


Further Results on Arc and Bar k -Visibility Graphs

Mehtaab Sawhney and Jonathan Weed

Massachusetts Institute of Technology



The Minnesota Journal of Undergraduate Mathematics

Volume 3 (2017-2018 Academic Year)

The Minnesota Journal of Undergraduate Mathematics

Volume 3 (2017-2018 Academic Year)

Further Results on Arc and Bar k -Visibility Graphs

Mehtaab Sawhney* and Jonathan Weed

Massachusetts Institute of Technology

ABSTRACT. We consider visibility graphs involving bars and arcs in which lines of sight can pass through at most k objects. We prove a new edge bound for arc k -visibility graphs, provide maximal constructions for arc visibility graphs and semi-arc k -visibility graphs, and give a complete characterization of semi-arc visibility graphs. We further show that the family of arc i -visibility graphs is never contained in the family of bar j -visibility graphs for any i and j , and that the family of bar i -visibility graphs is not contained in the family of bar j -visibility graphs for $i \neq j$. Finally, we give the first thickness bounds for arc and semi-arc k -visibility graphs.

1. INTRODUCTION

Visibility graphs are a general class of graphs that represent lines of sight between objects. For this paper we will concentrate on visibility graphs between one dimensional objects in the plane. Several visibility graphs of this type have attracted particular interest, such as bar k -visibility graphs (introduced by Dean et al. [5]), semi-bar k -visibility graphs (introduced by Felnsner and Massow [9]), and arc and semi-arc k -visibility graphs (introduced by Hutchinson [13] and Babbitt et al. [2]). In each case, standard questions include how many edges such a graph can have, how large the thickness of such a graph can be, and in general what properties of a graph guarantee that it is representable as a visibility graph. (See [1, 2, 5, 9, 12, 14, 18] for many results of this type.)

Babbitt et al. [2] proved the first such results for the class of arc and semi-arc k -visibility graphs. Specifically, they provided an upper bound for the maximum number of edges in an arc k -visibility graph on n vertices and lower and upper bounds on the maximum number of edges in a semi-arc k -visibility graph on n vertices. In the case of semi-arc k -visibility graphs, they conjectured their upper bound not to be optimal.

In this work, we extend the work of Babbitt et al. [2]. In Section 3, we give a complete characterization of semi-arc visibility graphs, which complements the characterization of arc visibility graphs given by Hutchinson [13]. This characterization implies that semi-arc visibility graphs are in fact in general planar. In Section 4, we give a stronger edge bound for arc k -visibility graphs and provide a construction showing that the upper bound on the number of edges of a semi-arc k -visibility graph given in Babbitt et al. [2] is tight,

* *Corresponding author*

thereby disproving their conjecture. In Section 5, we give the first nontrivial thickness bounds for arc and semi-arc k -visibility graphs. Finally, in Section 6, we consider the relationship between arc k -visibility graphs and the more common class of bar k -visibility graphs.

2. PRELIMINARIES

2.1. Visibility and k -visibility graphs. A visibility graph is a graph whose vertices correspond to objects in the plane. Two vertices are adjacent whenever the corresponding objects are connected by an unobstructed line of sight. A graph arising in this way is known as a *visibility graph*, and the corresponding arrangement of objects is known as a *visibility representation*.

Dean et al. [5] generalized this concept by defining *k -visibility graphs*, which are identical to visibility graphs except that a line of sight may intersect up to k objects in addition to the two objects it connects. Visibility graphs are obtained in the special case where $k = 0$.




FIGURE 1. Bar k -visibility


We consider the following four types of visibility graphs.

Definition 2.1. A *bar k -visibility graph* is a graph corresponding to an arrangement of nonintersecting closed horizontal line segments in the plane (“bars”). Each vertex of the graph corresponds to a bar, and two vertices are adjacent if and only if the corresponding bars are connected by a vertical line of sight passing through at most k other bars.

We note that two different classes of graphs are obtained if lines of sight are defined to be zero-width line segments (“strong visibility”) or rectangles of positive width (“ ϵ -visibility”). For example $K_{2,3}$ has no strong visibility representation but has an epsilon visibility representation. In this work we always adopt the notion of “strong visibility” for visibility graphs being considered.

Figure 1 shows a collection of bars and the corresponding visibility and 1-visibility graphs.

Definition 2.2. An *arc k -visibility graph* is a graph corresponding to an arrangement of nonintersecting concentric circular arcs. Each vertex of the graph corresponds to an arc, and two vertices are adjacent if and only if the corresponding arcs are connected by a radial line segment (which may pass through the center of the circle) intersecting at most k other arcs.

FIGURE 2. Arc k -visibility

Examples of arc k -visibility graphs appear in Figure 2.

Arc k -visibility graphs were introduced by Babbitt et al. [2] (although the $k = 0$ case was already introduced by Hutchinson [13] under the name polar visibility graphs) and possess a natural connection to the geometry of the projective plane.

There are two subtleties in defining arc k -visibility graphs. First, none of the arcs in an arc k -visibility representation are complete circles. (Graphs obtained by allowing both circles and arcs are called “circle visibility graphs” by Hutchinson [13].) The results and proofs below are stated only for arc k -visibility graphs. However, any circle k -visibility graph may be turned into an arc k -visibility graph by adding a small gap in any complete circles, at the price of possibly adding some edges. The upper bounds we prove for arc k -visibility graphs therefore hold for circle k -visibility graphs as well.

Second, a radial line of sight may intersect an obtuse arc more than once. We adopt the convention that for the purpose of counting visibilities these double intersections are counted only once.

Definition 2.3. A *semi-bar k -visibility graph* is a bar k -visibility graph where the left end-points of all the bars in the bar k -visibility representation lie on the same vertical line. A *semi-arc k -visibility graph* is an arc k -visibility graph in which all arcs in the arc k -visibility representation extend in a counterclockwise direction from the same radial ray.

Semi-bar and semi-arc k -visibility graphs were introduced by Felsner and Massow [9] and Babbitt et al. [2], respectively. Figure 3 gives examples of semi-bar and semi-arc visibility representations.

3. CHARACTERIZATION OF SEMI-ARC VISIBILITY GRAPHS

In this section, we obtain a full characterization of semi-arc visibility graphs. As a corollary, we obtain the fact that all semi-arc visibility graphs are planar. This is in contrast to the class of arc visibility graphs, which Hutchinson [13] observed contains the non-planar graph K_6 where K_n denotes the complete graph on n vertices (A proof of this fact appears in Section 6.)




FIGURE 3

Given a semi-arc visibility representation, we can divide the visibilities between arcs into two sets. Say that two arcs share a *central visibility* if all lines of sight between the two arcs pass through the center of the circle. Otherwise, if there exists a line of sight between two arcs not passing through the center of the circle, call the visibility between them *non-central*.

The non-central visibilities in a semi-arc visibility representation have the structure of a semi-bar visibility graph. Indeed, these are exactly the visibilities that exist in the semi-bar representation obtained by taking each arc in the semi-arc representation and straightening it into a line segment while maintaining the arcs' relative ordering and length. (See Figure 4.) We therefore refer to non-central visibilities as *semi-bar visibilities*.




FIGURE 4. Semi-arc visibility representation and corresponding semi-bar visibility representation. The arcs and bars have the same relative length and ordering, with the outermost arc corresponding to the topmost bar.

Cobos et al. [4] gave a complete characterization of semi-bar visibility graphs. We extend their result to semi-arc visibility graphs.

Definition 3.1. A graph G is *outerhamiltonian* if it has a planar embedding in which there is path through all the vertices and the edges on this path all lie on the outer face.

Theorem 3.2 (Cobos et al. [4]). *A graph is a semi-bar visibility graph if and only if it is outerhamiltonian. Moreover, the Hamiltonian path can be taken to be the path that visits each bar in order, from top to bottom.*

Given an outerhamiltonian graph and corresponding Hamiltonian path P , say that a vertex is *critical* if it is a cutpoint of the whole graph or is the first or last vertex in P . We first show that all central visibilities involve only critical vertices.

Proposition 3.3. *Consider a semi-arc visibility graph G and a fixed representation corresponding to G . If H is the outerhamiltonian graph corresponding to the semi-bar visibilities of the representation and P is the Hamiltonian path in H that visits each arc in order from outermost to innermost, then $H \subseteq G$ and each edge present in G but not in H connects two critical vertices of H .*

Proof. That $H \subseteq G$ is clear from the definition. If no arc has argument equal or greater than π , then there are no visibilities through the center and $H = G$.

Otherwise, consider the arc with the greatest argument. (If there is more than one such arc, take the innermost one.) Call this arc and its corresponding vertex c_m . Consider all arcs with smaller radius than c_1 , and label the vertex corresponding to the largest such arc c_2 . (If there are multiple candidates, take the innermost one.) Continuing in this way, construct the sequence of arcs c_1, \dots, c_m . Let $C = \{c_1, \dots, c_m\}$. (This procedure is identical to that done for semi-bar k -visibility graphs by Felsner and Massow [9].)

All visibilities through the center of the circle must be between two arcs in C , since by construction every other arc in the visibility representation is blocked from seeing the center of the circle by an element of C .

The vertices corresponding to the elements of C are critical vertices of H . Given $c \in C$, if c is the arc with the largest or smallest radius in the semi-arc visibility representation, then the vertex corresponding to c is the first or last vertex in P . Otherwise, the corresponding vertex must be a cutpoint in the associated semi-bar visibility graph, since the arc blocks all possible semi-bar visibilities between arcs with larger radius and arcs with smaller radius. In either case, the vertex corresponding to c is critical. \square

The proof of Proposition 3.3 establishes something stronger. Call a semi-arc visibility representation *monotone* if, for each pair of arcs, the arc with the smaller radius has strictly smaller argument.


Corollary 3.4. *Given a semi-arc visibility representation, there exists a monotone semi-arc visibility representation with the same central visibilities.*

Proof. Take the semi-arc visibility representation obtained by removing all arcs except those in the set C . This representation has the same central visibilities as the original representation, and by construction the arcs in C strictly decrease in argument as their radius decreases. \square

Corollary 3.4 implies that to characterize the central visibilities of semi-arc visibility representations, it suffices to characterize the central visibilities of *monotone* semi-arc visibility representations. To do so, we define the following family of graphs.

Definition 3.5. An *ordered matching* on $2n$ vertices is a graph with vertices labeled by v_1, \dots, v_{2n} and edges (v_i, v_{i+n}) for $1 \leq i \leq n$. For each vertex v_i , call the vertex v_{i+1} its *successor*.

An example of an ordered matching appears in Figure 5. For completeness, we also recall the following definition.

FIGURE 5. An ordered matching on $2n$ vertices

Definition 3.6. Given a graph G and two (not necessarily adjacent) vertices u and v , the (vertex) contraction of u into v is the graph that is obtained by removing the vertex u and replacing each edge of the form (u, x) for some vertex $x \neq v$ by an edge (v, x) , so long as this edge does not already exist.

Theorem 3.7. A graph G corresponds to the central visibilities in a monotone semi-arc visibility representation if and only if it can be obtained from an ordered matching by repeatedly contracting vertices into their successors.

Proof. We first show that a graph G corresponding to central visibilities in a monotone semi-arc visibility representation has the claimed form.

Fix a monotone semi-arc visibility representation, and fix a line of sight between each pair of arcs sharing a central visibility. Each such line of sight is a segment joining two arcs and passing through the center of the circle. Define the argument of each segment to be the angle in $[0, \pi)$ that the segment makes with the positive x -axis. Denote these segments by ℓ_1, \dots, ℓ_m , where the segments are indexed in order of increasing argument.

For each line of sight ℓ_i with $1 \leq i \leq m$, let a_i and b_i be the two arcs that it joins, where a_i is the arc with smaller radius and b_i the arc with larger radius. We thereby obtain a list $a_1, \dots, a_m, b_1, \dots, b_m$ of arcs in the semi-arc visibility representation. Call this list L .

We show that if an arc appears multiple times in L , then its appearances must be consecutive. To do so, we show that the radii of the arcs in L are nondecreasing. This implies that the appearances of any arc in the list must be consecutive, since if an arc x appears in L between two appearances of an arc y , then x and y have the same radius. Since different arcs in a semi-arc visibility representation do not intersect, the radii of different arcs are distinct. Hence $x = y$.

It therefore suffices to show that the radii of the arcs in L are nondecreasing. By construction, if $i < j$, then line of sight ℓ_i has a smaller argument than ℓ_j . Color the endpoint of ℓ_i that lies on a_i red, and color the endpoint that lies on b_i blue. Since the semi-arc visibility representation is monotone, as we rotate the line of sight ℓ_i counterclockwise to meet ℓ_j , the radii of the arcs that the red endpoint encounters do not decrease, and the radii of the arcs that the blue endpoint encounters also do not decrease. When ℓ_i has been rotated to meet ℓ_j , its red endpoint lies on a_j and its blue endpoint lies on b_j . Therefore a_i does not

have a larger radius than a_j and b_i does not have a larger radius than b_j . This implies that the radii of the arcs in the lists a_1, \dots, a_m and b_1, \dots, b_m are nondecreasing.

Moreover, the radius of a_m is not larger than the radius of b_1 . To see this, note that without loss of generality we can take ℓ_1 to be a segment with argument 0, which in turn implies that b_1 is the arc with the smallest argument greater than or equal to π . On the other hand, by definition the segment ℓ_m has argument strictly less than π . If a_m had a larger radius than b_1 , then the portion of the segment ℓ_m joining a_m to the center of the circle would intersect b_1 as well, which contradicts the definition of a line of sight. Hence the radius of a_m cannot be larger than the radius of b_1 .

We have shown that the radii of the arcs in L are nondecreasing, which implies that the appearances of any arc in L must be consecutive. Let H be an ordered matching on $2m$ vertices, and associate the vertices v_1, \dots, v_{2m} with the elements of L . That is, associate v_1, \dots, v_m with a_1, \dots, a_m and v_{m+1}, \dots, v_{2m} with b_1, \dots, b_m . Whenever an arc in L is the same as the following arc in L , contract the corresponding vertex of H into its successor. Perform this operation repeatedly until each vertex of the graph corresponds to a unique arc.

By construction, the resulting graph corresponds to the central visibilities of the arc visibility representation. Hence the resulting graph is G , and G can be obtained from an ordered matching by repeatedly contracting vertices into their successors, as claimed.

Conversely, suppose that G is obtained from an ordered matching on $2m$ vertices by contracting vertices into their successors. Define a semi-arc visibility representation with $2m$ arcs, labeled a_1, \dots, a_{2m} in order of increasing radius. For $1 \leq i \leq 2m$, let a_i have argument $(i-1)\pi/m$. Then the central visibilities of this representation are all of the form (a_i, a_{m+i}) for $1 \leq i \leq m$. Therefore the central visibilities of this representation form an ordered matching on $2m$ vertices, with arc a_i corresponding to vertex v_i for $1 \leq i \leq 2m$.

Removing the arc a_i from this representation corresponds to contracting the vertex v_i into v_{i+1} . If G is obtained from an ordered matching on $2m$ by contracting each of the vertices in $\{v_{j_1}, \dots, v_{j_k}\}$ into its successor, then by removing arcs $\{a_{j_1}, \dots, a_{j_k}\}$ from the given semi-arc visibility representation we obtain a monotone semi-arc visibility representation whose central visibilities correspond to G , as desired. \square

Corollary 3.8. *A graph G is a semi-arc visibility graph if and only if it can be written as the union of an outerhamiltonian graph H and a graph obtained from an ordered matching by contracting a subset of vertices into their successors whose vertices are a subset S of the critical vertices of H containing at least one endpoint of the Hamiltonian path in H .*

Proof. That G is necessarily of the claimed form follows from Proposition 3.3 and Theorem 3.7 noting that if a semi-arc visibility graph has any visibilities through the center of the circle then the innermost arc certainly has such a visibility.

Conversely, the proof of Theorem 3.2 given by Cobos et al. [4] establishes that any outerhamiltonian graph H can be written as a semi-bar visibility graph, and moreover that the critical points of H can be taken to correspond to bars which are longer than all non-critical bars.

If $H = G$, then there are no central visibilities and the claim follows upon embedding the semi-bar visibility graph in the upper half of a circle. Otherwise, Theorem 3.7 establishes that the edges in G but not in H correspond to the central visibilities of a monotone semi-arc visibility representation. Lengthen the arcs corresponding to critical vertices in the chosen subset S so that they form the required monotone semi-arc visibility representation. Note that the innermost arc corresponds to one of the endpoints of the Hamiltonian path and must lie in the subset S . For each critical vertex $v \notin S$, let w be the vertex in S nearest to v on the Hamiltonian path in H whose arc has smaller radius, and lengthen the arc corresponding to v so that it has the same argument as the arc corresponding to w . Finally, place the remaining arcs between these critical vertices, maintaining their relative lengths and positions. The resulting semi-arc visibility representation has semi-bar visibilities corresponding to H , since the arcs in the resulting visibility representation have the same relative ordering and length as the bars in the semi-bar visibility representation, and by construction they have central visibilities corresponding to edges in G but not in H . Hence the associated semi-arc visibility graph is G , and G is a semi-arc visibility graph, as claimed. \square

Corollary 3.9. *All semi-arc visibility graphs are planar.*

Proof. It suffices to consider the case where the semi-arc visibility representation is monotone. Indeed, label the arcs a_1, \dots, a_n , with indices increasing with increasing radius. Suppose there exists an arc a_j such that the argument of a_j is not larger than the argument of a_{j-1} . Choose the largest j for which this condition holds. Then the vertex corresponding to a_j has degree 1 (if a_j is the outermost arc) or degree 2 (if a_j is not the outermost arc and the argument of a_{j+1} is larger than the argument of a_j). Removing a_j has the effect of contracting the edge between the vertices corresponding to a_j and a_{j-1} , which does not affect the planarity of the graph.

We therefore suppose that the semi-arc visibility representation is monotone. The semi-arc visibility graph G can be written as the union of an outerhamiltonian graph H and a graph on the critical vertices of H . Since the representation is monotone, H is a path. The remaining edges are obtained from an ordered matching by contraction. Therefore G is a minor of the graph shown in Figure 6, where the dotted edges correspond to the subgraph H . This graph is planar, and therefore any of its minors is planar [6]. Hence G is planar, as claimed. \square

4. IMPROVED EDGE BOUNDS FOR ARC AND SEMI-ARC k -VISIBILITY GRAPHS

Babbitt et al. [2] established upper bounds on the total number of edges for arc and semi-arc k -visibility graphs. For arc k -visibility graphs with sufficiently many vertices, they proved that a graph with n vertices can have at most $(k+1)(3n-k-2)$ edges. We improve this bound to $(k+1)(3n - \frac{3k+6}{2})$. For semi-arc k visibility graphs with sufficiently many vertices, they proved that a graph with n vertices can have at most $(k+1)(2n - \frac{k+2}{2})$ edges but conjectured that the correct bound was smaller. We prove that in fact their original




FIGURE 6. Planar embedding of ordered matching (solid edges) on a path (dotted edges)

bound is tight. For comparison, semi-bar k -visibility graphs can have at most $(k + 1)(2n - 2k - 3)$ edges.

Since we seek to establish upper bounds on the number of edges, we assume in this section that each arc has a different radius and moreover that no two endpoints of any two arcs lie on the same radial segment. We can accomplish this without decreasing the number of edges by slightly perturbing arcs and their endpoints. These assumptions are essentially the same as those present in Babbitt et al. [2].

By definition, none of the arcs in an arc k -visibility graph is a circle. Moreover, we assume in this section that none of the arcs is a semi-circle: if an arc is a semi-circle, then since by assumption no two arcs have endpoints lying on the same radial line we can lengthen the arc slightly without decreasing the number of edges.

We begin by establishing some definitions. We begin by defining the *argument* of a point or ray in an arc visibility representation as its angular position, when measured with respect to the positive x -axis. Given an arc in such a representation, it is possible to choose arguments α and β for its endpoints such that $0 < \beta - \alpha < 2\pi$. Call the endpoint corresponding to β the *positive* endpoint and the endpoint corresponding to α the *negative* endpoint. Given an arc k -visibility representation, consider two arcs joined by some line of sight. Consider the set of all valid lines of sight between the two arcs. Two lines of sight are *contiguous* if we can rotate one into the other such that all intervening segments are also valid lines of sight, and we define a *region of visibility* to be the closure of a maximal set of contiguous lines of sight. For each region of visibility, we call the radial segment in it with the smallest argument the *limiting line* of the region.

Following Babbitt et al. [2], we associate edges in an arc k -visibility graph G with arcs according to the limiting lines of their regions of visibility, in the following way.¹

Fix an arc k -visibility representation of G . Suppose that two arcs a_u and a_v in the representation are connected by a line of sight, so that the corresponding vertices u and v are connected by an edge in G . We consider each region of visibility between a_u and a_v

¹Our definitions slightly differ from theirs in that we allow an edge to be associated with multiple arcs.

in turn. Given a region, if the limiting line contains an endpoint of a_u (respectively a_v), then we call the edge between u and v in G a *negative edge* of a_u (respectively a_v). If this limiting line extends inward (toward the center of the circle) from a_u (respectively a_v), then we call this an *inner negative edge*; if it extends outward, we call it an *outer negative edge*.

Otherwise, the limiting line must contain the endpoint of another arc, say a_w . In this case, we call the edge between u and v a *positive edge* of a_w . In this way, we assign each edge to (possibly many) arcs in the visibility representation.

In Figure 7, we give an example showing several limiting lines corresponding to regions of visibility in an arc k -visibility graph. In each case, the edge is assigned to the arc whose endpoint is contained in the limiting line.




FIGURE 7. Inner Negative (orange), Outer Negative (red), Positive (blue) edges

The following lemma establishes a link between the number of regions of visibility between two arcs and the number of arcs that the corresponding edge is assigned to.

Lemma 4.1. *Fix a pair of arcs a_u and a_v in an arc k -visibility representation corresponding to a pair of adjacent vertices u and v in G . If the arcs have at least two distinct regions of visibility, either both through the center of the circle or not through the center of the circle, then the edge between u and v is assigned to at least two different arcs.*

Proof. We assign the edge between u and v to an arc whenever that arc's endpoint lies on the limiting line of a region of visibility. By assumption, there are at least two distinct regions of visibility between a_u and a_v , hence at least two distinct limiting lines. Since these visibilities are either both through the center of the circle or both not, these limiting lines have different arguments. Moreover, it is impossible for an edge to be assigned as both a negative and a positive edge to the same arc because a negative edge is always assigned to an arc corresponding to one of its two vertices, whereas a positive edge is always assigned to a different arc. We therefore conclude that the two different limiting lines must be assigned to two different arcs, as claimed.

□

The following two lemmas are useful for bounding the number of negative and positive edges assigned to each arc.

Lemma 4.2. *If an arc a has ℓ distinct outer (inner) negative edges, then a radial ray extending outward (inward) from the negative endpoint of a intersects at least ℓ other arcs.*

Proof. If arc a has an outer (inner) negative edge with arc b , then a radial ray extending outward (inward) from the negative endpoint of a intersects b . \square

Lemma 4.3. *If an arc a has ℓ distinct positive edges, then a radial ray extending outward or inward from the positive endpoint of a intersects at least ℓ other arcs. Moreover, the radial line containing the positive endpoint of a makes at least $k + \ell + 1$ intersections with other arcs. (These may include multiple intersections with the same arc.)*

Proof. By definition, each positive edge of a involves two other arcs, and the positive endpoint of a lies on the limiting line of a region of visibility between those two arcs. Hence, if an edge between b and c is a positive edge of arc a , then a radial ray extending outward or inward from the positive endpoint of a intersects either b or c .

The radial segment between the two arcs corresponding to this limiting line must intersect exactly k arcs in addition to a . If a has ℓ distinct positive edges, then a radial ray extending outward from the positive endpoint of a intersects ℓ arcs involved in positive edges assigned to a . Let b be the arc in this set with the smallest radius and let c be the arc such that the edge between b and c is a positive edge assigned to a . The radial segment between b and c containing the positive endpoint of a intersects k arcs in addition to a . Therefore the radial line containing the positive endpoint of a intersects these k arcs, the arcs b and c , and the $\ell - 1$ arcs with radius larger than b . We obtain at least $k + \ell + 1$ intersections. \square

4.1. An improved bound for arc k -visibility graphs. Using Lemma 4.1 it is possible to improve the bound on the maximum number of edges given in Babbitt et al. [2] for arc k -visibility graphs.

Theorem 4.4. *The maximum number of edges in an arc k -visibility graph with n vertices is at most $\binom{n}{2}$ for $n \leq 4k + 4$ and $(k + 1)(3n - \frac{3k+6}{2})$ for $n > 4k + 4$.*

Proof. When $n \leq 4k + 4$, the bound is trivial, as any graph on n vertices has at most $\binom{n}{2}$ edges. It therefore suffices to consider $n > 4k + 4$.

Label the outermost $k + 1$ arcs a_{k+1}, \dots, a_1 with indices decreasing with decreasing radius. We first recall how to obtain the bound given in Babbitt et al. [2]. Note that there are at most $k + 1$ inner negative edges, $k + 1$ outer negative edges, and $k + 1$ positive edges associated with each arc. Moreover, for the outermost arcs a stronger bound holds: there are at most $0, \dots, k$ positive edges and $0, \dots, k$ outer negative edges for a_{k+1}, \dots, a_1 respectively. The total number of edges is therefore at most $(3k + 3)n - 2 \sum_{i=1}^{k+1} i = (k + 1)(3n - k - 2)$.

We will show that the above bound necessarily double-counts existing edges or counts non-existent edges. Each over-counted edge will involve an arc a_ℓ , where $1 \leq \ell \leq k + 1$. These over-counted edges are of two types: negative edges between a_ℓ and arcs of smaller radius than a_ℓ , and positive edges involving a_ℓ assigned to arcs of smaller radius than a_ℓ .

We consider in turn a_ℓ , for $1 \leq \ell \leq k + 1$. Say an arc a is in the *cone of visibility* of a_ℓ if any portion of it lies in the blue shaded region of Figure 8, where the left figure corresponds to the case where a_ℓ is obtuse and the right to the case where a_ℓ is acute.




FIGURE 8. Cone of visibility for obtuse and acute arcs. The arc a_ℓ is the solid arc, and the dotted lines indicate the boundaries of the cone of visibility.




FIGURE 9. Definition of S_ℓ for obtuse and acute arcs. We take the ℓ arcs encountered last in moving from the tail to the head of the arrow.

Consider all the arcs in the cone of visibility for a_ℓ in the order shown in Figure 9, where the two different cases correspond to whether a_ℓ is an obtuse or acute arc, and take the ℓ arcs encountered last in moving from the tail to the head of the arrow. If there are fewer than ℓ arcs, take them all. Call this set of arcs S_ℓ .

Let a be any arc in S_ℓ . For the Babbitt et al. [2] bound to be tight, the arc a should have $k + 1$ negative edges in each direction, or, if $a = a_m$ for $1 \leq m < \ell$, at least $k - m + 1$ outer negative edges and $k + 1$ inner negative edges.

Suppose that there is a radial ray originating at the negative endpoint of a and extending in the direction specified by Figure 9 that does not intersect a_ℓ . By construction, this radial ray intersects at most k arcs, $(\ell - 1)$ other arcs in S_ℓ and $(k + 1 - \ell)$ arcs with larger radius than a_ℓ . By Lemma 4.2, a has strictly fewer than $k + 1$ negative edges along this ray. If $a = a_m$ for $m < \ell$ and the ray extends outward from a , then a has strictly fewer than $k - m + 1$ outer negative edges, since there are only $k - m + 1$ arcs with larger radius than a and the ray does not intersect a_ℓ . In either case, the arc a has strictly fewer negative edges along this ray than it must for the Babbitt et al. [2] bound to be tight; in other words, that bound counts a non-existent edge between a and a_ℓ .

The argument for positive edges is extremely similar. Suppose there is a radial ray originating at the positive endpoint of a in the direction specified by Figure 9 that does not intersect a_ℓ . This radial ray intersects at most k arcs, $(\ell - 1)$ other arcs in S_ℓ and $(k + 1 - \ell)$ arcs with larger radius than a_ℓ . By Lemma 4.3, a has strictly fewer than $k + 1$ positive edges. If $a = a_m$ for $m < \ell$ and the ray extends outward from a , then a has strictly fewer than $k - m + 1$ positive edges, since there are only $k - m + 1$ arcs with larger radius than a and the ray does not intersect a_ℓ . If the ray extends inward from a , then it intersects at most k arcs, $(\ell - 1)$ other arcs in S_ℓ and $(k + 1 - \ell)$ arcs with larger radius than a_ℓ . The ray extending in the outward direction from the positive endpoint of a intersects at most $k - m + 1$ arcs because there are only $k - m + 1$ arcs with larger radius than a , so the radial line containing the positive endpoint of a makes at most $k + (k - m + 1)$ intersections, so by

Lemma 4.3, a has strictly fewer than $k - m + 1$ positive edges. Again, the Babbitt et al. [2] bound counts a non-existent positive edge of a between a_ℓ and some other arc.

Finally, if radial rays from both endpoints of a intersect a_ℓ , then a and a_ℓ have at least two distinct regions of visibility either both through the center of the circle or not. Lemma 4.1 then implies that the edge between a and a_ℓ is either assigned to two different arcs.

We have shown that each arc in S_ℓ is associated with at least one over-counted visibility, either one that does not exist in the original graph or is counted twice in the original bound. Moreover, if $|S_\ell| < \ell$, then the original bound counts at least $\ell - |S_\ell|$ non-existent inner negative visibilities assigned to a_ℓ , since a radial ray extending inwards from the negative endpoint of a_ℓ encounters at most $|S_\ell| + k + 1 - \ell$ arcs, at most $|S_\ell|$ with smaller radius and $k + 1 - \ell$ with larger radius.

In any case, we obtain that a_ℓ is associated with an over-counting of at least ℓ visibilities. Moreover, each non-existent edge or double-counting of an existing edge is uniquely associated with a_ℓ , so we never account for these missing visibilities more than once as ℓ ranges from 1 to $k + 1$.

Repeating this process for all ℓ with $1 \leq \ell \leq k + 1$ yields a total over-count of $\sum_{\ell=1}^{k+1} \ell = (k + 1) \binom{k+2}{2}$. We therefore obtain that the maximum number of edges is

$$(k + 1) \left(3n - k - 2 - \frac{k + 2}{2} \right) = (k + 1) \left(3n - \frac{3k + 6}{2} \right),$$

as desired. □

Corollary 4.5. *The maximum number of edges in an arc visibility graph with n vertices is $\binom{n}{2}$ for $n \leq 5$ and $3n - 3$ for $n \geq 6$.*

Proof. This bound can be achieved as shown in Figure 10 with the dots indicating any necessary additional arcs. (If $n < 5$ take the innermost n arcs in Figure 10.) □




FIGURE 10. Arc visibility representation with maximum number of edges

4.2. A tight construction for semi-arc k -visibility graphs. As noted above, the edge bound previously given in Babbitt et al. [2] for semi-arc k -visibility graph is actually optimal. By establishing optimality, we disprove their Conjecture 20, which posited that a semi-arc k -visibility graph had at most $(k+1)(2n - \frac{3k+6}{2})$ edges for a semi-arc k -visibility graph with more than $3k+3$ vertices.

Theorem 4.6. *The maximum number of edges in a semi-arc k -visibility graph with n vertices is exactly $(k+1)(2n - \frac{k+2}{2})$ for $n \geq 5k+5$.*

Proof. The maximum number of edges in a semi-arc k -visibility graph with n vertices is at most $(k+1)(2n - \frac{k+2}{2})$ for $n \geq 3k+3$. This is Theorem 13 in Babbitt et al. [2]. We claim




FIGURE 11. Semi-arc k -visibility representation with $5k+5$ arcs and the maximum number of edges. Each set of arcs has $k+1$ arcs.

that the semi-arc k -visibility representation in Figure 11 corresponds to a graph with $n = 5k+5$ vertices and exactly $(k+1)(2n - \frac{k+2}{2})$ edges. Let the arcs be marked a_1, \dots, a_{5k+5} with indices increasing with increasing radius in Figure 9. The arcs a_1, \dots, a_{5k+5} have arguments of $\frac{\pi}{5}, \frac{\pi}{5} + \epsilon, \dots, \frac{\pi}{5} + k\epsilon, \frac{3\pi}{5}, \dots, \frac{3\pi}{5} + k\epsilon, \pi, \dots, \pi + k\epsilon, \frac{7\pi}{5}, \dots, \frac{7\pi}{5} + k\epsilon, \frac{9\pi}{5}, \dots, \frac{9\pi}{5} + k\epsilon$ radians respectively for ϵ sufficiently small. This semi-arc k -visibility representation gives a total of $(k+1)(2(5k+5) - \frac{k+2}{2})$ edges, as there are $5(k+1)^2$ edges corresponding to visibilities through the center and $(k+1)(\frac{9k+8}{2})$ edges corresponding to visibilities not through the center. This establishes the desired edge count for $n = 5k+5$. For $n > 5k+5$, add an additional $n - 5k - 5$ arcs between a_{3k+3} and a_{3k+4} that have an argument less than that of a_1 . Notice that each new arc adds $2k+2$ edges, so the bound is optimal for $n > 5k+5$. \square

Babbitt et al. [2] also conjectured that the complete graph K_{3k+4} is not a semi-arc k -visibility graph. Using a construction similar to the one given above, we disprove this conjecture as well.

Theorem 4.7. *K_{3k+4} is a semi-arc k -visibility graph.*

Proof. The construction is given in Figure 12. As radius increases, the arcs have arguments of $\frac{\pi}{3}, \frac{\pi}{3} - \epsilon, \dots, \frac{\pi}{3} - k\epsilon, \frac{2\pi}{3}, \dots, \frac{2\pi}{3} - k\epsilon, \pi, \dots, \pi + k\epsilon, \frac{5\pi}{3}$ radians for ϵ sufficiently small. \square




FIGURE 12. Semi-arc k -visibility representation of K_{3k+4} . Each set of arcs has $k + 1$ arcs.

Combining Theorem 13 in Babbitt et al. [2] with Theorems 4.6 and 4.7, we have a construction of a semi-arc k -visibility graph on n vertices with the maximum number of edges for $n \leq 3k + 4$ and $n \geq 5k + 5$. This leaves open the question of finding a maximal construction (or proving an improved bound) for $3k + 4 < n < 5k + 5$. When $k = 0$, there is no gap; when $k = 1$ the only open cases are $n = 8$ and $n = 9$.

5. THICKNESS BOUNDS

In this section, we prove new bounds on the thickness of arc and semi-arc k -visibility graphs.

Definition 5.1. The *thickness* of a graph G , denoted $\theta(G)$, is the smallest number of planar graphs into which the edges of G can be partitioned.

Bounding the thickness of bar k -visibility graphs has been a main subject of interest ever since their introduction by Dean et al. [5]. This quantity is especially relevant to VLSI design, where graphs of low thickness correspond to circuit designs that are electrically practical [16].

Computing the thickness of a graph is NP-hard in general [15], and exact thickness results are still open for all but a few classes of visibility graphs. Recently, Chang et al. [3] proposed using a simpler quantity, arboricity, to obtain easier bounds on thickness purely in terms of extant edge bounds. The results of this section use this strategy and the results of Sections 3 and 4 to prove new thickness bounds for arc and semi-arc k -visibility graphs.

We first review some basic facts about arboricity.

Definition 5.2. The *arboricity* of a graph G , denoted $\text{arb}(G)$, is the smallest number of forests into which the edges of a graph can be partitioned.

Theorem 5.3 (Nash-Williams [17]). *For any graph G ,*

$$\text{arb}(G) = \max_{H \subseteq G} \left\lceil \frac{E_H}{N_H - 1} \right\rceil,$$

where N_H and E_H are the number of vertices and edges respectively in the subgraph H .

Unlike thickness, arboricity can be computed in polynomial time [10]. Moreover, it differs from thickness by at most a constant factor.

Proposition 5.4 (Mutzel et al. [16]). *If a graph G has thickness Θ , then*

$$\Theta \leq \text{arb}(G) \leq 3\Theta.$$

Proof. Any partition of G into forests is *a fortiori* a partition into planar graphs, so $\Theta \leq \text{arb}(G)$. Theorem 5.3 implies that if H is planar, then $\text{arb}(H) \leq 3$, as any subgraph of a planar graph is planar, and a planar graph with n vertices has at most $3n - 6$ edges [6]. Hence a partition of G into Θ planar graphs can always be subdivided further into a partition of G into at most 3Θ forests. Hence $\text{arb}(G) \leq 3\Theta$. \square

Combining Theorem 5.3 with Theorem 4.4, we obtain the following theorem.

Theorem 5.5. *The thickness of an arc k -visibility graph is at most $3k + 3$.*

Proof. Let G be an arc k -visibility graph on n vertices, and let $H \subseteq G$ have ℓ vertices. Removing all arcs from the visibility representation of G except those corresponding to the vertices of H yields an arc k -visibility graph G' on ℓ vertices such that $H \subseteq G'$. Thus we can assume that H is a subgraph of a arc k -visibility graph with the same number of vertices.

By Theorem 4.4, H has at most $(k + 1)\left(3N_H - \frac{3k+6}{2}\right)$ edges if $\ell > 4k + 4$ and $\binom{N_H}{2}$ otherwise. In the former case,

$$E_H \leq (k + 1)\left(3N_H - \frac{3k + 6}{2}\right) = (3k + 3)\left(N_H - 1 - \frac{k}{2}\right) \leq (3k + 3)(N_H - 1).$$

In the latter case, $E_H = \frac{N_H}{2}(N_H - 1) \leq (2k + 2)(N_H - 1)$.

Theorem 5.3 and Proposition 5.4 then yield

$$\theta(G) \leq \text{arb}(G) = \max_{H \subseteq G} \left\lceil \frac{E_H}{N_H - 1} \right\rceil \leq 3k + 3$$

as desired. \square

Corollary 5.6. *The thickness of an arc visibility graph is at most 3.*


Note that Corollary 5.6 is stronger than what could have been obtained by applying the Nash-Williams Theorem to the bound of $3n - 2$ for arc visibility graphs proved by Babbitt et al. [2], which yields a maximum thickness of 4.

Applying the above strategy to semi-arc k -visibility graphs using the edge bound in Bab-bitt et al. [2] (which we showed to be tight in Section 4) shows that the thickness of these graphs is at most $2k + 2$ for $k \geq 2$. Using the classification of semi-arc visibility graphs given in Section 3, we can obtain a stronger statement.

Theorem 5.7. *The thickness of a semi-arc k -visibility graph is at most $2k + 1$.*

Proof. Fix a semi-arc k -visibility graph G and an associated representation. As in Section 4, we can assume that all arcs have radially distinct positive endpoints and distinct radii since we can achieve this by small perturbations without decreasing the thickness of the graph.

Given such a representation, call the semi-arc (0-)visibility graph associated with the collection of arcs SA_0 . Note that $SA_0 \subseteq G$. Moreover, by Corollary 3.9 SA_0 is planar.



(A) Visibilities in a sample semi-arc 1-visibility representation. Red segments indicated visibilities corresponding to the subgraph SA_0 . Blue segments indicate the remaining visibilities.

(B) Choice of orientation for blue edges.

FIGURE 13. Partition of edges described in Theorem 5.7

Remove the edges in SA_0 from G and call the remaining graph G' . For every pair of adjacent vertices G' , the line of sight with largest argument between the corresponding arcs contains one of their endpoints. Direct all edges in G' from the arc whose endpoint is contained in the corresponding line of sight to the one whose endpoint is not contained in the line of sight. (See Figure 13.) Each vertex in this graph has outdegree at most $2k$, so this graph can be partitioned into $2k$ disjoint graphs in which each vertex has outdegree at most one. Any component of such a graph has at most the same number of edges as vertices, hence is either a tree or a tree plus a single edge [6]. Therefore, each of these graphs is planar.

Thus the thickness of G' is at most $2k$. We obtain $\theta(G) \leq \theta(G') + \theta(SA_0) \leq 2k + 1$, as desired. \square

6. COMPARISON OF FAMILIES OF ARC k -VISIBILITY GRAPHS AND BAR k -VISIBILITY GRAPHS

In this section, we consider the relationship between bar and arc k -visibility graphs. We first show several structural properties of the family of bar k -visibility graphs and then use these results to show that arc i -visibility graphs are not bar j -visibility graphs in general. On the other hand, we note that bar j -visibility graphs are a subset of arc j -visibility graphs, since any bar visibility representation can easily be converted into a corresponding arc visibility representation. (This observation appears in Hutchinson [13].)

We begin by analyzing the families of bar i -visibility and j -visibility graphs for $i \neq j$. A result of Hartke et al. [12] establishes that these families are incomparable under set inclusion when $j = i + 1$. We use a similar construction to generalize their argument to all $i \neq j$.

We require one definition.

Definition 6.1. An *interval graph* is a graph corresponding to an arrangement of nonintersecting closed horizontal line segments in the plane. Each vertex of the graph corresponds to a bar, and two vertices are adjacent if and only if the corresponding bars are connected by a vertical line of sight passing through an unlimited number of bars.

Informally, interval graphs are generalizations of bar k -visibility graphs where $k = \infty$. Following our practice above, we call the collection of bars corresponding to an interval graph an *interval representation*. Interval graphs are well studied and have been completely characterized. In particular, it is known that all interval graphs are chordal (that is, contain no induced cycle of length more than three) [11].

The following theorem provides a precise connection between k -visibility graphs and interval graphs.

Theorem 6.2. *Let G be a K_{k+2} -free graph. Then G is an interval graph if and only if it is a bar k -visibility graph.*

Proof. Suppose that G is an interval graph, and fix an interval representation. Since G is K_{k+2} free, no vertical line intersects $k + 2$ bars. Any line of sight in the interval representation therefore passes through at most $(k + 2) - 2 = k$ intervening bars. Thus this set of bars is also a representation of G as a bar k -visibility graph.

Conversely, assume G is a bar k -visibility graph and fix a representation. If there existed a vertical line intersecting $k + 2$ bars, then the corresponding vertices would form a copy of K_{k+2} , since every pair of bars would be separated by at most $(k + 2) - 2 = k$ intervening bars. Therefore any pair of bars intersected by a vertical line are separated by at most k bars, and hence the corresponding vertices are adjacent in G . Thus this set of bars is also an interval representation for G as an interval graph. \square

Evans et al. observed [8, Lemma 1] that triangle-free bar 1-visibility graphs are forests. Theorem 6.2 implies the following stronger corollary.

Definition 6.3. A *caterpillar* is a tree in which all vertices are within one edge of a central path.

Corollary 6.4. *Let $k \geq 1$. If a bar k -visibility graph is triangle-free, then it is a disjoint union of caterpillars.*

Proof. Theorem 6.2 implies that a triangle-free bar k -visibility graph with $k \geq 1$ is an interval graph, and any triangle-free interval graph is a union of caterpillars [7]. \square

We can now state the main theorem.

Theorem 6.5. *Let B_k be the family of bar k -visibility graphs for $k \geq 0$. Then $B_i \not\subseteq B_j$ and $B_j \not\subseteq B_i$ for $i \neq j$.*

Proof. Without loss of generality let $j < i$. Consider the graph $K_j + C_4$, the graph on $j + 4$ vertices formed by taking the union of a complete graph on j vertices and a cycle graph on 4 vertices and adding all $4j$ edges between the two graphs (where C_n is the cycle on n vertices). Note that $K_j + C_4$ is K_{j+3} free and hence K_{i+2} free.

Since $K_j + C_4$ contains an induced four-cycle, it is not chordal and in particular is not an interval graph. Theorem 6.2 then implies $K_j + C_4$ is not a bar i -visibility graph. However, Figure 14 shows that it is a bar j -visibility graph. Hence $B_i \not\subseteq B_j$.




FIGURE 14. Bar j -visibility representation of $K_j + C_4$

On the other hand, Hartke et al. [12] show that $K_{4j+4} \notin B_j$ and $K_{4j+4} \in B_i$. Therefore $B_j \not\subseteq B_i$. \square

Let A_k be the family of arc k -visibility graphs. Since $B_k \subset A_k$ for all k , Theorem 6.5 implies that $A_j \not\subseteq B_i$ for all $i \neq j$. A more careful analysis shows that in fact this claim holds for all i and j .

Theorem 6.6. *Let B_k be the family of bar k -visibility graphs and A_k be the family of arc k -visibility graphs. Then $A_j \not\subseteq B_i$ for all $i, j \geq 0$.*

Proof. Fix a nonnegative j . We first show that $A_j \not\subseteq B_0$. Figure 15 shows that K_5 is an arc j -visibility graph for any $j \geq 0$. Since all bar visibility graphs are planar, this implies $A_j \not\subseteq B_0$.

Suppose now that $i \geq 1$. The cycle graph C_4 is triangle free, so Theorem 6.2 implies that it is a bar i -visibility graph if and only if it is an interval graph. Since C_4 is not chordal, we conclude that C_4 is not a bar i -visibility graph for all $i \geq 1$. But Figure 16 shows it is an arc j -visibility graph. Therefore $A_j \not\subseteq B_i$, as desired. \square




FIGURE 15. Arc j -visibility representation of K_5




FIGURE 16. Arc j -visibility representation of C_4

We note that the analogous questions for semi-bar k -visibility graphs and semi-arc k -visibility graphs are far simpler. The family of semi-bar or semi-arc i -visibility graphs is never contained in the family of semi-bar or semi-arc j -visibility graphs for $i \neq j$ because a semi-bar or semi-arc k -visibility graph on n vertices has at least $(k+1)n - O(1)$ edges and at most $2(k+1)n + O(1)$ edges.

7. ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support of MIT PRIMES Program, especially the guidance of Tanya Khovanova, and the assistance of Jesse Geneson in suggesting this problem. The first author would also like to express his gratitude to his research teachers Richard Kurtz, Jeanette Collette, Dr. Lorraine Solomon, Dr. Daniel Kramer, and especially the second author for their guidance. The second author would like to acknowledge the generous support of the NSF Graduate Research Fellowship under Grant No. DGE-1122374.

REFERENCES

- [1] Thomas Andreae. Some results on visibility graphs. *Discrete Appl. Math.*, 40(1):5–17, 1992. Combinatorial methods in VLSI.
- [2] Matthew Babbitt, Jesse Geneson, and Tanya Khovanova. On k -visibility graphs. *J. Graph Algorithms Appl.*, 19(1):345–360, 2015.
- [3] Hsien-Chih Chang, Hsueh-I Lu, and Yen-Peng Sung. Asymptotically optimal thickness bounds of generalized bar visibility graphs. unpublished manuscript, 2009.
- [4] F. J. Cobos, J. C. Dana, F. Hurtado, A. Márquez, and F. Mateos. On a visibility representation of graphs. In *Graph drawing (Passau, 1995)*, volume 1027 of *Lecture Notes in Comput. Sci.*, pages 152–161. Springer, Berlin, 1996.
- [5] Alice M. Dean, William Evans, Ellen Gethner, Joshua D. Laison, Mohammad Ali Safari, and William T. Trotter. Bar k -visibility graphs. *J. Graph Algorithms Appl.*, 11(1):45–59, 2007.
- [6] Reinhard Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, Heidelberg, fourth edition, 2010.
- [7] Jürgen Eckhoff. Extremal interval graphs. *J. Graph Theory*, 17(1):117–127, 1993.
- [8] William Evans, Michael Kaufmann, William Lenhart, Tamara Mchedlidze, and Stephen Wismath. Bar 1-visibility graphs and their relation to other nearly planar graphs. *J. Graph Algorithms Appl.*, 18(5):721–739, 2014.
- [9] Stefan Felsner and Mareike Massow. Parameters of bar k -visibility graphs. *J. Graph Algorithms Appl.*, 12(1):5–27, 2008.

- [10] Harold N. Gabow and Herbert H. Westermann. Forests, frames, and games: algorithms for matroid sums and applications. *Algorithmica*, 7(5-6):465–497, 1992.
- [11] P. C. Gilmore and A. J. Hoffman. A characterization of comparability graphs and of interval graphs. *Canad. J. Math.*, 16:539–548, 1964.
- [12] Stephen G. Hartke, Jennifer Vandenbussche, and Paul Wenger. Further results on bar k -visibility graphs. *SIAM J. Discrete Math.*, 21(2):523–531 (electronic), 2007.
- [13] Joan P. Hutchinson. Arc- and circle-visibility graphs. *Australas. J. Combin.*, 25:241–262, 2002.
- [14] F. Luccio, S. Mazzone, and C. K. Wong. A note on visibility graphs. *Discrete Math.*, 64(2-3):209–219, 1987.
- [15] Anthony Mansfield. Determining the thickness of graphs is NP-hard. *Math. Proc. Cambridge Philos. Soc.*, 93(1):9–23, 1983.
- [16] Petra Mutzel, Thomas Odenthal, and Mark Scharbrodt. The thickness of graphs: a survey. *Graphs Combin.*, 14(1):59–73, 1998.
- [17] C. St. J. A. Nash-Williams. Decomposition of finite graphs into forests. *J. London Math. Soc.*, 39:12, 1964.
- [18] Roberto Tamassia and Ioannis G. Tollis. A unified approach to visibility representations of planar graphs. *Discrete Comput. Geom.*, 1(4):321–341, 1986.

STUDENT BIOGRAPHIES

Mehtaab Sawhney: (*Corresponding author:* msawhney@mit.edu) Mehtaab Sawhney is a sophomore studying mathematics at MIT. During Summer 2017, he completed the REU at the University of Minnesota at Duluth led by Professor Gallian focusing on problems with a combinatorial flavor.