigned angle between uw and um. There exists a path connecting u and w in the constrained $\theta_{(4k+5)}$-graph ($k \geq 1$) of length at most*

$$\left(\frac{\cos\alpha}{\cos\left(\frac{\theta}{2}\right)} + \frac{\left(\cos\alpha \cdot \tan\left(\frac{\theta}{2}\right) + \sin\alpha\right) \cdot \cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{3\theta}{4}\right)}\right) \cdot |uw|.$$

Proof. We apply Theorem 6.2 using $\boldsymbol{c} = \cos(\theta/4)/(\cos(\theta/2) - \sin(3\theta/4))$. The assumptions made in Theorem 6.2 still apply. Recall that c and d are the left and right corners of $T_{v_{j-1}v_j}$, opposite to v, and a_j is the intersection of the horizontal line through v_j and the left boundary of $C_0^{v_{j-1}}$. It remains to show that for the Type (i) configurations, we have that $\max\{|v_{j-1}c| + \boldsymbol{c} \cdot |cv_j|, |v_{j-1}d| + \boldsymbol{c} \cdot |dv_j|\} \leq |v_{j-1}a_j| + \boldsymbol{c} \cdot |a_j v_j|$. Let β be $\angle a_j v_j v_{j-1}$ and let γ be the angle between $v_j v_{j-1}$ and the bisector of $T_{v_{j-1}v_j}$.

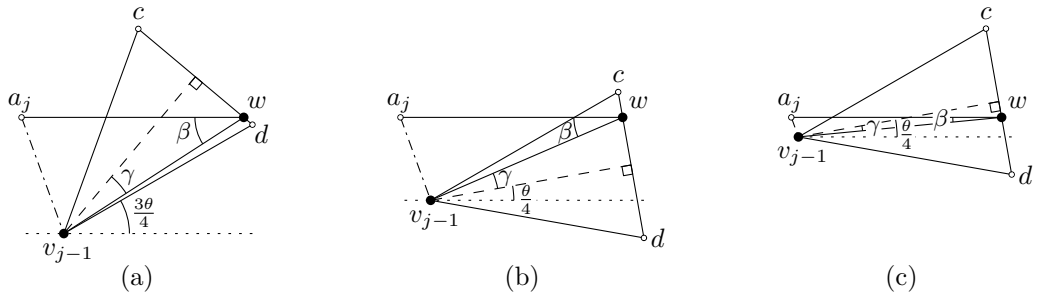

Figure 6.6: The remaining cases of the induction step for the $\theta_{(4k+5)}$-graph: (a) w lies in C_k^v and $|cw| > |dw|$, (b) w lies in C_{k+1}^v and $|cw| < |dw|$, (c) w lies in C_{k+1}^v and $|cw| \geq |dw|$

We distinguish two cases: (a) $v_j \in C_k^{v_{j-1}}$ and $|cw| > |dw|$, (b) $v_j \in C_{k+1}^{v_{j-1}}$.

Case (a): When $v_j \in C_k^{v_{j-1}}$ and $|cw| > |dw|$, the induction hypothesis for $T_{v_{j-1}v_j}$ gives $\delta(v_{j-1}, v_j) \leq |v_{j-1}c| + \boldsymbol{c} \cdot |cv_j|$ (see Figure 6.6a). We note that $\gamma = 5\theta/4 - \beta$. Hence Lemma 2.5 gives that the inequality holds when $\boldsymbol{c} \geq (\cos(5\theta/4 - \beta) - \sin\beta)/(\cos(\theta/2 - \beta) - \sin(7\theta/4 - \beta))$. As this function is decreasing in β for $3\theta/4 \leq \beta \leq 5\theta/4$, it is maximized when β equals $3\theta/4$. Hence $\boldsymbol{c}$ needs to be at least $(\cos(\theta/2) - \sin(3\theta/4))/(\cos(\theta/4) - \sin\theta)$, which is less than $\cos(\theta/4)/(\cos(\theta/2) - \sin(3\theta/4))$.

Case (b): When $v_j \in C_{k+1}^{v_{j-1}}$, the induction hypothesis for T_{vw} gives $\delta(v_{j-1}, v_j) \leq \max\{|v_{j-1}c| + \boldsymbol{c} \cdot |cv_j|, |v_{j-1}d| + \boldsymbol{c} \cdot |dv_j|\}$. If $|cw| < |dw|$ (see Figure 6.6b), we note that $\gamma = \beta - \theta/4$. Hence Lemma 2.5 gives that the inequality holds when $\boldsymbol{c} \geq (\cos(\beta - \theta/4) - \sin\beta)/(\cos(\theta/2 - \beta) - \sin(\theta/4 + \beta))$, which is equal to $\cos(\theta/4)/(\cos(\theta/2) - \sin(3\theta/4))$.

If $|dw| < |cw|$ (see Figure 6.6c), we note that $\gamma = \theta/4 - \beta$. Hence Lemma 2.5 gives that the inequality holds when $\boldsymbol{c} \geq (\cos(\theta/4 - \beta) - \sin\beta)/(\cos(\theta/2 - \beta) - \sin(3\theta/4 - \beta))$. As this function is decreasing in β for $0 \leq \beta \leq \theta/4$, it is maximized when β equals 0. Hence $\boldsymbol{c}$ needs to be at least $\cos(\theta/4)/(\cos(\theta/2) - \sin(3\theta/4))$. □

When looking at two vertices u and w in the constrained $\theta_{(4k+3)}$-graph and $\theta_{(4k+5)}$-graph, we notice that when the angle between uw and the bisector of T_{uw} is α, the angle between wu and the bisector of T_{wu} is $\theta/2 - \alpha$. Hence the worst case spanning ratio becomes the minimum of the spanning ratio when looking at T_{uw} and the spanning ratio when looking at T_{wu}.

Theorem 6.6 *The constrained $\theta_{(4k+3)}$-graph and $\theta_{(4k+5)}$-graph ($k \geq 1$) are $\frac{\cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{3\theta}{4}\right)}$-spanners of $Vis(P, S)$.*

Proof. The spanning ratio of the constrained $\theta_{(4k+3)}$-graph and $\theta_{(4k+5)}$-graph is at most:

$$\min \left\{ \begin{array}{l} \frac{\cos\alpha}{\cos\left(\frac{\theta}{2}\right)} + \frac{\left(\cos\alpha \cdot \tan\left(\frac{\theta}{2}\right) + \sin\alpha\right) \cdot \cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{3\theta}{4}\right)}, \\ \frac{\cos\left(\frac{\theta}{2} - \alpha\right)}{\cos\left(\frac{\theta}{2}\right)} + \frac{\left(\cos\left(\frac{\theta}{2} - \alpha\right) \cdot \tan\left(\frac{\theta}{2}\right) + \sin\left(\frac{\theta}{2} - \alpha\right)\right) \cdot \cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{3\theta}{4}\right)} \end{array} \right\}$$

Since $\cos\alpha / \cos\left(\frac{\theta}{2}\right) + \boldsymbol{c} \cdot \left(\cos\alpha \cdot \tan\left(\frac{\theta}{2}\right) + \sin\alpha\right)$ is increasing for $\alpha \in [0, \theta/2]$, for $\theta \leq 2\pi/7$, the minimum of these two functions is maximized when the two functions are equal, i.e. when $\alpha = \theta/4$. Thus the constrained $\theta_{(4k+3)}$-graph and $\theta_{(4k+5)}$-graph has spanning ratio at most:

$$\frac{\cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right)} + \frac{\left(\cos\left(\frac{\theta}{4}\right) \cdot \tan\left(\frac{\theta}{2}\right) + \sin\left(\frac{\theta}{4}\right)\right) \cdot \cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{3\theta}{4}\right)} = \frac{\cos\left(\frac{\theta}{4}\right) \cdot \cos\left(\frac{\theta}{2}\right)}{\cos\left(\frac{\theta}{2}\right) \cdot \left(\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{3\theta}{4}\right)\right)}. \quad \square$$

6.5 Conclusion

We showed that the constrained $\theta_{(4k+2)}$-graph has a tight spanning ratio of $1 + 2\sin(\theta/2)$. This is the first time tight spanning ratios have been found for a large family of constrained θ-graphs. Previously, the only constrained θ-graph for which tight bounds were known was the constrained θ_6-graph. We also gave improved upper bounds on the spanning ratio of the constrained $\theta_{(4k+3)}$-graph, the constrained $\theta_{(4k+4)}$-graph, and the constrained $\theta_{(4k+5)}$-graph.

There remain a number of open problems, such as finding tight spanning ratios for the constrained $\theta_{(4k+3)}$-graph, the constrained $\theta_{(4k+4)}$-graph, and the constrained $\theta_{(4k+5)}$-graph. Another set of open problems concerns constrained θ-graphs with few cones. In the unconstrained setting, it is known that θ_4 and θ_5-graphs are spanners, but this question remains unanswered in the constrained setting.

Finally, a major open problem involves routing on constrained θ-graphs. Where in the unconstrained setting a simple routing algorithm, called θ-routing, is known to exist, there is currently no equivalent in the constrained setting. In other words, even though we know that there exists a short path between any pair of visible vertices, there is no local routing algorithm to actually traverse such a path.

Chapter 7

Ordered Theta-Graphs

In this chapter, we present tight bounds on the spanning ratio of a large family of ordered θ-graphs. Given some ordering of the vertices, the ordered θ_m-graph is constructed as follows: we insert the vertices in the order given by the ordering. When a vertex u is inserted, for each cone C_i of u, we add an edge from u to the closest previously-inserted vertex in that cone, where distance is measured along the bisector of the cone. As was the case for unordered θ-graphs, our general position assumption implies that each vertex adds at most one edge per cone to the graph. As the ordered θ-graph depends on the ordering of the vertices, different orderings can produce different θ-graphs.

In this chapter, we show that for any integer $k \geq 1$, ordered θ-graphs with $4k+4$ cones have a tight spanning ratio of $1 + 2\sin(\theta/2)/(\cos(\theta/2) - \sin(\theta/2))$. We also show that for any integer $k \geq 2$, ordered θ-graphs with $4k+2$ cones have a tight spanning ratio of $1/(1 - 2\sin(\theta/2))$. We provide lower bounds for ordered θ-graphs with $4k+3$ and $4k+5$ cones. For ordered θ-graphs with $4k+2$ and $4k+5$ cones these lower bounds are strictly greater than the worst case spanning ratios of their unordered counterparts. These are the first results showing that ordered θ-graphs have worse spanning ratios than unordered θ-graphs. Finally, we show that, unlike their unordered counterparts, the ordered θ-graphs with 4, 5, and 6 cones are not spanners.

The results presented in this chapter are joint work with Prosenjit Bose and Pat Morin and have been accepted to the 25th International Symposium on Algorithm and Computation (ISAAC 2014) [26].

7.1 The Ordered $\theta_{(4k+4)}$-Graph

In this section, we give tight bounds on the spanning ratio of the ordered $\theta_{(4k+4)}$-graph, for any integer $k \geq 1$. We start by improving the current upper bound of $1/(1-2\sin(\theta/2))$ provided by Bose *et al.* [21].

Theorem 7.1 *Let u and w be two vertices in the plane such that w was inserted before u. Let m be the midpoint of the side of T_{uw} opposite u and let α be the unsigned angle between uw and um. There exists a path connecting u and w in the ordered $\theta_{(4k+4)}$-graph ($k \geq 1$) of length at most*

$$\left(\frac{\cos\alpha}{\cos\left(\frac{\theta}{2}\right)} + \boldsymbol{c}\cdot\left(\cos\alpha\cdot\tan\left(\frac{\theta}{2}\right)+\sin\alpha\right)\right)\cdot|uw|,$$

where $\boldsymbol{c}$ equals $1/(\cos(\theta/2)-\sin(\theta/2))$.

Proof. We assume without loss of generality that $w \in C_0^u$. We prove the theorem by induction on the rank, when ordered by area, of the canonical triangles T_{xy} for all pairs of vertices where y was inserted before x. Let a and b be the upper left and right corners of T_{uw}. Our inductive hypothesis is $\delta(u,w) \leq \max\{|ua|+\boldsymbol{c}\cdot|aw|, |ub|+\boldsymbol{c}\cdot|bw|\}$, where $\delta(u,w)$ denotes the length of the shortest path from u to w in the ordered $\theta_{(4k+4)}$-graph and $\boldsymbol{c}$ equals $1/(\cos(\theta/2)-\sin(\theta/2))$.

We first show that this induction hypothesis implies the theorem. Basic trigonometry gives us the following equalities: $|um| = |uw|\cdot\cos\alpha$, $|mw| = |uw|\cdot\sin\alpha$, $|am| = |bm| = |uw|\cdot\cos\alpha\cdot\tan(\theta/2)$, and $|ua| = |ub| = |uw|\cdot\cos\alpha/\cos(\theta/2)$. Thus the induction hypothesis gives that $\delta(u,w)$ is at most $|uw|\cdot(\cos\alpha/\cos(\theta/2) + \boldsymbol{c}\cdot(\cos\alpha\cdot\tan(\theta/2)+\sin\alpha))$.

Base case: T_{uw} has rank 1. Since this triangle is a smallest triangle where w was inserted before u, it is empty: if it is not empty, let x be a vertex in T_{uw}. Since T_{ux} and T_{xu} are both smaller than T_{uw}, the existence of x contradicts that T_{uw} is the smallest triangle where w was inserted before u. Since T_{uw} is empty, w is the closest vertex to u in C_0^u. Hence, since w was inserted before u, u adds an edge to w when it is inserted. Therefore, the edge uw is part of the ordered $\theta_{(4k+4)}$-graph, and $\delta(u,w) = |uw|$. From the triangle inequality and the fact that $\boldsymbol{c} \geq 1$, we have $|uw| \leq \max\{|ua|+\boldsymbol{c}\cdot|aw|, |ub|+\boldsymbol{c}\cdot|bw|\}$, so the induction hypothesis holds.

Induction step: We assume that the induction hypothesis holds for all pairs of vertices with canonical triangles of rank up to j. Let T_{uw} be a canonical triangle of rank $j+1$.

If uw is an edge in the ordered $\theta_{(4k+4)}$-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 be the vertex in T_{uw} that u connected to when it was inserted, let a' and b' be the upper left and right corners of T_{uv}, and let a'' be the intersection of the side of T_{uw} opposite u and the left boundary of C_0^v (see Figure 7.1).

We need to perform case distinction on whether w was inserted before or after v, to determine whether we can apply induction on T_{vw} or T_{wv}. Let c and d be the left and right corners of T_{vw} and let c' and d' be the left and right corner of T_{wv}. We note that since the ordered $\theta_{(4k+4)}$-graph has an even number of cones, $vcwc'$ and $vdwd'$ form two parallelograms. Thus, we have that $|vc| + \boldsymbol{c} \cdot |cw| = |wc'| + \boldsymbol{c} \cdot |c'v|$ and $|vd| + \boldsymbol{c} \cdot |dw| = |wd'| + \boldsymbol{c} \cdot |d'v|$. Hence, we can assume without loss of generality that the canonical triangle we need to look at is T_{vw}.

Without loss of generality, we assume that v lies to the left of or has the same x-coordinate as w. Since we need to show that $\delta(u, w) \leq \max\{|ua| + \boldsymbol{c} \cdot |aw|, |ub| + \boldsymbol{c} \cdot |bw|\}$, it suffices to show that $\delta(u, w) \leq |ua| + \boldsymbol{c} \cdot |aw|$. We perform a case analysis based on the cone of v that contains w: (a) $w \in C_0^v$, (b) $w \in C_i^v$ where $1 \leq i \leq k-1$, or $i = k$ and $|cw| \leq |dw|$, (c) $w \in C_k^v$ and $|cw| > |dw|$, (d) $w \in C_{k+1}^v$. To prove that $\delta(u, w) \leq |ua| + \boldsymbol{c} \cdot |aw|$, it suffices to show that $\delta(v, w) \leq |va''| + \boldsymbol{c} \cdot |a''w|$, as $|uv| \leq |ua'| + \boldsymbol{c} \cdot |a'v|$ and v, a'', a, and a' form a parallelogram (see Figure 7.1).

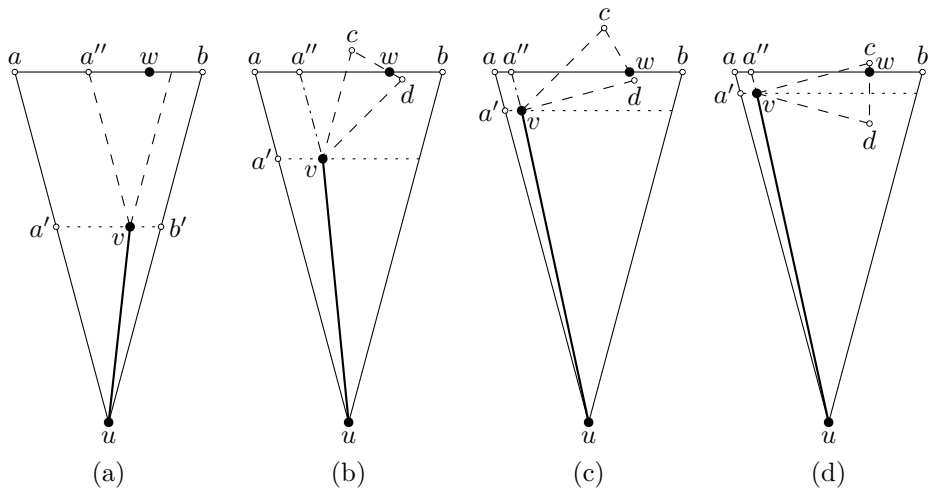

Figure 7.1: The four cases based on the cone of v that contains w

Case (a): Vertex w lies in C_0^v (see Figure 7.1a). Since T_{vw} has smaller area than T_{uw}, we apply the inductive hypothesis to T_{vw}. Since v lies to the left of or has the

same x-coordinate as w, we have $\delta(v,w) \leq |va''| + \boldsymbol{c} \cdot |a''w|$.

Case (b): Vertex w lies in C_i^v, where $1 \leq i \leq k-1$, or $i = k$ and $|cw| \leq |dw|$. Since T_{vw} is smaller than T_{uw}, by induction we have $\delta(v,w) \leq \max\{|vc|+\boldsymbol{c}\cdot|cw|, |vd|+\boldsymbol{c}\cdot|dw|\}$ (see Figure 7.1b). Since $w \in C_i^v$ where $1 \leq i \leq k-1$, or $i = k$ and $|cw| \leq |dw|$, we can apply Lemma 2.4. Note that point a in Lemma 2.4 corresponds to point a'' in this proof. Hence, we get that $\max\{|vc|+|cw|, |vd|+|dw|\} \leq |va''| + |a''w|$ and $\max\{|cw|, |dw|\} \leq |a''w|$. Since $\boldsymbol{c} \geq 1$, this implies that $\max\{|vc| + \boldsymbol{c} \cdot |cw|, |vd| + \boldsymbol{c} \cdot |dw|\} \leq |va''| + \boldsymbol{c} \cdot |a''w|$.

Case (c) Vertex w lies in C_k^v and $|cw| > |dw|$. Since T_{vw} is smaller than T_{uw} and $|cw| > |dw|$, the induction hypothesis for T_{vw} gives $\delta(v,w) \leq |vc| + \boldsymbol{c} \cdot |cw|$ (see Figure 7.1c). Let β be $\angle a''wv$ and let γ be the angle between vw and the bisector of T_{vw}. We note that $\gamma = \theta - \beta$. Hence Lemma 2.5 gives that $|vc|+\boldsymbol{c}\cdot|cw| \leq |va''|+\boldsymbol{c}\cdot|a''w|$ holds when $\boldsymbol{c} \geq (\cos(\theta-\beta) - \sin\beta)/(\cos(\theta/2-\beta) - \sin(3\theta/2-\beta))$. As this function is decreasing in β for $\theta/2 \leq \beta \leq \theta$, it is maximized when β equals $\theta/2$. Hence $\boldsymbol{c}$ needs to be at least $(\cos(\theta/2) - \sin(\theta/2))/(1-\sin\theta)$, which can be rewritten to $1/(\cos(\theta/2) - \sin(\theta/2))$.

Case (d) Vertex w lies in C_{k+1}^v (see Figure 7.1d). Since T_{vw} is smaller than T_{uw}, we can apply induction on it. Since w lies above the bisector of T_{vw}, the induction hypothesis for T_{vw} gives $\delta(v,w) \leq |vd| + \boldsymbol{c} \cdot |dw|$. Let β be $\angle a''wv$ and let γ be the angle between vw and the bisector of T_{vw}. We note that $\gamma = \beta$. Hence Lemma 2.5 gives that $|vd| + \boldsymbol{c} \cdot |dw| \leq |va''| + \boldsymbol{c} \cdot |a''w|$ holds when $\boldsymbol{c} \geq (\cos\beta - \sin\beta)/(\cos(\theta/2 - \beta) - \sin(\theta/2+\beta))$, which is equal to $1/(\cos(\theta/2) - \sin(\theta/2))$. □

Since $\cos\alpha/\cos(\theta/2) + (\cos\alpha \cdot \tan(\theta/2) + \sin\alpha)/(\cos(\theta/2) - \sin(\theta/2))$ is increasing for $\alpha \in [0, \theta/2]$, for $\theta \leq \pi/4$, it is maximized when $\alpha = \theta/2$, and we obtain the following corollary:

Corollary 7.1 *The ordered $\theta_{(4k+4)}$-graph ($k \geq 1$) is a $\left(1 + \frac{2\sin\left(\frac{\theta}{2}\right)}{\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{\theta}{2}\right)}\right)$-spanner.*

Next, we provide a matching lower bound on the spanning ratio of the ordered $\theta_{(4k+4)}$-graph.

Lemma 7.1 *The ordered $\theta_{(4k+4)}$-graph ($k \geq 1$) has spanning ratio at least $1 + \frac{2\sin\left(\frac{\theta}{2}\right)}{\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{\theta}{2}\right)}$.*

Proof. To prove the lower bound, we first construct a point set, after which we specify the order in which they are inserted into the graph. We place a vertex u and we place

a vertex w arbitrarily close to the right boundary of C_0^u. Next, we place a vertex v_1 arbitrarily close to the left corner of T_{uw}, followed by a vertex v_2 arbitrarily close to the upper corner of T_{wv_1}. Finally, we repeat the following two steps an arbitrary number of times: we place a vertex v_i arbitrarily close to the left corner of $T_{v_{i-2}v_{i-1}}$, followed by a vertex v_{i+1} arbitrarily close to the upper corner of $T_{v_{i-1}v_i}$. Let v_n be the last vertex placed in this fashion. We insert the vertices in the following order: v_n, v_{n-1}, ..., v_2, v_1, w, u. The resulting ordered $\theta_{(4k+4)}$-graph consists of a single path between u and w and is shown in Figure 7.2. Note that when a vertex v is inserted, all previously-inserted vertices lie in the same cone of v. This ensures that no shortcuts are introduced when inserting v.

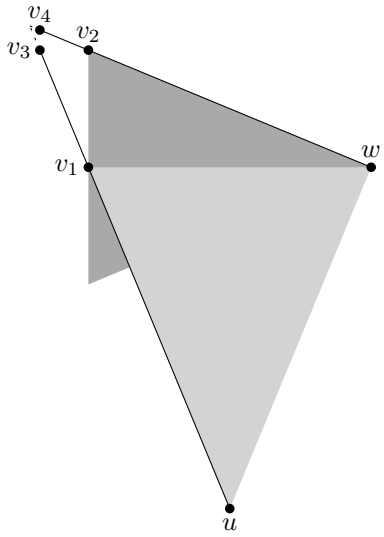

Figure 7.2: A lower bound for the ordered $\theta_{(4k+4)}$-graph

We note that edges uv_1 and edges of the form v_iv_{i+2} (for odd $i \geq 1$) lie on a line. We also note that edges wv_2 and edges of the form v_iv_{i+2} (for even $i \geq 2$) lie on a line. Let x be the intersection of these two lines and let β be $\angle xwv_1$. Hence, as the number of vertices approaches infinity, the total length of the path approaches $|ux| + |xw|$. Using that $\angle uxw = (\pi - \theta)/2 - \beta$, we compute the following edge lengths:

$$|ux| = |uw| \cdot \frac{\sin\left(\frac{\pi-\theta}{2} + \beta\right)}{\sin\left(\frac{\pi-\theta}{2} - \beta\right)}$$

$$|xw| = |uw| \cdot \frac{\sin\theta}{\sin\left(\frac{\pi-\theta}{2} - \beta\right)}$$

Since for the ordered $\theta_{(4k+4)}$-graph $\beta = \theta/2$, the sum of these equalities is $1/(\cos\theta + \tan\theta)$, which can be rewritten to $1 + 2\sin(\theta/2)/(\cos(\theta/2) - \sin(\theta/2))$. □

Theorem 7.2 *The ordered $\theta_{(4k+4)}$-graph ($k \geq 1$) has a tight spanning ratio of*

$$1 + \frac{2\sin\left(\frac{\theta}{2}\right)}{\cos\left(\frac{\theta}{2}\right) - \sin\left(\frac{\theta}{2}\right)}.$$

7.2 Lower Bounds

Next, we provide lower bounds for the ordered $\theta_{(4k+2)}$-graph, the ordered $\theta_{(4k+3)}$-graph, and the ordered $\theta_{(4k+5)}$-graph. For the ordered $\theta_{(4k+2)}$-graph, this lower bound implies that the current upper bound on the spanning ratio is tight. For the ordered $\theta_{(4k+2)}$-graph and the ordered $\theta_{(4k+5)}$-graph, these lower bounds are strictly larger than the upper bound on the worst case spanning ratio of its unordered counterpart.

Lemma 7.2 *The ordered $\theta_{(4k+2)}$-graph ($k \geq 2$) has spanning ratio at least $\frac{1}{1-2\sin(\frac{\theta}{2})}$.*

Proof. To prove the lower bound, we first construct a point set, after which we specify the order in which they are inserted into the graph. We place a vertex u, we place a vertex w arbitrarily close to the right boundary of C_0^u, and we place a vertex v_1 arbitrarily close to the left corner of T_{uw}. Next, we place the following configuration an arbitrary number of times: place a vertex l_1 in T_{v_1w} arbitrarily close to v_1, place a vertex v_2 in the right corner of T_{wv_1}, place a vertex l_2 close to the right boundary of $T_{v_1v_2}$ arbitrarily close to v_1, place a vertex r_1 in the intersection of $T_{v_2l_2}$ and $C_0^{l_2}$ arbitrarily close to v_2, place a vertex v_3 in the left corner of the intersection of $T_{l_1r_1}$ and $T_{l_2r_1}$, and place a vertex r_2 in the intersection of $T_{v_2v_3}$ and $T_{v_3v_2}$ such that v_3r_2 is parallel to v_1w. Since v_3r_2 is parallel to v_1w, we can repeat placing this configuration, constructing a staircase of vertices (see Figure 7.3). When we place the i-th configuration, we place vertices l_{2i-1}, v_{2i}, l_{2i}, r_{2i-1}, v_{2i+1}, and r_{2i}.

We insert these vertices into the ordered $\theta_{(4k+2)}$-graph in the following order: starting from the last configuration down to the first one, insert the vertices of the i-th configuration in the order r_{2i}, r_{2i-1}, v_{2i+1}, l_{2i}, l_{2i-1}, v_{2i}. Finally, we insert w, v_1, and u. The resulting ordered $\theta_{(4k+2)}$-graph is essentially a path between u and w and is shown in Figure 7.3.

We note that edges uv_1 and edges of the form v_iv_{i+2} (for odd $i \geq 1$) lie on a line. We also note that edges wv_2 and edges of the form v_iv_{i+2} (for even $i \geq 2$) lie on a

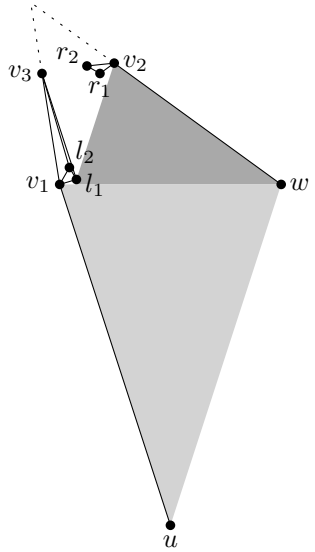

Figure 7.3: A lower bound for the ordered $\theta_{(4k+2)}$-graph

line. Let x be the intersection of these two lines. Hence, as the number of vertices approaches infinity, the total length of the path approaches $|ux| + |xw|$. Using that $\angle xuw = \theta$, $\angle xwu = (\pi + \theta)/2$, $\angle uxw = (\pi - 3\theta)/2$, and the law of sines, we compute the following edge lengths:

$$|ux| = |uw| \cdot \frac{\sin\left(\frac{\pi+\theta}{2}\right)}{\sin\left(\frac{\pi-3\theta}{2}\right)}$$

$$|xw| = |uw| \cdot \frac{\sin\theta}{\sin\left(\frac{\pi-3\theta}{2}\right)}$$

Hence, the spanning ratio of the ordered $\theta_{(4k+2)}$-graph is at least $(\sin((\pi + \theta)/2) + \sin\theta)/\sin((\pi - 3\theta)/2)$, which can be rewritten to $1/(1 - 2\sin(\theta/2))$. □

Since it is known that the $\theta_{(4k+2)}$-graph has a spanning ratio of at most $1/(1 - 2\sin(\theta/2))$ [21], this lower bound implies the following theorem.

Theorem 7.3 *The ordered $\theta_{(4k+2)}$-graph ($k \geq 2$) has a tight spanning ratio of $\frac{1}{1-2\sin(\frac{\theta}{2})}$.*

We also note that since the worst case spanning ratio of the unordered $\theta_{(4k+2)}$-graph is $1 + 2\sin(\theta/2)$ [15], this shows that the ordered $\theta_{(4k+2)}$-graph has a worse

worst case spanning ratio.

Lemma 7.3 *The ordered $\theta_{(4k+3)}$-graph ($k \geq 1$) has spanning ratio at least* $\frac{\cos\left(\frac{\theta}{4}\right)+\sin\theta}{\cos\left(\frac{3\theta}{4}\right)}$.

Proof. The proof is analogous to the proof of Lemma 7.1, where $\beta = \theta/4$, and shows that the spanning ratio of the ordered $\theta_{(4k+3)}$-graph is at least $(\sin(\pi/2 - \theta/4) + \sin\theta)/\sin(\pi/2 - 3\theta/4)$, which can be rewritten to $(\cos(\theta/4) + \sin\theta)/\cos(3\theta/4)$. □

Lemma 7.4 *The ordered $\theta_{(4k+5)}$-graph ($k \geq 1$) has spanning ratio at least* $1 + \frac{2\sin\left(\frac{\theta}{2}\right)\cdot\cos\left(\frac{\theta}{4}\right)}{\cos\left(\frac{\theta}{2}\right)-\sin\left(\frac{3\theta}{4}\right)}$.

Proof. The proof is analogous to the proof of Lemma 7.1, where $\beta = 3\theta/4$, and shows that the spanning ratio of the ordered $\theta_{(4k+5)}$-graph is at least $(\sin(\pi/2 + \theta/4) + \sin\theta)/\sin(\pi/2 - 5\theta/4)$, which can be rewritten to $1 + 2\sin(\theta/2)\cdot\cos(\theta/4)/(\cos(\theta/2) - \sin(3\theta/4))$. □

We note that this lower bound on the spanning ratio of the ordered $\theta_{(4k+5)}$-graph is the same as the current upper bound on θ-routing on the unordered $\theta_{(4k+5)}$-graph, which is strictly greater than the current upper bound on the spanning ratio of the unordered $\theta_{(4k+5)}$-graph.

7.3 Ordered Theta-Graphs with Few Cones

In this section we show that ordered θ-graphs with 3, 4, 5, or 6 cones are not spanners. For the ordered θ_4, θ_5, and θ_6-graph, this is surprising, since their unordered counterparts were recently show to be spanners [5, 6, 27].

For each of these ordered θ-graphs, we build a *tower* similar to the ones from the previous section. However, unlike the towers in the previous section, the towers of ordered θ-graphs that have at most 6 cones do not converge, thus giving rise to point sets where the spanning ratio depends on the size of these sets.

Lemma 7.5 *The ordered θ_4-graph is not a spanner.*

Proof. To prove that the ordered θ_4-graph is not a spanner, we first construct a point set, after which we specify the order in which they are inserted into the graph. We place a vertex u and we place a vertex w slightly to the right of the bisector of C_0^u.

Next, we place a vertex v_1 arbitrarily close to the left corner of T_{uw} and we place a vertex v_2 arbitrarily close to the upper corner of T_{wv_1}. Note that the placement of v_2 implies that it lies slightly to the right of the bisector of $C_0^{v_1}$. Because of this, we can repeat placing pairs of vertices in a similar fashion, constructing a staircase of vertices (see Figure 7.4). Let v_n denote the last vertex that was placed.

We insert these vertices into the ordered θ_4-graph in the following order: v_n, v_{n-1}, v_{n-2}, v_{n-3}, ..., v_2, v_1, w, u. The resulting ordered θ_4-graph consists of a single path between u and w and is shown in Figure 7.4.

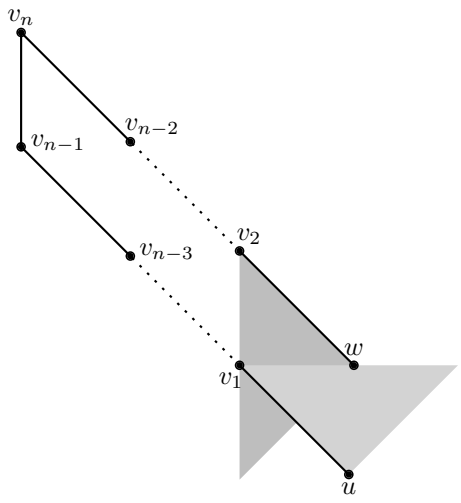

Figure 7.4: The ordered θ_4-graph is not a spanner

When we take $|uw|$ to be 1, all diagonal edges have length $\boldsymbol{c} = \sqrt{2}$ and the total length of the path is $1+n\cdot\sqrt{2}$. Hence, we have a graph whose spanning ratio depends on the number of vertices, implying that there does not exist a constant t, such that it is a t-spanner. □

Lemma 7.6 *The ordered θ_3-graph is not a spanner.*

Proof. The proof is analogous to the proof of Lemma 7.5, where $\boldsymbol{c} = \cos(\pi/6) = \sqrt{3}/2$, and shows that the total length of the path is $1 + n \cdot \sqrt{3}/2$. Hence, we have a graph whose spanning ratio depends on the number of vertices, implying that there does not exist a constant t, such that it is a t-spanner. □

Lemma 7.7 *The ordered θ_5-graph is not a spanner.*

Proof. The proof is analogous to the proof of Lemma 7.5, where vertex w is placed such that the angle between uw and the bisector of C_0^u is $\theta/4 = \pi/10$ and $\boldsymbol{c} = \cos(\pi/10)/\cos(\pi/5)$, and shows that the total length of the path is $1 + n \cdot \cos(\pi/10)/\cos(\pi/5)$. Hence, we have a graph whose spanning ratio depends on the number of vertices, implying that there does not exist a constant t, such that it is a t-spanner. We note that the placement of v_i (for even i) implies that the angle between $v_{i-1}v_i$ and the bisector of $C_0^{v_{i-1}}$ is $\theta/4$. Hence, every pair v_{i-1}, v_i of the staircase has the same relative configuration as the pair u, w. □

Lemma 7.8 *The ordered θ_6-graph is not a spanner.*

Proof. To prove that the ordered θ_6-graph is not a spanner, we first construct a point set, after which we specify the order in which they are inserted into the graph. We place a vertex u, we place a vertex w arbitrarily close to the right boundary of C_0^u, and we place a vertex v_1 arbitrarily close to the left corner of T_{uw}. Next, we place the following configuration an arbitrary number of times: place a vertex l_1 in T_{v_1w} arbitrarily close to v_1, place a vertex v_2 in the right corner of T_{wl_1}, place a vertex r_1 in $T_{v_2l_1}$ arbitrarily close to v_2, and place a vertex v_3 in the left corner of $T_{l_1r_1}$. Note that the line segment v_3r_1 is parallel to v_1w. Because of this, we can repeat placing

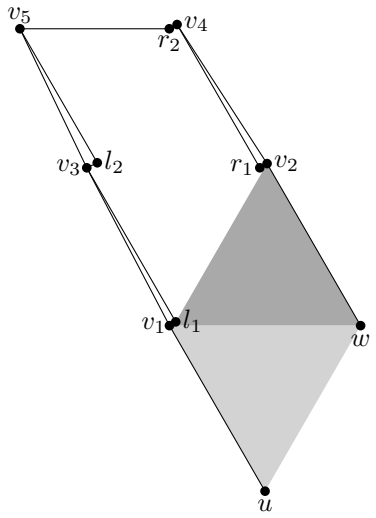

Figure 7.5: A lower bound for the ordered θ_6-graph

four vertices in a similar fashion, constructing a staircase of vertices (see Figure 7.5). When we place the i-th configuration, we place vertices l_i, v_{2i}, r_i, and v_{2i+1}.

We insert these vertices into the ordered θ_6-graph in the following order: starting from the last configuration down to the first one, insert the vertices of the i-th configuration in the order r_i, v_{2i+1}, l_i, v_{2i}. Finally, we insert w, v_1, and u. The resulting ordered θ_6-graph is essentially a path between u and w and is shown in Figure 7.5.

When we take $|uw|$ to be 1, we note that every configuration of four vertices extends the path length by 2. Hence, we have a graph whose spanning ratio depends on the number of vertices, implying that there does not exist a constant t, such that it is a t-spanner. □

7.4 Conclusion

We have provided tight spanning ratios for ordered θ-graphs with $4k + 2$ or $4k + 4$ cones. We also provided lower bounds for ordered θ-graphs with $4k + 3$ or $4k + 5$ cones. The lower bounds for ordered θ-graphs with $4k+2$ or $4k+5$ cones are strictly greater than those of their unordered counterparts. Furthermore, we showed that ordered θ-graphs with fewer than 7 cones are not spanners. For the ordered θ_4, θ_5, and θ_6-graph, this is surprising, since their unordered counterparts were show to be spanners [5, 6, 27]. Thus we have shown for the first time that the nice properties obtained when using ordered θ-graphs come at a price.

A number of open problems remain with respect to ordered θ-graphs. For starters, though we provided lower bounds for ordered θ-graphs with $4k + 3$ or $4k + 5$ cones, they do not match the current upper bound of $1/(1 - 2\sin(\theta/2))$. Hence, the obvious open problem is to find tight matching bounds for these graphs.

However, more importantly, there is currently no routing algorithm known for ordered θ-graphs. The θ-routing algorithm used for unordered θ-graphs cannot be used, since it assumes that when there exist vertices in a cone of the current vertex, there also exists an edge to a vertex in that cone. This assumption need not hold for ordered θ-graphs, since whether or not an edge is present depends on the order of insertion as well.

Chapter 8

Constrained Yao-Graphs

In this chapter, we present upper bounds on the spanning ratio of constrained Yao-graphs with at least seven cones. Given a set of points in the plane, a Yao-graph partitions the plane around each vertex into k disjoint cones, each having aperture $\theta = 2\pi/k$, and adds an edge to the closest vertex in each cone. Constrained Yao-graphs have the additional property that no edge properly intersects any of the given line segment constraints. We show that constrained Yao-graphs with an even number of cones have spanning ratio at most $1/(1-2\sin(\theta/2))$ and constrained Yao-graphs with an odd number of cones have spanning ratio at most $1/(1-2\sin(3\theta/8))$. These bounds match the current upper bounds in the unconstrained setting.

The results presented in this chapter appear in the 26th Canadian Conference on Computational Geometry (CCCG 2014) [57].

8.1 Preliminaries

As was the case with θ-graphs, 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. When constructing a (constrained) Y_k-graph, for each vertex u consider the rays originating from u with the angle between consecutive rays being $\theta = 2\pi/k$. Each pair of consecutive rays defines a cone. The cones are oriented such that the bisector of some cone coincides with the vertical halfline through u that lies above u. Let this cone be C_0 of u and number the cones in clockwise order around u (see Figure 8.1). The cones around the other vertices have the same orientation as the ones around u. We write C_i^u to indicate the i-th cone of a vertex u. For ease of exposition, we only consider point sets in general position: no two points lie on a line parallel to one of the rays

that define the cones and no three points are collinear.

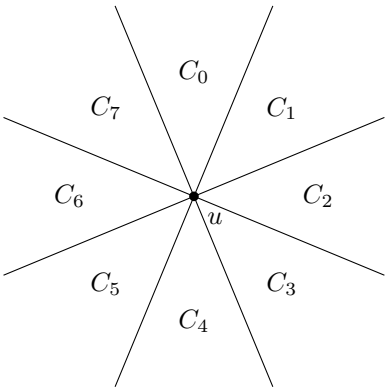

Figure 8.1: The cones having apex u in the Y_8-graph

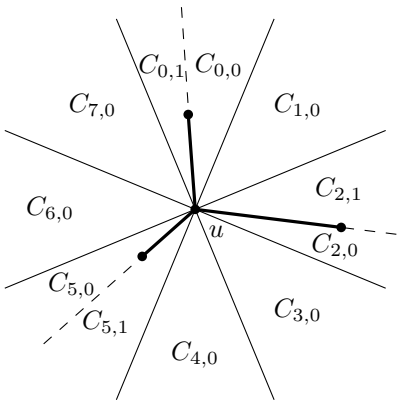

Figure 8.2: The subcones having apex u in the Y_8-graph. Constraints are shown as thick segments

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 *subcones*. We use $C_{i,j}^u$ to denote the j-th subcone of C_i^u (see Figure 8.2). 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 consider a cone that is not split to be a single subcone.

We now introduce the constrained Y_k-graph: for each subcone $C_{i,j}$ of each vertex u, add an edge from u to the closest vertex in that subcone that can see u (see Figure 8.3). When there exist multiple closest vertices in a subcone, we add an edge

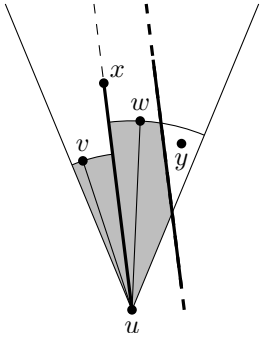

Figure 8.3: Vertex v is the closest visible vertex to u in the left subcone and w is the closest visible vertex to u in the right subcone, since y is not visible to u

to only a single one of them. 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 points $w \in C_{i,j}^u$ that can see u, $|uv| \leq |uw|$, where $|xy|$ denotes the length of the line segment between two points x and y and ties are broken arbitrarily.

8.2 Spanning Ratio

In this section, we prove that constrained Yao-graphs with at least 7 cones are spanners of the visibility graph.

Theorem 8.1 *The constrained Y_k-graph ($k \geq 7$) is a $1/\left(1 - 2\sin\left(\frac{\theta}{2}\right)\right)$-spanner of $Vis(P, S)$.*

Proof. Let u and w be two vertices that can see each other. We show that there exists a path connecting u and w in the constrained Y_k-graph ($k \geq 7$) of length at most $t \cdot |uw|$ for $t = 1/(1 - 2\sin(\theta/2))$, by induction on the distance between every pair of vertices u and w that can see each other. For ease of exposition, we assume without loss of generality that $w \in C_0^u$.

Base case: Vertices u and w are a closest visible pair. Since the closest visible pair need not be unique, we proceed to show that the subcone of C_0^u that contains w does not contain any vertices visible to u at distance at most $|uw|$: If there were such a vertex x, since ux and xw are visibility edges that lie in the same subcone, by Lemma 2.1 there exists a convex chain of visibility edges connecting x to w. Since we have at least 7 cones, the vertex adjacent to w along this chain is strictly closer to w than u, contradicting that $|uw|$ is a closest visible pair. Hence, since w is the closest visible vertex, uw is an edge in the constrained Y_k-graph and thus there exists a path between u and w of length $|uw| < t \cdot |uw|$.

Induction step: We assume that the induction hypothesis holds for all pairs of vertices that can see each other and whose distance is less than $|uw|$.

If uw is an edge in the constrained Y_k-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 be the closest visible vertex to u in the subcone of u that contains w, and let x be the point along uw such that $|uv| = |ux|$ (see Figure 8.4). Since x lies on uw, both ux and xw are visibility edges.

Next, we show that vx is also a visibility edge: If vx is not a visibility edge, that implies that it crosses some constraint. Since uv and ux are visibility edges, this constraint cannot cross them. Therefore, one endpoint of the constraint is contained

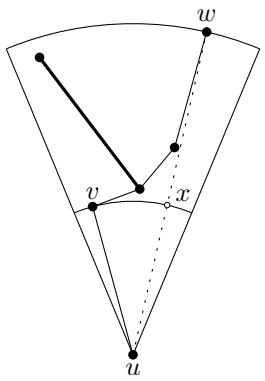

Figure 8.4: A convex chain from v to w

in triangle uvx. Let y be this endpoint. Since v and w lie in the same subcone of u, u is not the endpoint of a constraint intersecting the interior of uvx. Hence, we can apply Lemma 2.1 and obtain a convex chain of visibility edges from v and x and the polygon defined by uv, ux, and the convex chain is empty and does not contain any constraints. This implies that u can see every vertex along the convex chain, each of which is closer to it than v, contradicting that v was the closest visible vertex to u.

Since vx and xw are visibility edges, we can apply Lemma 2.1 to triangle vxw and we obtain a convex chain of visibility edges $v = p_0, ..., p_j = w$ connecting v and w (see Figure 8.4). Since we have at least 7 cones, the distance between any two consecutive vertices is strictly less than $|uw|$. Hence, since every pair of consecutive vertices along this convex chain can see each other, we can apply induction on each of them. Therefore, there exists a path from u to w via v of length at most

$$|uv| + t \cdot \sum_{i=0}^{j-1} |p_i p_{i+1}|.$$

Since the chain between v and w is contained in triangle vxw and the chain is convex, it follows that the total length of the chain is at most $|vx| + |xw|$. Thus, we can upper bound the length of the path by

$$|uv| + t \cdot (|vx| + |xw|)\,.$$

Since $|uv| = |ux|$, triangle uvx is an isosceles triangle and we can express $|vx|$ as $2\sin(\angle vux/2) \cdot |uv|$. Since this function is increasing for $\angle vux \in [0, 2\pi/7]$ and $\angle vux$ is at most θ, it follows that $|vx| \leq 2\sin(\theta/2) \cdot |uv|$. Next, we look at $|xw|$: Since x

lies on uw and $|uv| = |ux|$, it follows that $|xw| = |uw| - |ux| = |uw| - |uv|$. Hence, the path between u and w has length at most

$$\begin{aligned} & |uv| + t \cdot (|vx| + |xw|) \\ \leq\ & |uv| + t \cdot \left(2 \sin\left(\frac{\theta}{2}\right) \cdot |uv| + |uw| - |uv|\right) \\ =\ & t \cdot |uw| + \left(1 + 2 \sin\left(\frac{\theta}{2}\right) \cdot t - t\right) \cdot |uv|. \end{aligned}$$

Hence, for the length of the path to be at most $t \cdot |uw|$, we need that

$$1 + 2 \sin\left(\frac{\theta}{2}\right) \cdot t - t \leq 0,$$

which can be rewritten to

$$t \geq \frac{1}{1 - 2 \sin\left(\frac{\theta}{2}\right)},$$

completing the proof. □

For odd values of k, the spanning ratio can be decreased a bit: Let C_i^u be the cone of u that contains w and let C_j^w be the cone of w that contains u. When we look at two vertices u and w in the constrained Y_k-graph, we notice that when the angle between uw and the bisector of C_i^u is α, the angle between wu and the bisector of C_j^w is $\theta/2 - \alpha$ (see Figure 8.5). Hence, when bounding the worst case spanning ratio of constrained Y_k-graphs with an odd number of cones, we can assume without loss of generality that the angle between the bisector of the cone and uw is at most $\theta/4$.

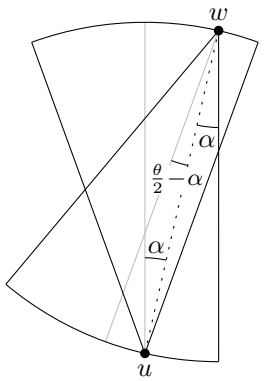

Figure 8.5: The angle between uw and the bisector of C_i^u is α and the angle between wu and the bisector of C_j^w is $\theta/2 - \alpha$

Theorem 8.2 *For odd $k \geq 7$, the constrained Y_k-graph is a $1/\left(1-2\sin\left(\frac{3\theta}{8}\right)\right)$-spanner of $Vis(P,S)$.*

Proof. Let u and w be two vertices that can see each other. We show that there exists a path connecting u and w in the constrained Y_k-graph ($k \geq 7$) of length at most $t \cdot |uw|$ for $t = 1/(1-2\sin(3\theta/8))$, by induction on the distance between every pair of vertices u and w that can see each other. For ease of exposition, we assume without loss of generality that $w \in C_0^u$. We also assume without loss of generality that the angle between the bisector of C_0^u and uw is at most $\theta/4$.

Base case: Vertices u and w are a closest visible pair. Using the same argument as in Theorem 8.1, it follows that uw is an edge of the constrained Y_k-graph and thus there exists a path between u and w of length $|uw| < t \cdot |uw|$.

Induction step: We assume that the induction hypothesis holds for all pairs of vertices that can see each other and whose distance is less than $|uw|$.

If uw is an edge in the constrained Y_k-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 be the closest visible vertex to u in the subcone of u that contains w, and let x be the point along uw such that $|uv| = |ux|$ (see Figure 8.6). Since x lies on uw, both ux and xw are visibility edges.

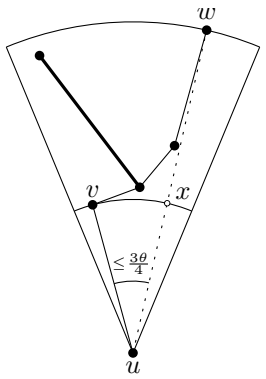

Figure 8.6: A convex chain from v to w

Using the same argument as in Theorem 8.1, it follows that vx is also a visibility edge. Hence, we can apply Lemma 2.1 to triangle vxw and we obtain a convex chain of visibility edges $v = p_0, ..., p_j = w$ connecting v and w (see Figure 8.6). Since we have at least 7 cones, the distance between any two consecutive vertices is strictly less than $|uw|$. Hence, since every pair of consecutive vertices along this convex chain

can see each other, we can apply induction on each of them. Therefore, there exists a path from u to w via v of length at most

$$|uv| + t \cdot \sum_{i=0}^{j-1} |p_i p_{i+1}|.$$

Analogous to Theorem 8.1, this expression can be upper bounded by

$$|uv| + t \cdot (|vx| + |xw|)\,.$$

Since $|uv| = |ux|$, triangle uvx is an isosceles triangle and we can express $|vx|$ as $2\sin(\angle vux/2) \cdot |uv|$. Since this function is increasing for $\angle vux \in [0, 2\pi/7]$ and $\angle vux$ is at most $3\theta/4$, it follows that $|vx| \leq 2\sin(3\theta/8) \cdot |uv|$. Analogous to Theorem 8.1, it holds that $|xw| = |uw| - |uv|$. Hence, the path between u and w has length at most

$$\begin{aligned} & |uv| + t \cdot (|vx| + |xw|) \\ \leq\ & |uv| + t \cdot \left(2\sin\left(\frac{3\theta}{8}\right) \cdot |uv| + |uw| - |uv|\right) \\ =\ & t \cdot |uw| + \left(1 + 2\sin\left(\frac{3\theta}{8}\right) \cdot t - t\right) \cdot |uv|. \end{aligned}$$

Hence, for the length of the path to be at most $t \cdot |uw|$, we need that

$$1 + 2\sin\left(\frac{3\theta}{8}\right) \cdot t - t \leq 0,$$

which can be rewritten to

$$t \geq \frac{1}{1 - 2\sin\left(\frac{3\theta}{8}\right)},$$

completing the proof. □

8.3 Conclusion

We showed that constrained Yao-graphs with at least 7 cones are spanners of the visibility graph and the upper bounds on the spanning ratio we obtained match those of the unconstrained Yao-graphs. This raises a number of new questions, the obvious one being whether we can reduce the upper bounds or find matching lower bound constructions.

Another set of open problems involves constrained Yao-graphs with at most 6 cones. In the unconstrained setting, it is known that the Y_k-graph is a spanner if and only if $k \geq 4$. Since the proof presented in this chapter can be applied only to Yao-graphs with at least 7 cones, it remains unknown whether this is also true in the constrained setting.

Finally, though we have upper bounds on the spanning ratio of constrained Yao-graphs with at least 7 cones, we do not have a local competitive routing algorithm to actually route messages between any two visible vertices. The main difficulty stems from the inductive steps along the convex chain, since these steps make it unclear where the routing algorithm should forward the message to. In particular, we cannot assume that there exists an edge in the subcone that contains the destination, since visibility may be blocked by a constraint. Hence, routing remains a major open problem in this area.

Chapter 9

Constrained Generalized Delaunay Graphs

In this chapter, we look at generalized Delaunay graphs in the constrained setting by introducing line segments which the edges of the graph are not allowed to cross. Given an arbitrary convex shape C (with non-zero area), an unconstrained Delaunay graph is constructed by adding an edge between two vertices p and q if and only if there exists a scaled translate of C with p and q on its boundary that does not contain any other vertices. We show that, regardless of the convex shape used to construct the constrained Delaunay graph G, it is a spanner whose spanning ratio depends solely on the properties of the empty convex shape C used to create it: We show that G satisfies the α_C-diamond property and the visible-pair κ_C-spanner property (defined in Section 9.2.3), which implies that it is a t-spanner for:

$$t = \begin{cases} 2\kappa_C \cdot \max\left(\frac{3}{\sin(\alpha_C/2)}, \kappa_C\right), & \textit{if } G \textit{ is a triangulation} \\ 2\kappa_C^2 \cdot \max\left(\frac{3}{\sin(\alpha_C/2)}, \kappa_C\right), & \textit{otherwise.} \end{cases}$$

To this end, we first provide two simple proofs showing that these graphs are plane and connected. Though the spanning proof is very general, since it holds for arbitrary convex shapes, its implied spanning ratio is far from tight. To improve on this, we also consider the special case where the empty convex shape C is a rectangle and show that it has spanning ratio at most $\sqrt{2} \cdot (2l/s + 1)$, where l and s are the length of the long and short side of C.

The results presented in this chapter are joint work with Prosenjit Bose and Jean-Lou De Carufel.

9.1 Preliminaries

Throughout this chapter, we fix a convex shape C. We assume without loss of generality that the origin lies in the interior of C. A *homothet* of C is obtained by scaling C with respect to the origin, followed by a translation. Thus, a homothet of C can be written as

$$x + \lambda C = \{x + \lambda z : z \in C\},$$

for some scaling factor $\lambda > 0$ and some point x in the interior of C after translation. We refer to x as the *center* of the homothet $x + \lambda C$. C is called *strictly convex* if every straight line intersects the boundary of C at most twice, i.e. C does not contain line segments as part of its boundary.

C defines a convex distance function d_C: given two visible vertices x and y, the distance from x to y is the smallest λ, such that $x + \lambda C$ contains y.

$$d_C(x, y) = \begin{cases} \min\{\lambda \geq 0 : y \in x + \lambda C\}, & \text{if } x \text{ can see } y, \\ \infty, & \text{otherwise.} \end{cases}$$

Note that in general, this distance function is not symmetric, i.e. $d_C(x, y) \neq d_C(y, x)$. If C is symmetric with respect to the origin, d_C is symmetric.

Using this convex distance function, we define the constrained Voronoi diagram. For each vertex p in the point set P, we define

$$V'_C(p) = \{x \in \mathbb{R}^2 : \text{for all } q \in P, d_C(x, p) \leq d_C(x, q)\}.$$

If C is not strictly convex, $V'_C(p)$ can consist of a closed region with an infinite halfline attached to it (see Figure 9.1a). Also, the intersection between two regions $V'_C(p)$ and $V'_C(q)$, with $p \neq q$, can have positive area. As a result, the collection of all these regions need not be a subdivision of the plane in which the interior of each cell is associated with exactly one vertex.

In order to obtain such a subdivision, we follow the approach described by Klein and Wood [48]. Infinite rays attached to regions of positive area are not considered to be part of that region. To assign a unique region to each point in the plane, we define a lexicographical ordering on the vertices of P. The Voronoi cell of the first vertex p_1 is now defined as

$$V_C(p_1) = cl(int(V'_C(p_1)))$$

and those of subsequent vertices p_i, with $1 < i \leq n$, we define

$$V_C(p_i) = cl\left(int\left(V'_C(p_i) \setminus \left(\bigcup_{j<i} V_C(p_j)\right)\right)\right),$$

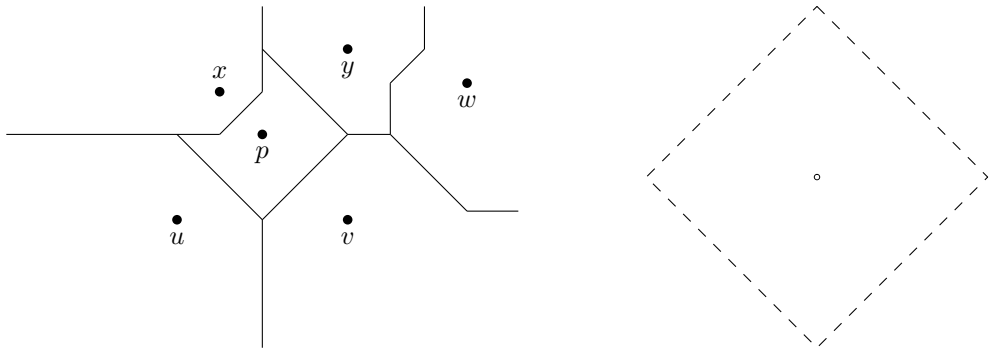

Figure 9.1: A closed region $V'_C(p)$ with an infinite half-line attached to it: (a) the infinite half-line on the boundary of $V'_C(u)$ and $V'_C(v)$ also lies on the boundary of $V'_C(p)$, (b) the convex shape used to construct the (constrained) Voronoi diagram

where $cl(X)$ and $int(X)$ denote the closure and interior of a set $X \subseteq \mathbb{R}^2$.

The constrained Voronoi diagram of a point set P with respect to a set of constraints S is defined as the collection of Voronoi cells $V_C(p)$, where p ranges over all vertices of P. The dual of the constrained Voronoi diagram has vertex set P and contains an edge between two visible vertices p and q if and only if the Voronoi cell of p and the Voronoi cell of q share a boundary. Note that this definition does not require that constraints are edges of the dual of the constrained Voronoi diagram.

Finally, for a given set of vertices P and a set of constraints S, we now define the constrained generalized Delaunay graph. Given any two visible vertices p and q, let $C(p, q)$ be any homothet of C with p and q on its boundary. The constrained generalized Delaunay graph contains an edge between p and q if and only if there exists a $C(p, q)$ such that $C(p, q)$ does not contain any vertices visible to both p and q. Note that this implies that constraints are *not* necessarily edges of the constrained generalized Delaunay graph. Joe and Wang showed that the constrained generalized Delaunay graph is not necessarily the dual of the constrained Voronoi diagram [43]. We later show that the dual of the constrained Voronoi diagram is a subgraph of the the constrained generalized Delaunay graph.

9.1.1 Auxiliary Lemmas

Next, we present two auxiliary lemmas that are needed to prove our main results. First, we reformulate a lemma that appears in several papers [33, 50, 51, 56].

Lemma 9.1 *Let C be a convex closed curve in the plane. The intersection of two distinct homothets of C is the union of two sets, each of which is either a segment, a singleton, or empty.*

Given two vertices p and q that can see each other and a convex polygon $C(p, q)$ with p and q on its boundary, we look at the constraints that have p as an endpoint and the edge(s) of $C(p, q)$ on which p lies and extend them to half-lines that have p as an endpoint (see Figure 9.2a). Given the cyclic order of these half-lines around p and the line segment pq, we define the clockwise neighbor of pq to be the half-line that minimizes the strictly positive clockwise angle with pq. Analogously, we define the counterclockwise neighbor of pq to be the half-line that minimizes the strictly positive counterclockwise angle with pq. We define the *cone* C_q^p that contains q to be the region between the clockwise and counterclockwise neighbor of p. Finally, let $C(p, q)_q^p$, the *region of $C(p, q)$ that contains q with respect to p*, be the intersection of $C(p, q)$ and C_q^p (see Figure 9.2b).

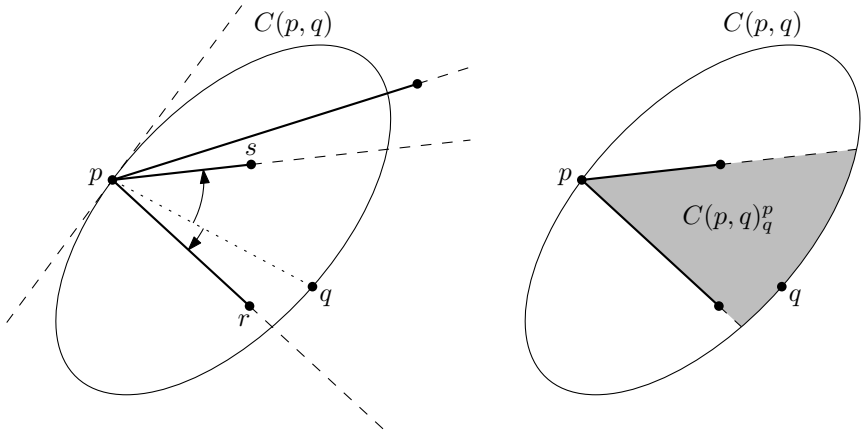

Figure 9.2: Defining the region of $C(p, q)$ that contains q with respect to p: (a) The clockwise and counterclockwise neighbor of pq are the half-lines through pr and ps, (b) $C(p, q)_q^p$ is marked in gray

Lemma 9.2 *Let p and q be two vertices that can see each other and let $C(p, q)$ be any convex polygon with p and q on its boundary. If $C(p, q)$ contains a vertex x in $C(p, q)_q^p$ that is visible to p, then $C(p, q)$ contains a vertex y that is visible to both p and q and triangle pyq is empty.*

Proof. We have two visibility edges, namely pq and px. Since x lies in $C(p,q)^p_q$, p is not the endpoint of a constraint such that q and x lie on opposite sides of the line through this constraint. Hence, we can apply Lemma 2.1 and we obtain a convex chain of visibility edges from x to q and the polygon defined by pq, px and the convex chain is empty and does not contain any constraints. Furthermore, since the convex chain is contained in pxq, which in turn is contained in $C(p,q)$, every vertex along the convex chain is contained in $C(p,q)$ (see Figure 9.3).

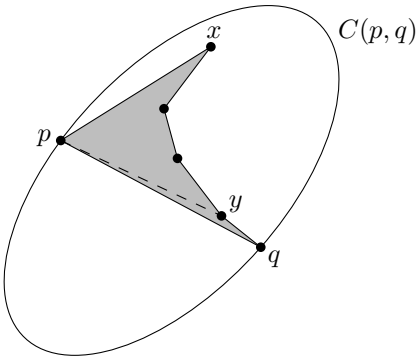

Figure 9.3: Vertex y lies in $C(p,q)$ and is visible to both p and q

Let y be the neighbor of q along this convex chain. Hence, y is visible to q and contained in $C(p,q)$. Furthermore, p can see y, since the line segment py is contained in the polygon defined by pq, px and the convex chain, which is empty and does not contain any constraints. This also implies that triangle pyq is empty. □

9.2 The Constrained Generalized Delaunay Graph

Before we show that every constrained generalized Delaunay graph is a spanner, we first show that they are plane and connected.

9.2.1 Planarity

In order to show that the constrained generalized Delaunay graph is plane, we first show that no edge of the graph can contain a vertex.

Lemma 9.3 *Let pq be an edge of the constrained generalized Delaunay graph. The line segment pq does not contain any vertices other than p and q.*

Proof. Since pq is an edge, p and q can see each other, which implies that they can see every point along pq as well. Furthermore, since $C(p,q)$ is convex, pq is contained in $C(p,q)$. Hence, if there were a vertex on line segment pq, it would lie in $C(p,q)$ and be visible to both p and q, contradicting that pq is an edge of the constrained generalized Delaunay graph. □

Lemma 9.4 *The constrained generalized Delaunay graph is plane.*

Proof. We prove this by contradiction, so assume that there exist two edges pq and rs that intersect properly, i.e. not at their endpoints. It follows from Lemma 9.3 that neither p nor q lies on rs and that neither r nor s lies on pq, so the edges intersect properly. Since pq is contained in $C(p,q)$ and rs is contained in $C(r,s)$, $C(p,q)$ and $C(r,s)$ intersect.

We first show that this implies that $p \in C(r,s)$, $q \in C(r,s)$, $r \in C(p,q)$, or $s \in C(p,q)$. If either $p \in C(r,s)$ or $q \in C(r,s)$, we are done, so assume that neither p nor q lies in $C(r,s)$. Lemma 9.1 states that $C(p,q)$ and $C(r,s)$ intersect each other at most twice. These intersections split the boundary of $C(p,q)$ into two parts: one that is contained in $C(r,s)$ and one that is not. Since $p \notin C(r,s)$ and $q \notin C(r,s)$, p and q lie on the arc of $C(p,q)$ that is not contained in $C(r,s)$ (see Figure 9.4). However, pq intersects $C(r,s)$, since otherwise pq cannot intersect rs. Let x and y be the two intersections of pq with the boundary of $C(r,s)$ (if the boundary of $C(r,s)$ is parallel to pq, x and y are the two endpoints of the interval of this intersection). We note that x and y split $C(r,s)$ into two parts, one of which is contained in $C(p,q)$, and that r and s cannot lie on the same part. In particular, one of r and s lies on the part that is contained in $C(p,q)$, proving that $r \in C(p,q)$, or $s \in C(p,q)$.

In the remainder of the proof, we assume without loss of generality that $r \in C(p,q)$ (see Figure 9.4). Let z be the intersection of pq and rs. Hence, z can see both p and r. Also, z is not the endpoint of a constraint intersecting the interior of triangle pzr. Therefore, it follows from Lemma 2.1 that there exists a convex chain of visibility edges from p to r. Let v be the neighbor of p along this convex chain. Since v is part of the convex chain, which is contained in pzr, which in turn is contained in $C(p,q)$, it follows that v is a vertex visible to p contained in $C(p,q)$. Furthermore, since the polygon defined by pz, zr and the convex chain does not contain any constraints, v lies in $C(p,q)^p_q$. Thus, it follows from Lemma 9.2 that there exists a vertex in $C(p,q)$ that is visible to both p and q, contradicting that pq is an edge of the constrained generalized Delaunay graph. □

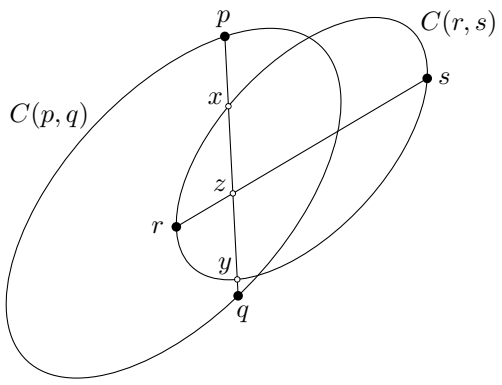

Figure 9.4: $C(p, q)$ and $C(r, s)$ intersect and p and q intersect $C(r, s)$ at x and y

9.2.2 Connectedness

In this section, we prove that the constrained generalized Delaunay graph is connected.

Lemma 9.5 *The constrained generalized Delaunay graph is connected.*

Proof. We first show that there exists a path between any two visible vertices. We show this by contradiction, so assume there exists a pair of vertices such that they are part of different components of the constrained generalized Delaunay graph. We call such a pair of vertices a *disconnected pair*. Let p and q be a disconnected pair and let $C(p, q)$ be the smallest homothet of C with p and q on its boundary. Let $C(p, q)$ be the minimum among all such smallest homothets among all disconnected pairs. Since p and q are disconnected, there cannot be an edge between them, hence there exists a vertex x in $C(p, q)$ that is visible to both p and q. Since p and q are not part of the same component, x is not part of at least one of these components. Assume without loss of generality that x is not part of the component that contains p. We look at the line segment pq and grow a homothet C' of $C(p, q)$ such that the relative area above and below pq is the same as that of $C(p, q)$. We grow C' until we hit x. Since x lies in the interior of $C(p, q)$, C' is contained in $C(p, q)$. However, this contradicts that $C(p, q)$ was the minimum such homothet, proving that there exists a path between any two visible vertices.

Since it is known that the visibility graph of line segments (the visibility graph where P consists of all endpoints of S and no other vertices) is connected, we can construct a path between any pair of (non-visible) vertices in the constrained generalized Delaunay graph by taking the path in the visibility graph and replacing each

visibility edge by a path of the constrained generalized Delaunay graph. Hence, the constrained generalized Delaunay graph is connected as well. □

9.2.3 Spanning Ratio

We use the dual of the constrained Voronoi diagram to show that the constrained generalized Delaunay graph is a spanner. However, since Joe and Wang showed that the constrained generalized Delaunay graph is not necessarily the dual of the constrained Voronoi diagram [43], we first prove that the dual of the constrained Voronoi diagram is a subgraph of the constrained generalized Delaunay graph. Joe and Wang prove a similar property, but since we do not require the constraints to be edges in either of these graphs, we cannot simply reuse their proof.

Lemma 9.6 *The dual of the constrained Voronoi diagram is a subgraph of the constrained generalized Delaunay graph.*

Proof. We prove this by contradiction, so let pq be an edge of the dual of the constrained Voronoi diagram that is not part of the constrained generalized Delaunay graph. Since pq is an edge of the dual of the constrained Voronoi diagram, the Voronoi cells of p and q share a boundary. Let y be a point on this shared boundary and let $C(p, q)$ be the convex shape centered at y that has p and q on its boundary. Since pq is not part of the constrained generalized Delaunay graph, there exists a vertex $x \in C(p, q)$ that is visible to both p and q. In case pq is a constraint, the remainder of the proof focuses on the part of $C(p, q)$ that contains y.

If px or qx intersect py or qy, let z be this intersection (see Figure 9.5a). Since z can see both x and y. Also, z is not the endpoint of a constraint intersecting the interior of triangle xyz, it follows from Lemma 2.1 that there exists a convex chain of visibility edges from x to y. Hence there exists a vertex v in xyz that is visible to y. Note that this vertex can be x. Since v is contained in xyz, which in turn is contained in $C(p, q)$, v is contained in $C(p, q)$. However, every vertex visible to y in $C(p, q)$ is closer to y than p and q, contradicting that y lies on the shared boundary of the Voronoi cells of p and q.

Next, we look at the case where px and qx do not intersect py and qy. By definition, p can see x and y. Hence, if p is not the endpoint of a constraint intersecting the interior of triangle pxy, it follows from Lemma 2.1 that there exists a convex chain of visibility edges from x to y (see Figure 9.5b). Since the neighbor of y along this

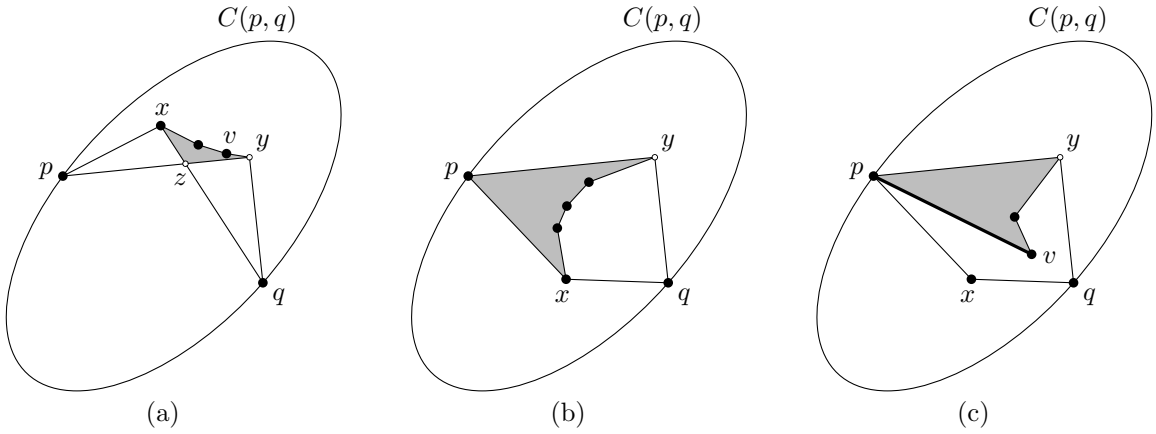

Figure 9.5: The three cases of Lemma 9.6: (a) px or qx intersect py or qy, (b) px and qx do not intersect py and qy and p is not the endpoint of a constraint intersecting the interior of triangle pxy, and (c) px and qx do not intersect py and qy and p is the endpoint of a constraint intersecting the interior of triangle pxy

chain is visible to y and contained in $C(p, q)$, this contradicts that y lies on the shared boundary of the Voronoi cells of p and q.

If p is the endpoint of a constraint intersecting the interior of triangle pxy, let v be the other endpoint of the constraint that minimizes the angle between pv and py (see Figure 9.5c). Hence, we have that p can see v and y and p is not the endpoint of a constraint intersecting the interior of triangle pvy. Hence, it follows from Lemma 2.1 that there exists a convex chain of visibility edges from x to y. Since the neighbor of y along this chain is visible to y and contained in $C(p, q)$, this contradicts that y lies on the shared boundary of the Voronoi cells of p and q. □

Let x and y be two distinct points on the boundary ∂C of C. These two points split ∂C into two parts. For each of these parts, there exists an isosceles triangle with base xy such that the third vertex lies on that part of ∂C. We denote the base angles of these two triangles by $\alpha_{x,y}$ and $\alpha'_{x,y}$. We define α_C as follows:

$$\alpha_C = \min\{\max(\alpha_{x,y}, \alpha'_{x,y}) : x, y \in \partial C, x \neq y\}.$$

Given a graph G and an angle $0 < \alpha < \pi/2$, we say that an edge pq of G satisfies the *α-diamond property*, when at least one of the two isosceles triangles with base pq and base angle α does not contain any vertex visible to both p and q. A graph G satisfies the α-diamond property when all of its edges satisfy this property [39].

Lemma 9.7 *The constrained generalized Delaunay graph satisfies the α_C-diamond property.*

Proof. Let pq be any edge of the constrained generalized Delaunay graph. Since pq is an edge, there exists a $C(p,q)$ such that $C(p,q)$ does not contain any vertices that are visible to both p and q. The vertices p and q split the boundary $\partial C(p,q)$ of $C(p,q)$ into two parts and each of these parts defines an isosceles triangle with base pq. Let β and γ be the base angles of these two isosceles triangles and assume without loss of generality that $\beta \geq \gamma$ (see Figure 9.6). Let x be the third vertex of the isosceles triangle having base angle β.

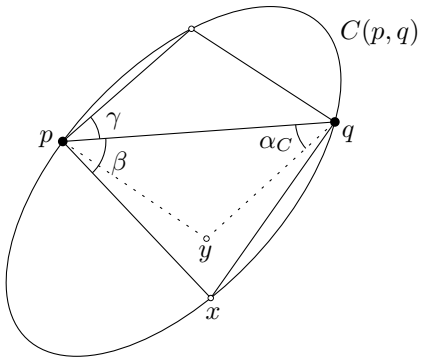

Figure 9.6: The constrained generalized Delaunay graph satisfies the α_C-diamond property

Translate and scale $C(p,q)$ such that it corresponds to C. This transformation does not affect the angles β and γ. Recall that $\alpha_C = \min\{\max(\alpha_{x,y}, \alpha'_{x,y}) : x, y \in \partial C, x \neq y\}$. Hence, since $p \neq q$ and both lie on the boundary of $C(p,q)$, the pair $\{\beta, \gamma\}$ is one of the pairs considered when determining α_C. Hence, since $\beta \geq \gamma$, it follows that $\alpha_C \leq \beta$. Let y be the third vertex of the isosceles triangle having base pq and base angle α_C that lies on the same side of pq as triangle pxq (see Figure 9.6). Since $\alpha_C \leq \beta$, triangle pyq is contained in triangle pxq. By convexity of $C(p,q)$, pxq is contained in $C(p,q)$. Hence, since $C(p,q)$ does not contain any vertices visible to both p and q, triangle pyq does not contain any vertices visible to both p and q either. Hence, pq satisfies the α_C-diamond property. □

For the next property, let O be a point in the interior of C and let x and y be two distinct points on the boundary ∂C of C, such that x, y, and O are collinear.

Again, x and y split ∂C into two parts. Let $\ell_{x,y}$ and $\ell'_{x,y}$ denote the length of these two parts. We define $\kappa_{C,O}$ as follows:

$$\kappa_{C,O} = \max\left\{\frac{\max(\ell_{x,y}, \ell'_{x,y})}{|xy|} : x, y \in \partial C, x \neq y, \text{ and } x, y, \text{ and } O \text{ are collinear}\right\}.$$

We note that the constrained generalized Delaunay graph does not depend on the location of O inside C, as the presence of any edge pq is defined in terms of $C(p, q)$, which does not depend on the location of O. Therefore, we define κ_C as follows:

$$\kappa_C = \min\{\kappa_{C,O} : O \text{ is in the interior of } C\}.$$

Throughout the rest of this section, we assume that O is picked such that $\kappa_C = \kappa_{C,O}$. Some examples of κ_C are the following: When C is a circle, $\kappa_C = \pi/2$ with O being the center of C, when C is a rectangle where l and s are the length of its long and short side, $\kappa_C = (l + s)/s = l/s + 1$ with O being the center of C, and when C is an equilateral triangle, $\kappa_C = \sqrt{3}$ with O being the center of C.

Given a constrained generalized Delaunay graph G, let p and q be two vertices on the boundary of a face f of the constrained generalized Delaunay graph, such that p can see q and the line segment pq does not intersect the exterior of f. If for every such pair p and q on every face f, there exists a path in G of length at most $\kappa \cdot |pq|$, then G satisfies the *visible-pair κ-spanner property*. We show that the constrained generalized Delaunay graph satisfies the visible-pair κ_C-spanner property. To do this, we look at the sequence of Voronoi cells intersected by the line segment pq.

Let p and q be two visible vertices, such that pq is not an edge of the constrained generalized Delaunay graph. Consider the sequence of Voronoi cells intersected by the line segment pq. If this line segment does not intersect any Voronoi vertices (points equidistant to at least three vertices), the vertices corresponding to the cells along this sequence form a path from p to q. However, since these Voronoi cells need not be convex, this path can contain the same cell multiple times (see Figure 9.7). To avoid this, we modify the path by 'skipping ahead' when this happens (a more precise definition follows after the next paragraph).

In the remainder of this section, we assume that pq is horizontal, with p to the left of q. This implies a left-to-right order of the intersected Voronoi cells. If pq contains a Voronoi vertex, we imagine moving pq up by an infinitesimal amount, such that this is no longer the case. Therefore, we can assume that pq does not contain any Voronoi vertices.

We are now ready to redefine the path from p to q. The first vertex p_0 of this path is p. We define x_1 to be the point on pq, such that $x_1 \in V_C(p_0)$ and x_1 is closest

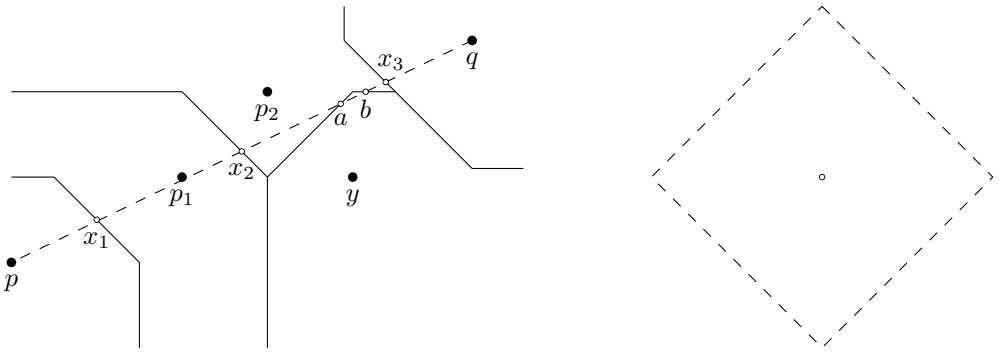

Figure 9.7: A (constrained) Voronoi diagram: (a) cells are not necessarily convex and pq intersects the boundary of $V_C(p_2)$ at x_2, a, b, and x_3, (b) the convex shape used to construct the (constrained) Voronoi diagram

to q. Next, for $i \geq 1$, assume that vertices $p_0, ..., p_{i-1}$ and points $x_1, ..., x_i$ have been defined, where x_i is the point on pq, such that $x_i \in V_C(p_{i-1})$ and x_i is closest to q. If $p_{i-1} = q$, the construction is completed. Otherwise, note that x_i lies on the boundary of two Voronoi cells. We define p_i to be the vertex whose Voronoi cell has x_i on its boundary (and $p_i \neq p_{i-1}$). We also define x_{i+1} to be the point on pq, such that $x_{i+1} \in V_C(p_i)$ and x_{i+1} is closest to q (see Figure 9.7). Note that since $x_i \in V_C(p_{i-1})$ is the point on pq that is closest to q, no vertex p_{i-1} is visited multiple times.

Let $p = p_0, p_1, ..., p_k = q$ be the sequence of vertices constructed in this way. Since each pair of consecutive vertices share a boundary of their Voronoi cells, the dual of the constrained Voronoi diagram contains an edge between them. It follows from Lemma 9.6 that these edges are also part of the constrained generalized Delaunay graph. Thus, $p = p_0, p_1, ..., p_k = q$ defines a path from p to q in the constrained generalized Delaunay graph. We refer to this path as the *direct path*. If all vertices along this part are above pq or all vertices are below pq, we call the direct path *one-sided*. We now repeat a lemma proven by Bose *et al.* [12].

Lemma 9.8 *Let p and q be two vertices on the boundary of a face f of the constrained generalized Delaunay graph, such that p can see q and the line segment pq does not intersect the exterior of f. If the direct path between p and q is one-sided, then its length is at most $\kappa_C \cdot |pq|$.*

Formally, Lemma 9.8 has been proven only for the unconstrained setting. However, due to the structure of this proof, which is based on upper bounding the lengths of

the edges of the direct path by partial boundaries of C, the proof applies to the constrained setting as well. Thus, proving that the direct path between p and q is one-sided, implies that the constrained generalized Delaunay graph satisfies the visible-pair κ_C-spanner property.

Lemma 9.9 *Let p and q be two vertices on the boundary of a face f of the constrained generalized Delaunay graph, such that p can see q and the line segment pq does not intersect the exterior of f. The direct path between p and q is one-sided.*

Proof. Recall that p and q are part of the same face f, p can see q, and the line segment pq does not intersect the exterior of f. Furthermore, since the constrained generalized Delaunay graph is plane (Lemma 9.4), no edge outside of f can cross the line segment pq. Hence, if the direct path is not one-sided, the path crosses the line through p and q either to the left of p or to the right of q. We show that the path cannot cross the line through p and q to the right of q. An analogous argument shows that the path cannot cross the line through p and q to the left of p.

We prove this by contradiction, so let p_ip_{i+1} be an edge of the direct path such that p_ip_{i+1} crosses the line through p and q to the right of q (see Figure 9.8). Let x_i be the point on pq on the boundary of the Voronoi cells of p_i and p_{i+1}. Let $C(p_i, p_{i+1})$ be the convex shape centered at x_i that has p_i and p_{i+1} on its boundary. Since x_i lies on the boundary of the Voronoi cells of p_i and p_{i+1}, $C(p_i, p_{i+1})$ cannot contain any vertices visible to x_i, since this vertex would be closer to x_i than p_i and p_{i+1}.

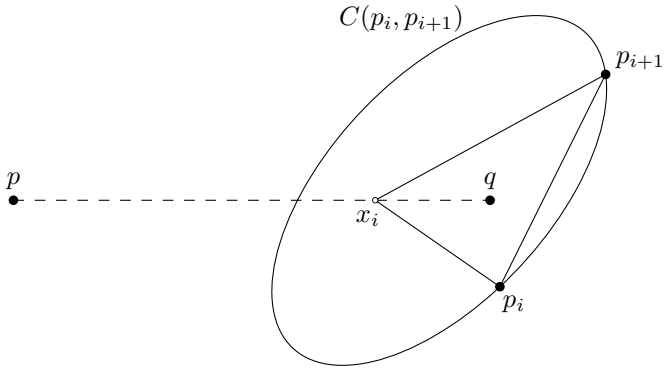

Figure 9.8: If the direct path is not one-sided, $C(p_i, p_{i+1})$ contains vertex q

However, since x_i lies on the line segment pq and the edge p_ip_{i+1} crosses the line through p_i and p_{i+1} to the right of q, q is contained in triangle $x_ip_ip_{i+1}$. Furthermore, since $C(p_i, p_{i+1})$ is convex, this triangle is contained in $C(p_i, p_{i+1})$. Finally, since x_i

on pq and p can see q, x_i can see q. Hence, $C(p_i, p_{i+1})$ contains a vertex visible to x_i that is closer than p_i and p_{i+1}, contradicting that x_i lies on the boundary of their Voronoi cells. □

Corollary 9.1 *The constrained generalized Delaunay graph satisfies the visible-pair κ_C-spanner property.*

We now have everything in place to prove that the constrained generalized Delaunay graph is a spanner. Das and Joseph [39] showed that any plane graph that satisfies the diamond property and the good polygon property (similar to the visible-pair κ-spanner property) is a spanner. Subsequently, Bose *et al.* [24] improved slightly on the spanning ratio. They showed that a geometric (constrained) graph G is a spanner of the visibility graph when it satisfies the following properties:

1. G is plane.
2. G satisfies the α-diamond property.
3. The spanning ratio of any one-sided path in G is at most κ.
4. G satisfies the visible-pair κ'-spanner property.

In particular, G is a t-spanner for

$$t = 2\kappa\kappa' \cdot \max\left(\frac{3}{\sin(\alpha/2)}, \kappa\right).$$

It follows from Lemmas 9.4, 9.7, and 9.8, and Corollary 9.1 that the constrained generalized Delaunay graph satisfies these four properties. Moreover, even though in general the constrained generalized Delaunay graph is not a triangulation, if for a specific convex shape it is, it satisfies the visible-pair 1-spanner property. Therefore, we obtain the following theorem:

Theorem 9.1 *The constrained generalized Delaunay graph G is a t-spanner for*

$$t = \begin{cases} 2\kappa_C \cdot \max\left(\frac{3}{\sin(\alpha_C/2)}, \kappa_C\right), & \text{if } G \text{ is a triangulation} \\ 2\kappa_C^2 \cdot \max\left(\frac{3}{\sin(\alpha_C/2)}, \kappa_C\right), & \text{otherwise.} \end{cases}$$

Though this theorem holds for all convex shapes, the bound it provides is rather loose when we look at a specific shape. For example, for the constrained Delaunay graph that uses an equilateral triangle, i.e. the constrained half-θ_6-graph, the above theorem implies an upper bound of $2 \cdot \sqrt{3}^2 \cdot 3/\sin(\pi/12) \approx 17.6$, which is far greater than the tight bound of 2 obtained from Theorem 3.1.

9.3 The Constrained Empty-Rectangle Delaunay Graph

In this section, we look at the special case where the empty convex shape is an arbitrary rectangle to improve on the spanning ratio shown in the previous section. We assume without loss of generality that the rectangle is axis-aligned. We do not, however, assume anything about the ratio between the height and width of the rectangle. We use p_x and p_y denote the x- and y-coordinate of a point p. We first show that if two visible vertices cannot see any vertices in $C(p,q)$ on one side of pq, then no vertex in $C(p,q)$ on the opposite side of pq can see any vertices beyond pq either.

Lemma 9.10 *Let p and q be two vertices that can see each other, such that pq is not vertical, and let $C(p,q)$ be any convex polygon with p and q on its boundary. If the region of $C(p,q)$ below pq does not contain any vertices visible to p and q, then no point x in $C(p,q)$ above pq can see any vertices in $C(p,q)$ below pq.*

Proof. We prove the lemma by contradiction, so assume that there exists a vertex y in $C(p,q)$ below pq that is visible to x, but not to p and q. Since $C(p,q)$ is a convex polygon and x and y lie on opposite sides of pq, the visibility edge xy intersects pq. Let z be this intersection (see Figure 9.9).

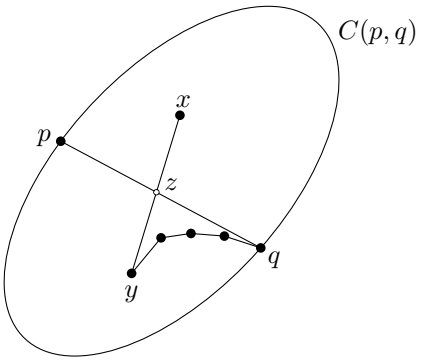

Figure 9.9: If x can see a vertex below pq, then so can q

Hence, zy and zq are visibility edges. Since z is not a vertex, it is not the endpoint of any constraints intersecting the interior of triangle yzq. It follows from Lemma 2.1 that there exists a convex chain of visibility edges between y and q and this chain is contained in yzq. However, this implies that the neighbor of q along this chain is visible to q and lies in $C(p,q)$ below pq, contradicting that this region does not

contain any vertices visible to p and q. □

Next, we introduce some notation for the following lemma. Let p and q be two vertices of the constrained generalized Delaunay graph that can see each other. Let $C(p, q)$ be a rectangle with p and q on its West and East boundary and let a and b be the Northwest and Northeast corner of $C(p, q)$. Let $m_1, ..., m_{k-1}$ be any $k - 1$ points on pq in the order they are visited when walking from p to q (see Figure 9.10). Let $m_0 = p$ and $m_k = q$. Consider the homothets S_i of $C(p, q)$ with m_i and m_{i+1} on their respective boundaries, for $0 \leq i < k$, such that the relative area above and below pq is the same as that of $C(p, q)$, and let a_i and b_i be the Northwest and Northeast corner of S_i.

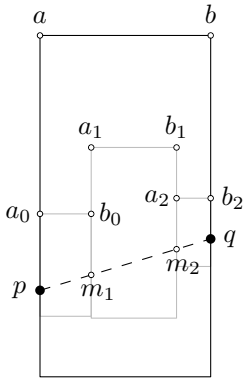

Figure 9.10: The total length of the sides of the rectangles S_i equals that of $C(p, q)$

Lemma 9.11 *We have*

$$\sum_{i=0}^{k-1} \big(|m_i a_i| + |a_i b_i| + |b_i m_{i+1}|\big) = |pa| + |ab| + |bq|.$$

Proof. Let $c = (|pa| + |ab| + |bq|)/|pq|$. Since for every S_i we have that the relative area above and below pq is the same as that of $C(p, q)$, we have $(|m_i a_i| + |a_i b_i| + |b_i m_{i+1}|)/|m_i m_{i+1}| = c$, for $0 \leq i < k$. Hence, we get

$$\begin{aligned}\sum_{i=0}^{k-1} \big(|m_i a_i| + |a_i b_i| + |b_i m_{i+1}|\big) &= \sum_{i=0}^{k-1} \big(c \cdot |m_i m_{i+1}|\big) \\ &= c \cdot |pq| \\ &= |pa| + |ab| + |bq|,\end{aligned}$$

proving the lemma. □

We first bound the length of the spanning path for the case where the rectangle is partially empty. We call a rectangle $C(p,q)$ *half-empty* when $C(p,q)$ contains no vertices in $C(p,q)^p_q$ below pq that are visible to p and $C(p,q)$ contains no vertices in $C(p,q)^q_p$ below pq that are visible to q.

Lemma 9.12 *Let p and q be two vertices that can see each other. Let $C(p,q)$ be a rectangle with p and q on its boundary, such that it is half-empty. Let a and b be the corners of $C(p,q)$ on the non-half-empty side. The constrained generalized Delaunay graph contains a path between p and q of length at most $|pa| + |ab| + |bq|$.*

Proof. We prove the lemma by induction on the size of $C(x,y)$, for any two visible vertices x and y, such that $C(x,y)$ is half-empty. We assume without loss of generality that p lies on the West boundary, q lies on the East boundary and that $C(p,q)$ is half-empty below pq. This implies that a and b are the Northwest and Northeast corner of $C(p,q)$. We also assume without loss of generality that the slope of pq is non-negative, i.e. $p_x < q_x$ and $p_y \leq q_y$ (see Figure 9.11).

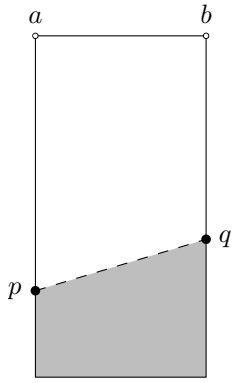

Figure 9.11: A rectangle with p on its West boundary and q on its East boundary

Base case: If $C(p,q)$ is a rectangle of smallest area, then $C(p,q)$ does not contain any vertices visible to both p and q: Assume this is not the case and grow a rectangle R (a homothet of $C(p,q)$ such that the relative area above and below pq is the same as that of $C(p,q)$) from p to q. Let x be the first vertex hit by R that is visible to p and lies in $C(p,q)^p_q$. Note that this implies that R is contained in $C(p,q)$. Therefore, R is smaller than $C(p,q)$. Furthermore, R is half-empty: By Lemma 9.10, the part

below the line through p and q does not contain any vertices visible to p or x in $C(p,q)^p_q$, and the part between the line through p and x and the line through p and q does not contain any vertices visible to p or x since x is the first visible vertex hit while growing R. However, this contradicts that $C(p,q)$ is the smallest half-empty rectangle.

Hence, $C(p,q)$ does not contain any vertices visible to both p and q, which implies that pq is an edge of the constrained generalized Delaunay graph. By the the length of the shortest path from p to q is at most $|pq| \leq |pa| + |ab| + |bq|$.

Induction step: We assume that for all half-empty rectangles $C(x,y)$ smaller than $C(p,q)$ the lemma holds. If pq is an edge of the constrained generalized Delaunay graph, by the the length of the shortest path from p to q is at most $|pq| \leq |pa| + |ab| + |bq|$.

If there is no edge between p and q, there exists a vertex in $C(p,q)$ that is visible from both p and q. We grow a rectangle R (a homothet of $C(p,q)$, such that the relative area above and below pq is the same as that of $C(p,q)$) from p to q. Let x be the first vertex hit by R that is visible to p and lies in $C(p,q)^p_q$ and let a' and b' be the Northwest and Northeast corner of R (see Figure 9.12). Note that this implies that R is contained in $C(p,q)$. We note that px is not necessarily an edge in the constrained generalized Delaunay graph, since if it is a constraint, there can be vertices visible to both p and x above px. However, since R is half-empty and smaller than $C(p,q)$, we can apply induction on it and we obtain that the path from p to x has length at most $|pa'| + |a'b'| + |b'x|$ when x lies on the East boundary of R, and that the path from p to x has length at most $|pa'| + |a'x|$ when x lies on the North boundary of R.

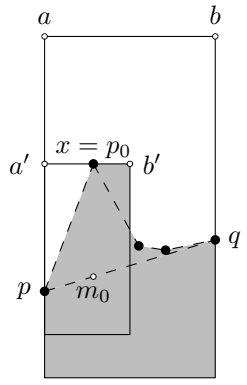

Figure 9.12: An inductive path from p to q

Let m_0 be the projection of x along the vertical axis onto pq. Since m_0 is contained in R, x can see m_0. Since xm_0 and m_0q are visibility edges and m_0 is not the endpoint of a constraint intersecting the interior of triangle xm_0q, we can apply Lemma 2.1 and obtain a convex chain $x = p_0, p_1, ..., p_k = q$ of visibility edges (see Figure 9.12). For each of these visibility edges p_ip_{i+1}, there is a homothet R_i of $C(p, q)$ that falls in one of the following three types (see Figure 9.13): (i) p_i lies on the North boundary and p_{i+1} lies in the Southeast corner, (ii) p_i lies on the West boundary and p_{i+1} lies on the East boundary and the slope of p_ip_{i+1} is negative, (iii) p_i lies on the West boundary and p_{i+1} lies on the East boundary and the slope of p_ip_{i+1} is not negative. Let a_i and b_i be the Northwest and Northeast corner of R_i. We note that by convexity, these three types occur in the order Type (i), Type (ii), and Type (iii).

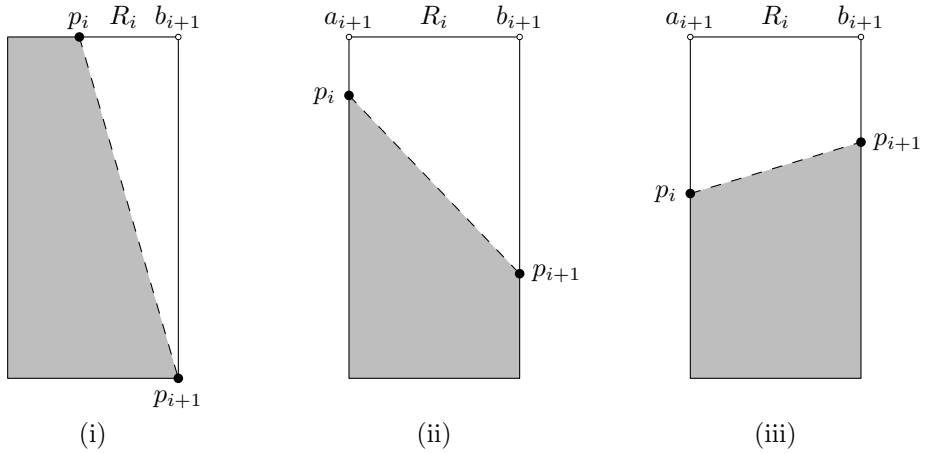

Figure 9.13: The three types of rectangles along the convex chain

Let m_i be the projection of p_i along the vertical along the vertical axis onto pq, let C_i be the homothet of $C(p, q)$ with m_i and m_{i+1} on its boundary, such that the relative area above and below pq is the same as that of $C(p, q)$, and let a'_i and b'_i be the Northwest and Northeast corner of C_i. Using these C_i, we shift Type (ii) and Type (iii) rectangles down as far as possible: We shift R_i down until either p_i or p_{i+1} lies in one of the North corners or the South boundary corresponds to the South boundary of C_i. In the latter case, R_i and C_i are the same rectangle.

Since all rectangles R_i are smaller than $C(p, q)$, we can apply induction, provided that we can show that R_i is half-empty. For Type (i) visibility edges, the part of the rectangle that lies below the line through p_i and p_{i+1} is contained in R, which does not contain any visible vertices, and the region below the convex chain, which

is empty. For Type (ii) and Type (iii) visibility edges, the part of the rectangle that lies below the line through p_i and p_{i+1} is contained in the region below the convex chain, which is empty, and the region of $C(p,q)$ below the line through p and q, which does not contain any visible vertices by Lemma 9.10. Hence, all R_i are half-empty and we obtain an inductive path of length at most: (i) $|p_ib_i| + |b_ip_{i+1}|$, (ii) $|p_ia_i| + |a_ib_i| + |b_ip_{i+1}|$, (iii) $|p_ia_i| + |a_ib_i| + |b_ip_{i+1}|$.

To bound the total path length, we perform case distinction on the location of x on R and whether the convex path from x to q goes down: (a) x lies on the East boundary of R and the convex path does not go down, (b) x lies on the East boundary of R and the convex path goes down, (c) x lies on the North boundary of R and the convex path does not go down, (d) x lies on the North boundary of R and the convex path goes down.

Case (a): Vertex x lies on the East boundary of R and the convex path does not go down. Recall that the length of the path from p to x is at most $|pa'| + |a'b'| + |b'x|$, which is at most $|pa'| + |a'b'| + |b'm_0|$. Since the convex chain does not go down, it cannot contain any Type (i) or Type (ii) visibility edges. Furthermore, since x lies on the East boundary of R, R and all C_i are disjoint. Thus, Lemma 9.11 implies that the boundaries above pq of R and all C_i sum up to $|pa| + |ab| + |bq|$. Hence, if we can show that, for all R_i, $|p_ia_i| + |a_ib_i| + |b_ip_{i+1}| \leq |m_ia_i'| + |a_i'b_i'| + |b_i'm_{i+1}|$, the proof of this case is complete.

By convexity, the slope of p_ip_{i+1} is at most that of pq and m_im_{i+1}. Hence, when p_{i+1} lies in the Northeast corner of R_i, we have $p_{i+1} = b_i$ and $|p_ia_i| + |a_ip_{i+1}| \leq |m_ia_i'| + |a_i'b_i'| + |b_i'm_{i+1}|$. If p_{i+1} does not lies in the Northeast corner, $R_i = C_i$. Hence, since p_i and p_{i+1} lie above pq, we have that $|p_ia_i| + |a_ib_i| + |b_ip_{i+1}| \leq |m_ia_i'| + |a_i'b_i'| + |b_i'm_{i+1}|$.

Case (b): Vertex x lies on the East boundary of R and the convex path goes down. Recall that the length of the path from p to x is at most $|pa'| + |a'b'| + |b'x|$. Let p_j be the lowest vertex along the convex chain. Since p_j lies above pq and pq has non-negative slope, the convex path cannot have gone down more than $|xm_0|$. Hence, when we charge this to R, we used $|pa'| + |a'b'| + |b'm_0|$ of its boundary (see Figure 9.14).

Like in the Case (a), since x lies on the East boundary of R, R and all C_i are disjoint. Thus, Lemma 9.11 implies that the boundaries above pq of R and all C_i sum up to $|pa| + |ab| + |bq|$. Hence, if we can show that, for all R_i, the inductive path length at most $|m_ia_i'| + |a_i'b_i'| + |b_i'm_{i+1}|$, the proof of this case is complete.

For Type (i) visibility edges, we have already charged $|b_ip_{i+1}|$ to R, so it remains to show that $|p_ib_i| \leq |m_ia_i'| + |a_i'b_i'| + |b_i'm_{i+1}|$. This follows, since m_i and m_{i+1} are

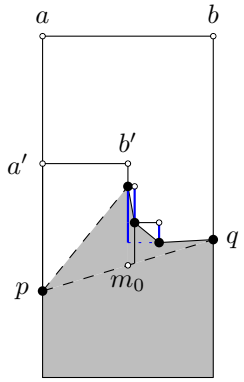

Figure 9.14: Going down along the convex chain (blue segments) is charged to R

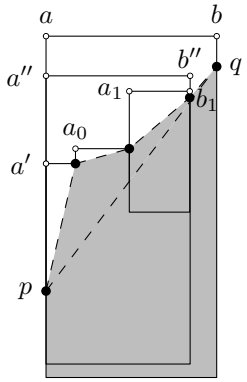

Figure 9.15: Charging the path from p to p_j to $C(p, p_j)$

the vertical projections of p_i and p_{i+1}, which implies that $|p_ib_i| = |a'_ib'_i|$.

For Type (ii) visibility edges, we already charged $|b_ip_{i+1}| - |p_ia_i|$ to R, so we can consider p_ip_{i+1} to be horizontal and it remains to charge the remaining $2 \cdot |p_ia_i| + |a_ib_i|$. If p_i lies in the Northwest corner of R_i, it follows that $|p_ia_i| = 0$ and we have that $|p_ib_i| = |a'_ib'_i| \leq |m_ia'_i| + |a'_ib'_i| + |b'_im_{i+1}|$. If p_i does not lie in the Northwest corner, R_i is the same as C_i. Hence, since we can consider p_ip_{i+1} to be horizontal and p_i and p_{i+1} lie above pq, it follows that $2 \cdot |p_ia_i| + |a_ib_i| \leq |m_ia'_i| + |a'_ib'_i| + |b'_im_{i+1}|$.

Finally, Type (iii) visibility edges are charged as in Case (a), hence we have that $|p_ia_i| + |a_ib_i| + |b_ip_{i+1}| \leq |m_ia'_i| + |a'_ib'_i| + |b'_im_{i+1}|$, completing the proof of this case.

Case (c): Vertex x lies on the North boundary of R and the convex path does not go down. Recall that the length of the path from p to x is at most $|pa'| + |a'x|$. Since the convex chain does not go down, it cannot contain any Type (i) or Type (ii) visibility edges. Let p_j be the first vertex along the chain, such that R_{j-1} is the same as C_{j-1}. Since q lies on the East boundary of $C(p, q)$, this condition is satisfied for the last visibility edge along the convex chain, hence p_j exists.

Let $C(p, p_j)$ be the homothet of $C(p, q)$ that has p and p_j on its boundary, such that the relative area above and below pq is the same as that of $C(p, q)$. Let a'' and b'' be the Northwest and Northeast corners of $C(p, p_j)$ (see Figure 9.15). Since p_j is first vertex along the convex chain that does not lie in the Northeast corner of R_{j-1}, we have that along the path from p to p_j the projections of $a'x$, all a_ip_{i+1}, and $a_{j-1}b_{j-1}$ onto $a''b''$ are disjoint and the projections of pa', all p_ia_i, and $p_{j-1}a_{j-1}$ onto pa'' are disjoint. Hence, their total lengths sum up to at most $|pa''| + |a''b''|$. Finally, since $|b_{j-1}p_j| \leq |b''p_j|$, the total length of the path from p to p_j is at most

$|pa''| + |a''b''| + |b''p_j|$, which is at most $|pa''| + |a''b''| + |b''m_j|$.

All Type (iii) visibility edges following p_j are charged as in Case (a), hence we have that $|p_ia_i|+|a_ib_i|+|b_ip_{i+1}| \leq |m_ia_i'|+|a_i'b_i'|+|b_i'm_{i+1}|$. We now apply Lemma 9.11 to $C(p, p_j)$ and all C_i following p_j and obtain that the total length of the path from p to q is at most $|pa| + |ab| + |bq|$.

Case (d): Vertex x lies on the North boundary of R and the convex path goes down. Recall that the length of the path from p to x is at most $|pa'| + |a'x|$ and that p_1 is the neighbor of x along the convex chain. Let $C(p, p_1)$ be the homothet of $C(p, q)$ that has p and p_1 on its boundary, such that the relative area above and below pq is the same as that of $C(p, q)$. Let a'' and b'' be the Northwest and Northeast corners of $C(p, p_1)$. Since p_1 lies to the right of R and lower than x, it lies on the East boundary of $C(p, p_1)$. We first show that the length of the path from p to p_1 is at most $|pa''| + |a''b''| + |b''p_1|$.

If xp_1 is a Type (i) visibility edge, the length of the path from x to p_1 is at most $|xb_0| + |b_0p_1|$. Hence we have a path from p to p_1 of length at most $|pa'| + |a'x| + |xb_0| + |b_0p_1| = |pa'| + |a''b''| + |b_0p_1|$. Since $|pa'| \leq |pa''|$ and $|b_0p_1| \leq |b''p_1|$, this implies that the path has length at most $|pa''| + |a''b''| + |b''p_1|$. If xp_1 is a Type (ii) visibility edge and x lies in the Northwest corner an analogous argument shows that the path from p to p_1 is at most $|pa''| + |a''b''| + |b''p_1|$. If xp_1 is a Type (ii) visibility edge and $R_0 = C_0$, we have that the projections of $a'x$ and a_0b_0 onto $a''b''$ are disjoint and the projections of pa' and xa_0 onto pa'' are disjoint. Hence, their total lengths sum up to at most $|pa''| + |a''b''|$. Finally, since $|b_0p_1| \leq |b''p_1|$, the total length of the path from p to p_1 is at most $|pa''| + |a''b''| + |b''p_1|$.

Next, we observe, like in Case (b), that starting from p_1 the convex path cannot go down more than $|p_1m_1|$. Hence, when we charge this to $C(p, p_1)$, we used $|pa''| + |a''b''| + |b''m_1|$ of its boundary. Finally, we use arguments analogous to the ones in Case (b) to show that each inductive path after p_1 has length at most $|m_ia_i'| + |a_i'b_i'| + |b_i'm_{i+1}|$. We now apply Lemma 9.11 to $C(p, p_1)$ and all C_i following p_1 and obtain that the total length of the path from p to q is at most $|pa|+|ab|+|bq|$. □

Lemma 9.13 *Let p and q be two vertices that can see each other. Let $C(p, q)$ be the rectangle with p and q on its boundary, such that p lies in a corner of $C(p, q)$. Let l and s be the length of the long and short side of $C(p, q)$. The constrained generalized Delaunay graph contains a path between p and q of length at most $\left(\frac{2l}{s} + 1\right) \cdot (|p_x - q_x| + |p_y - q_y|)$.*

Proof. We assume without loss of generality that p lies on the Southwest corner and q lies on the East boundary. Note that this implies that the slope of pq is non-negative, i.e. $p_x < q_x$ and $p_y \leq q_y$. We prove the lemma by induction on the size of $C(x, y)$, for any two visible vertices x and y, such that x lies in a corner of $C(x, y)$. In fact, we show that the constrained generalized Delaunay graph contains a path between x and y of length at most $c \cdot (q_x - p_x) + d \cdot (q_y - p_y)$ and derive bounds on c and d.

Base case: If $C(p, q)$ is the smallest rectangle with p in a corner, then $C(p, q)$ does not contain any vertices visible to both p and q: Let u be a vertex in $C(p, q)$ that is visible to both p and q. Let $C(p, u)$ be the rectangle with p in a corner and u on its boundary. Since u lies in $C(p, q)$, $C(p, u)$ is smaller than $C(p, q)$, contradicting that $C(p, q)$ is the smallest rectangle with p in a corner. Hence, $C(p, q)$ does not contain any vertices visible to both p and q, which implies that pq is an edge of the constrained generalized Delaunay graph. Hence, the constrained generalized Delaunay graph contains a path between p and q of length at most $|pq| \leq (q_x - p_x) + (q_y - p_y) \leq c \cdot (q_x - p_x) + d \cdot (q_y - p_y)$, provided that $c \geq 1$ and $d \geq 1$.

Induction step: We assume that for all rectangles $C(x, y)$, with x in some corner of $C(p, q)$, smaller than $C(p, q)$ the lemma holds. If pq is an edge of the constrained generalized Delaunay graph, by the triangle inequality, the length of the shortest path from p to q is at most $|pq| \leq |p_x - q_x| + |p_y - q_y|$.

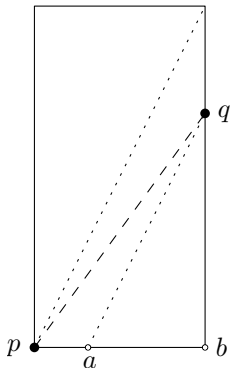

Figure 9.16: There exists a vertex u visible to both p and q below pq

If there is no edge between p and q, there exists a vertex u in $C(p, q)$ that is visible from both p and q. We first look at the case where u lies below pq. Let a be the intersection of the South boundary of $C(p, q)$ and the line though q parallel to the diagonal of $C(p, q)$ through p, and let b be the Southeast corner of $C(p, q)$ (see Figure 9.16). If u lies in triangle paq, by induction we have that the path from p to u has length at most $c \cdot (u_x - p_x) + d \cdot (u_y - p_y)$ and the path from u to q has length

at most $c \cdot (q_x - u_x) + d \cdot (q_y - u_y)$. Hence, there exists a path from p to q via u of length at most $c \cdot (q_x - p_x) + d \cdot (q_y - p_y)$.

If u lies in triangle abq, by induction we have that the path from p to u has length at most $c \cdot (u_x - p_x) + d \cdot (u_y - p_y)$ and the path from q to u has length at most $d \cdot (q_x - u_x) + c \cdot (q_y - u_y)$. When we take c and d to be equal, this implies that there exists a path from p to q via u of length at most $c \cdot (q_x - p_x) + d \cdot (q_y - p_y)$.

If there does not exist a vertex below pq that is visible to both p and q, than Lemma 9.2 implies that there are no vertices in $C(p,q)_q^p$ below pq that are visible to p and that there are no vertices in $C(p,q)_p^q$ below pq that are visible to q. Hence, we can apply Lemma 9.12 and obtain that there exists a path between p and q of length at most $|pa| + |ab| + |bq|$, where a and b are the Northwest and Northeast corner of $C(p,q)$. Since $|ab|$ is $(q_x - p_x)$ and $|bq| \leq |pa| \leq \frac{l}{s} \cdot (q_x - p_x)$, we can upper bound $|pa| + |ab| + |bq|$ by $c \cdot (q_x - p_x)$ when c is at least $\left(\frac{2l}{s} + 1\right)$. Hence, since c and d need to be equal, we obtain that the all cases work out when $c = d = \left(\frac{2l}{s} + 1\right)$, completing the proof. □

Corollary 9.2 *The constrained generalized Delaunay graph using an empty rectangle as empty convex shape has spanning ratio at most* $\sqrt{2} \cdot \left(\frac{2l}{s} + 1\right)$.

This reduces the dependency on the ratio between l and s to linear, instead of quadratic or even cubic if the graph is not a triangulation, as implied by Theorem 9.1. We also observe a difference compared to the tight spanning ratio of $\sqrt{4 + 2\sqrt{2}} \approx 2.61$ for the unconstrained Delaunay graph that uses a square as its convex shape, since the above corollary implies a spanning ratio of at most $3\sqrt{2} \approx 4.25$ in the constrained setting. It is currently not known whether the constrained setting has a strictly larger spanning ratio than the unconstrained setting.

9.4 Conclusion

We showed that every constrained generalized Delaunay graph is a plane spanner, whose spanning ratio depends on the α-diamond property and the visible-pair κ-spanner property. In the special case where the empty convex shape is a rectangle, we reduce the spanning ratio by showing that it depends linearly on the aspect ratio of the rectangle used to construct the graph.

In light of other recent results in the constraint setting, such as the fact that Yao- and θ-graphs with sufficiently many cones are spanners, the result presented in this

chapter raises a tantalizing question: What conditions need to hold for a graph to be a spanner in the constrained setting? In particular, these and previous results show a number of sufficient conditions, but do not immediately give rise to a set of necessary conditions.

Chapter 10

Conclusion

We studied a number of geometric spanners and showed that three well-known types of graphs that are spanners in the unconstrained setting are also spanners in the constrained setting. Specifically, we showed that constrained θ-graphs with at least 6 cones, constrained Yao-graphs with at least 7 cones, and constrained generalized Delaunay graphs are spanners. It is not too surprising that the spanning ratio of θ- and Yao-graphs depends on the number of cones and that the spanning ratio of constrained generalized Delaunay graphs depends on the properties of the empty convex shape used to construct it.

We studied θ-graphs intensively and provided tight bounds on the spanning ratio of large families of θ-graphs in the unconstrained, constrained, and ordered setting. We also showed that though ordered θ-graphs have some additional nice properties, such as logarithmic diameter and logarithmic maximum degree, the spanning ratio obtained in these graphs is worse than that of its unordered counterparts. Most surprisingly, we showed that in the ordered setting we need at least 7 cones for the graph to be a spanner, whereas in the unordered setting 4 cones suffice.

We also showed how to construct a bounded degree spanner in the constrained setting, by starting from the constrained half-θ_6-graph, i.e. the constrained Delaunay triangulation that uses an equilateral triangle as its empty convex shape. We also showed how to route locally and competitively on both the constrained θ_6-graph and the constrained half-θ_6-graph. These are the first local competitive routing algorithms in the constrained setting.

Though a lot of progress was made on the problems we studied, a number of open problems remain. The most obvious one is the question whether the upper bounds on the spanning ratio of the various graphs are tight, i.e. whether there exist matching

lower bounds. This problem can be approached from two angles: one being the search for stronger lower bounds, while the other focuses on improving the current upper bounds.

It is also interesting to see whether the technique to bound the maximum degree of the half-θ_6-graph in the constrained setting can be applied to other constrained graphs. Furthermore, in light of the recent results in the unconstrained setting [9] that show how to construct a spanner of maximum degree 4, it may be possible to reduce the maximum degree in the constrained setting as well.

Since one of the main purposes of building spanners is to use them to route efficiently between vertices, however, one of the most fundamental open problems is that of routing locally and competitively in the constrained setting. Though we showed how to do this on the constrained θ_6-graph and the constrained half-θ_6-graph, it remains unknown how to do this on any of the other graphs we studied. Specifically, unlike in the unconstrained setting, there is currently no local competitive routing algorithm known that works on all θ- and Yao-graphs. In the unconstrained setting, θ-routing is a local and competitive routing algorithm, but this algorithm fails in the constrained setting: Due to the presence of constraints, there need not be an edge in the cone of the current vertex that contains the destination, which implies that the algorithm may be unable to proceed.

Finally, though we have shown for a number of different graphs that they are spanners in the constrained setting, we do not have a nice formulation of the precise conditions a spanner in the unconstrained setting needs to meet in order to also be a spanner in the constrained setting. We showed a number of sufficient conditions, such as the different distance functions used in the various graphs. However, as these are not all required to hold at the same time, they do not form a set of necessary conditions.

Bibliography

[1] O. Aichholzer, S. W. Bae, L. Barba, P. Bose, M. Korman, A. van Renssen, P. Taslakian, and S. Verdonschot. Theta-3 is connected. In *Proceedings of the 25th Canadian Conference on Computational Geometry (CCCG 2013)*, pages 205–210, 2013.

[2] I. Althöfer, G. Das, D. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry (DCG)*, 9(1):81–100, 1993.

[3] R. A. Baeza-Yates, J. C. Culberson, and G. J. Rawlins. Searching in the plane. *Information and Computation*, 106(2):234–252, 1993.

[4] L. Barba, P. Bose, M. Damian, R. Fagerberg, W. L. Keng, J. O'Rourke, A. van Renssen, P. Taslakian, S. Verdonschot, and G. Xia. New and improved stretch factors of Yao graphs. In *Proceedings of the 30th ACM Symposium on Computational Geometry (SoCG 2014)*, pages 30–39, 2014.

[5] L. Barba, P. Bose, J.-L. De Carufel, A. van Renssen, and S. Verdonschot. On the stretch factor of the theta-4 graph. In *Proceedings of the 13th Algorithms and Data Structures Symposium (WADS 2013)*, volume 8037 of *Lecture Notes in Computer Science*, pages 109–120, 2013.

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

[7] N. Bonichon, C. Gavoille, N. Hanusse, and L. Perkovic. Plane spanners of maximum degree six. In *Proceedings of the 37th International Colloquium on Automata, Languages and Programming (ICALP 2010 (1))*, pages 19–30, 2010.

[8] N. Bonichon, C. Gavoille, N. Hanusse, and L. Perković. The stretch factor of L_1- and L_∞-Delaunay triangulations. In *Proceedings of the 20th Annual European Symposium on Algorithms (ESA 2012)*, pages 205–216, 2012.

[9] N. Bonichon, I. Kanj, L. Perkovic, and G. Xia. There are plane spanners of maximum degree 4. In *Proceedings of the 30th ACM Symposium on Computational Geometry (SoCG 2014)*, pages 20–29, 2014.

[10] P. Bose, A. Brodnik, S. Carlsson, E. D. Demaine, R. Fleischer, A. López-Ortiz, P. Morin, and I. J. Munro. Online routing in convex subdivisions. *International Journal of Computational Geometry & Applications (IJCGA)*, 12(04):283–295, 2002.

[11] P. Bose, P. Carmi, and L. Chaitman-Yerushalmi. On bounded degree plane strong geometric spanners. *Journal of Discrete Algorithms (JDA)*, 15:16–31, 2012.

[12] P. Bose, P. Carmi, S. Collette, and M. Smid. On the stretch factor of convex Delaunay graphs. *Journal of Computational Geometry (JoCG)*, 1(1):41–56, 2010.

[13] P. Bose, M. Damian, K. Douïeb, J. O'Rourke, B. Seamone, M. Smid, and S. Wuhrer. $\pi/2$-angle Yao graphs are spanners. *ArXiv e-prints*, 2010. arXiv:1001.2913 [cs.CG].

[14] P. Bose, M. Damian, K. Douïeb, J. O'Rourke, B. Seamone, M. Smid, and S. Wuhrer. $\pi/2$-angle Yao graphs are spanners. *International Journal of Computational Geometry & Applications (IJCGA)*, 22(1):61–82, 2012.

[15] 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 the 24th Canadian Conference on Computational Geometry (CCCG 2012)*, pages 305–310, 2012.

[16] P. Bose, J.-L. De Carufel, P. Morin, A. van Renssen, and S. Verdonschot. Towards tight bounds on theta-graphs. *Submitted to Theoretical Computer Science (TCS)*, 2014.

[17] P. Bose, L. Devroye, M. Löffler, J. Snoeyink, and V. Verma. Almost all Delaunay triangulations have stretch factor greater than $\pi/2$. *Computational Geometry: Theory and Applications (CGTA)*, 44(2):121–127, 2011.

[18] P. Bose, R. Fagerberg, A. van Renssen, and S. Verdonschot. Competitive routing in the half-θ_6-graph. In *Proceedings of the 23rd ACM-SIAM Symposium on Discrete Algorithms (SODA 2012)*, pages 1319–1328, 2012.

[19] P. Bose, R. Fagerberg, A. van Renssen, and S. 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.

[20] P. Bose, R. Fagerberg, A. van Renssen, and S. Verdonschot. Competitive local routing with constraints. *Submitted to the 26th ACM-SIAM Symposium on Discrete Algorithms (SODA 2015)*, 2015.

[21] P. Bose, J. Gudmundsson, and P. Morin. Ordered theta graphs. *Computational Geometry: Theory and Applications (CGTA)*, 28(1):11–18, 2004.

[22] P. Bose, J. Gudmundsson, and M. Smid. Constructing plane spanners of bounded degree and low weight. *Algorithmica*, 42(3-4):249–264, 2005.

[23] P. Bose and J. M. Keil. On the stretch factor of the constrained Delaunay triangulation. In *Proceedings of the 3rd International Symposium on Voronoi Diagrams in Science and Engineering (ISVD 2006)*, pages 25–31, 2006.

[24] P. Bose, A. Lee, and M. Smid. On generalized diamond spanners. In *Proceedings of the 10th Workshop on Algorithms and Data Structures (WADS 2007)*, volume 4619 of *Lecture Notes in Computer Science*, pages 325–336, 2007.

[25] P. Bose, A. Maheshwari, G. Narasimhan, M. Smid, and N. Zeh. Approximating geometric bottleneck shortest paths. *Computational Geometry: Theory and Applications (CGTA)*, 29(3):233–249, 2004.

[26] P. Bose, P. Morin, and A. van Renssen. The price of order. *Accepted to the 25th International Symposium on Algorithm and Computation (ISAAC 2014)*, 2014.

[27] P. Bose, P. Morin, A. van Renssen, and S. Verdonschot. The θ_5-graph is a spanner. In *Proceedings of 39th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2013)*, volume 8165 of *Lecture Notes in Computer Science*, pages 100–114, 2013.

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

[29] P. Bose, M. Smid, and D. Xu. Delaunay and diamond triangulations contain spanners of bounded degree. *International Journal of Computational Geometry & Applications (IJCGA)*, 19(02):119–140, 2009.

[30] P. Bose and A. 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.

[31] P. Bose and A. van Renssen. Upper bounds on the spanning ratio of constrained theta-graphs. *ArXiv e-prints*, 2014. arXiv:1401.2127 [cs.CG].

[32] P. Bose, A. van Renssen, and S. Verdonschot. On the spanning ratio of theta-graphs. In *Proceedings of the 13th Workshop on Algorithms and Data Structures (WADS 2013)*, volume 8037 of *Lecture Notes in Computer Science*, pages 182–194, 2013.

[33] J. Buter. Überkonvexe Mengen in der Ebene. In *Proceedings of the Koninklijke Nederlandse Akademie van Wetenschappen*, volume 41, pages 756–762, 1938.

[34] L. P. Chew. There is a planar graph almost as good as the complete graph. In *Proceedings of the 2nd Annual Symposium on Computational Geometry (SoCG 1986)*, pages 169–177, 1986.

[35] L. P. Chew. There are planar graphs almost as good as the complete graph. *Journal of Computer and System Sciences (JCSS)*, 39(2):205–219, 1989.

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

[37] M. Damian and K. Raudonis. Yao graphs span theta graphs. *Discrete Mathematics, Algorithms and Applications (DMAA)*, 4(02):1250024, 16 pages, 2012.

[38] G. 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.

[39] G. Das and D. Joseph. Which triangulations approximate the complete graph? In *Proceedings of the International Symposium on Optimal Algorithms*, volume 401 of *Lecture Notes in Computer Science*, pages 168–192, 1989.

[40] D. P. Dobkin, S. J. Friedman, and K. J. Supowit. Delaunay graphs are almost as good as complete graphs. *Discrete & Computational Geometry (DCG)*, 5(1):399–407, 1990.

[41] N. M. El Molla. Yao spanners for wireless ad hoc networks. Master's thesis, Villanova University, 2009.

[42] B. E. Flinchbaugh and L. K. Jones. Strong connectivity in directional nearest-neighbor graphs. *SIAM Journal on Algebraic and Discrete Methods (JADM)*, 2(4):461–463, 1981.

[43] B. Joe and C. A. Wang. Duality of constrained Voronoi diagrams and Delaunay triangulations. *Algorithmica*, 9(2):142–155, 1993.

[44] I. A. Kanj and L. Perkovic. On geometric spanners of Euclidean and unit disk graphs. In *Proceedings of the 25th Annual Symposium on the Theoretical Aspects of Computer Science (STACS 2008)*, pages 409–420, 2008.

[45] I. A. Kanj and G. Xia. Improved local algorithms for spanner construction. *Theoretical Computer Science (TCS)*, 453:54–64, 2012.

[46] J. Keil. Approximating the complete Euclidean graph. In *Proceedings of the 1st Scandinavian Workshop on Algorithm Theory (SWAT 1988)*, pages 208–213, 1988.

[47] J. M. Keil and C. A. Gutwin. Classes of graphs which approximate the complete Euclidean graph. *Discrete & Computational Geometry (DCG)*, 7(1):13–28, 1992.

[48] R. Klein and D. Wood. Voronoi diagrams based on general metrics in the plane. In *Proceedings of the 5th Annual Symposium on Theoretical Aspects of Computer Science (STACS 1988)*, volume 294 of *Lecture Notes in Computer Science*, pages 281–291, 1988.

[49] X.-Y. Li and Y. Wang. Efficient construction of low weighted bounded degree planar spanner. *International Journal of Computational Geometry & Applications (IJCGA)*, 14(01-02):69–84, 2004.

[50] H. Martini, K. J. Swanepoel, and G. Weiß. The geometry of Minkowski spaces - A survey. Part I. *Expositiones Mathematicae*, 19(2):97–142, 2001.

[51] A. E. Mayer. Eine Überkonvexität. *Mathematische Zeitschrift*, 39(1):511–531, 1935.

[52] S. Misra, S. C. Misra, and I. Woungang. *Guide to Wireless Sensor Networks*. Springer, 2009.

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

[54] H. Räcke. Survey on oblivious routing strategies. In *Mathematical Theory and Computational Practice*, volume 5635 of *Lecture Notes in Computer Science*, pages 419–429, 2009.

[55] J. Ruppert and R. Seidel. Approximating the d-dimensional complete Euclidean graph. In *Proceedings of the 3rd Canadian Conference on Computational Geometry (CCCG 1991)*, pages 207–210, 1991.

[56] K. Swanepoel. Helly-type theorems for homothets of planar convex curves. *Proceedings of the American Mathematical Society*, 131(3):921–932, 2003.

[57] A. van Renssen. On the spanning ratio of constrained Yao-graphs. In *Proceedings of the 26th Canadian Conference on Computational Geometry (CCCG 2014)*, 2014.

[58] G. Xia. The stretch factor of the Delaunay triangulation is less than 1.998. *SIAM Journal on Computing*, 42(4):1620–1659, 2013.

[59] G. Xia and L. Zhang. Toward the tight bound of the stretch factor of Delaunay triangulations. In *Proceedings of the 23rd Canadian Conference on Computational Geometry (CCCG 2011)*, pages 175–180, 2011.

[60] A. C. C. Yao. On constructing minimum spanning trees in k-dimensional spaces and related problems. *SIAM Journal on Computing*, 11(4):721–736, 1982.