

On the Connectivity of Visibility Graphs

Michael S. Payne · Attila Pór · Pavel Valtr ·
David R. Wood

Received: 25 August 2011 / Revised: 16 April 2012 / Accepted: 5 July 2012 /
Published online: 18 July 2012
© Springer Science+Business Media, LLC 2012

Abstract The visibility graph of a finite set of points in the plane has the points as vertices and an edge between two vertices if the line segment between them contains no other points. This paper establishes bounds on the edge- and vertex-connectivity of visibility graphs.

Unless all its vertices are collinear, a visibility graph has diameter at most 2, and so it follows by a result of Plesník (Acta Fac. Rerum Nat. Univ. Comen. Math. 30:71–93, 1975) that its edge-connectivity equals its minimum degree. We strengthen the result of Plesník by showing that for any two vertices v and w in a graph of diameter 2, if $\deg(v) \leq \deg(w)$ then there exist $\deg(v)$ edge-disjoint vw -paths of length at most 4.

For vertex-connectivity, we prove that every visibility graph with n vertices and at most ℓ collinear vertices has connectivity at least $\frac{n-1}{\ell-1}$, which is tight. We also prove the qualitatively stronger result that the vertex-connectivity is at least half the minimum degree. Finally, in the case that $\ell = 4$ we improve this bound to two thirds of the minimum degree.

Keywords Visibility graph · Connectivity · Bivisibility graph · Elliptic curve

M.S. Payne (✉) · D.R. Wood
Department of Mathematics and Statistics, The University of Melbourne, Melbourne, Australia
e-mail: m.payne3@pgrad.unimelb.edu.au

D.R. Wood
e-mail: woodd@unimelb.edu.au

A. Pór
Department of Mathematics, Western Kentucky University, Bowling Green, KY, USA
e-mail: attila.por@wku.edu

P. Valtr
Department of Applied Mathematics and Institute for Computer Science (ITI), Charles University,
Prague, Czech Republic

1 Introduction

Let P be a finite set of points in the plane. Two distinct points v and w in the plane are *visible* with respect to P if no point in P is in the open line segment vw . The *visibility graph* of P has vertex set P , and an edge between two vertices if and only if they are visible with respect to P .

Visibility graphs have interesting graph-theoretic properties and give rise to several challenging open problems. For example, Kára, Pór and Wood [5] showed that K_4 -free visibility graphs are 3-colourable. They also conjectured that visibility graphs with bounded collinearities and bounded clique number have a bounded number of vertices. See [5, 10, 13] for further results and conjectures about the clique and chromatic number of visibility graphs. Further related results can be found in [1, 3, 8]. The study of visibility graphs is a natural extension of point-line incidence geometry where the order of points on a line is taken into account. Visibility graphs appear implicitly in Székely’s celebrated paper [14]. For example, Székely’s proof of the Szemerédi–Trotter theorem could be summarised as, “apply the crossing lemma to a certain subgraph of the visibility graph”. We therefore consider the fundamental graph-theoretic properties of visibility graphs to be worthy of study.

The purpose of this paper is to study the edge- and vertex-connectivity of visibility graphs. A graph G on at least $k + 1$ vertices is k -vertex-connected (k -edge-connected) if G remains connected whenever fewer than k vertices (edges) are deleted. Menger’s theorem says that this is equivalent to the existence of k vertex-disjoint (edge-disjoint) paths between each pair of vertices. Let $\kappa(G)$ and $\lambda(G)$ denote the vertex- and edge-connectivity of a graph G . Let $\delta(G)$ denote the minimum degree of G . Then $\kappa(G) \leq \lambda(G) \leq \delta(G)$.

If a visibility graph G has n vertices, at most ℓ of which are collinear, then $\delta(G) \geq \frac{n-1}{\ell-1}$. We will show that both edge- and vertex-connectivity are at least $\frac{n-1}{\ell-1}$ (Theorem 4 and Corollary 11). Since there are visibility graphs with $\delta = \frac{n-1}{\ell-1}$ these lower bounds are best possible.

We will refer to visibility graphs whose vertices are not all collinear as *non-collinear visibility graphs*. Non-collinear visibility graphs have diameter 2 [5], and it is known that graphs of diameter 2 have edge-connectivity equal to their minimum degree [11]. We strengthen this result to show that if a graph has diameter 2 then for all vertices v and w with $\deg(v) \leq \deg(w)$, there are $\deg(v)$ edge-disjoint vw -paths of length at most 4 (Theorem 2).

With regard to vertex-connectivity, our main result is that $\kappa \geq \frac{\delta}{2} + 1$ for all non-collinear visibility graphs (Theorem 9). This bound is qualitatively stronger than the bound $\kappa \geq \frac{n-1}{\ell-1}$ since it is always within a factor of 2 of being optimal. In the special case of at most four collinear points, we improve this bound to $\kappa \geq \frac{2\delta+1}{3}$ (Theorem 13). We conjecture that $\kappa \geq \frac{2\delta+1}{3}$ for all visibility graphs.

This bound would be best possible since, for each integer k , there is a visibility graph with a vertex cut of size $2k + 1$, but minimum degree $\delta = 3k + 1$. Therefore the vertex-connectivity is at most $2k + 1 = \frac{2\delta+1}{3}$. Figure 1 shows the case $k = 4$.

A central tool in this paper, which is of independent interest, is a kind of bipartite visibility graph. Let A and B be disjoint sets of points in the plane. The *bivisibility graph* $\tilde{B}(A, B)$ of A and B has vertex set $A \cup B$, where points $v \in A$ and $w \in B$ are




Fig. 1 A visibility graph with vertex-connectivity $\frac{2\delta+1}{3}$. The black vertices are a cut set. The minimum degree $\delta = 3k + 1$ is achieved, for example, at the top left vertex. Not all edges are drawn

adjacent if and only if they are visible with respect to $A \cup B$. The following simple observation is used several times in this paper.

Observation 1 *Let G be a visibility graph. Let $\{A, B, C\}$ be a partition of $V(G)$ such that C separates A and B . If $\mathcal{B}(A, B)$ contains t pairwise non-crossing edges, then $|C| \geq t$ since there must be a distinct vertex in C on each such edge.*

2 Edge-Connectivity

Non-collinear visibility graphs have diameter at most 2 [5]. This is because even if two points cannot see each other, they can both see the point closest to the line containing them. Plesník [11] proved that the edge-connectivity of a graph with diameter at most 2 equals its minimum degree. Thus the edge-connectivity of a non-collinear visibility graph equals its minimum degree. There are several other known conditions that imply that the edge-connectivity of a graph is equal to the minimum degree; see for example [2, 12, 16]. Here we prove the following strengthening of the result of Plesník.

Theorem 2 *Let G be a graph with diameter 2. Then the edge-connectivity of G equals its minimum degree. Moreover, for all distinct vertices v and w in G , if $d := \min\{\deg(v), \deg(w)\}$ then there are d edge-disjoint vw -paths of length at most 4, including at least one of length at most 2.*

Proof First suppose that v and w are not adjacent. Let C be the set of common neighbours of v and w . For each vertex $c \in C$, take the path (v, c, w) . Let A be a set of $d - |C|$ neighbours of v not in C . Let B be a set of $d - |C|$ neighbours of w not in C . Let M_1 be a maximal matching in the bipartite subgraph of G induced by A and B . Call these matched vertices A_1 and B_1 . For each edge $ab \in M_1$, take the path (v, a, b, w) . Let A_2 and B_2 , respectively, be the subsets of A and B consisting of the unmatched vertices. Let $D := V(G) \setminus (A_2 \cup B_2 \cup \{v, w\})$. Let M_2

be an arbitrary pairing of vertices in A_2 and B_2 . For each pair $ab \in M_2$, take the path (v, a, x, b, w) , where x is a common neighbour of a and b (which exists since G has diameter 2). Since x is adjacent to a , $x \neq w$, and by the maximality of M_1 , $x \notin B_2$. Similarly, $x \neq v$ and $x \notin A_2$, and so $x \in D$. Thus there are three types of path, namely (v, C, w) , (v, A_1, B_1, w) , and (v, A_2, D, B_2, w) . Paths within each type are edge-disjoint. Even though D contains A_1 and B_1 , edges between each pair of sets from $\{A_1, B_1, A_2, B_2, C, D, \{v\}, \{w\}\}$ occur in at most one of the types, and all edges are between distinct sets from this collection. Hence no edge is used twice, so all the paths are edge-disjoint. The total number of paths is $|C| + |A_1| + |A_2| = d$. This finishes the proof if v and w are not adjacent. If G does contain the edge vw then take this as the first path, then remove it and find $d - 1$ paths in the same way as above. \square

Note that the lengths of the paths found in Theorem 2 cannot be improved, as shown by the following example. For integers $\gamma \geq 1$ and $\delta \geq 3$, let G be the graph obtained from a 5-cycle (v, w, x, y, z) by replacing x by a $(\delta - 1)$ -clique X , replacing y by a γ -clique Y , replacing z by a $(\delta - 1)$ -clique Z , and replacing edges between these vertices with complete bipartite subgraphs. Each vertex in X is adjacent to w and to each vertex in Y . Each vertex in Z is adjacent to v and to each vertex in Y . Thus G has minimum degree δ and diameter 2. Note that $\deg(v) = \deg(w) = \delta$. In fact, by choosing γ large, we can make v and w the only vertices of degree δ and every other vertex have arbitrarily large degree. Consider a set S of δ edge-disjoint paths between v and w . One path in S is the edge vw , while every other path has length at least 4. Thus the paths found in Theorem 2 are best possible. A further example, in which u and v are not adjacent, can be constructed by taking two disjoint $(\delta + 1)$ -cliques and identifying a vertex from each. Suppose u and v come from different cliques and are not the identified vertex. Then there are $\delta - 1$ vw -paths of length 4 and one of length 2. Alternatively, one can take $\delta - 2$ vw -paths of length 4 and two of length 3.

Theorem 2 implies the following corollary for visibility graphs.

Corollary 3 *Let G be a non-collinear visibility graph. Then the edge-connectivity of G equals its minimum degree. Moreover, for distinct vertices v and w , there are $\min\{\deg(v), \deg(w)\}$ edge-disjoint vw -paths of length at most 4, including at least one of length at most 2.*

The full version of this paper [9] extends Corollary 3 to prove that every minimum edge-cut in a non-collinear visibility graph is the set of edges incident to some vertex.


We now show that not only is the edge-connectivity as high as possible, but it is realised by paths with at most one bend.

Theorem 4 *Let G be a visibility graph with n vertices, at most ℓ of which are collinear. Then G is $\lceil \frac{n-1}{\ell-1} \rceil$ -edge-connected, which is best possible. Moreover, between each pair of vertices, there are $\lceil \frac{n-1}{\ell-1} \rceil$ edge-disjoint 1-bend paths.*

Proof Let v and w be distinct vertices of G . Let V^* be the set of vertices of G not on the line vw . Let $m := |V^*|$. Thus $m \geq n - \ell$.

Let \mathcal{L} be the pencil of lines through v and the vertices in V^* . Let \mathcal{M} be the pencil of lines through w and the vertices in V^* . Let H be the bipartite graph with vertex

Fig. 2 (a) If each ray from v through $V(G)$ contains ℓ vertices, the degree of v is $\frac{n-1}{\ell-1}$. (b) Two properly coloured non-crossing geometric graphs with no black-white edge between them



set $\mathcal{L} \cup \mathcal{M}$, where $L \in \mathcal{L}$ is adjacent to $M \in \mathcal{M}$ if and only if $L \cap M$ is a vertex in V^* .

Thus H has m edges, and maximum degree at most $\ell - 1$. Hence, by König’s theorem [6], H is $(\ell - 1)$ -edge-colourable. Thus H contains a matching of at least $\frac{m}{\ell-1}$ edges. This matching corresponds to a set S of at least $\frac{m}{\ell-1}$ vertices in V^* , no two of which are collinear with v or w .

For each vertex $x \in S$, take the path in the visibility graph from v straight to x and then straight to w . These paths are edge-disjoint. Adding the path straight from v to w , we get at least $\frac{m}{\ell-1} + 1$ paths, which is at least $\frac{n-1}{\ell-1}$. Figure 2(a) shows that this bound is best possible. □

3 A Key Lemma

We call a plane graph drawn with straight edges a *non-crossing geometric graph*. The following interesting fact about non-crossing geometric graphs will prove useful. It says that, except for some degenerate cases, two properly coloured non-crossing geometric graphs that are separated by a line can be joined by an edge such that the union is a properly coloured non-crossing geometric graph. Note that this is false if the two graphs are not separated by a line, as demonstrated by the example in Fig. 2(b).

Lemma 5 *Let G_1 and G_2 be two properly coloured non-crossing geometric graphs with at least one edge each. Suppose their convex hulls are disjoint and that $V(G_1) \cup V(G_2)$ is not collinear. Then there exists an edge $e \in V(G_1) \times V(G_2)$ such that $G_1 \cup G_2 \cup \{e\}$ is a properly coloured non-crossing geometric graph.*

Proof Let h be a line separating G_1 and G_2 . Assume that h is vertical with G_1 to the left. Let $G := G_1 \cup G_2$.

Call a pair of vertices $v_1 \in V(G_1)$ and $v_2 \in V(G_2)$ a *visible pair* if the line segment between them does not intersect any vertices or edges of G . We aim to find a visible pair with different colours, so assume for the sake of contradiction that every visible pair is monochromatic.

We may assume that G_1 and G_2 are edge maximal with respect to the colouring, since the removal of an edge only makes it easier to find a bichromatic visible pair.

Suppose the result holds when there are no isolated vertices in G . Then, if there are isolated vertices, we can ignore them and find a bichromatic visible pair (v_1, v_2)




Fig. 3 Proof of Lemma 5. The shaded areas are empty. (a) A type-1 visible pair. (b) A type-2 visible pair. (c) The highest visible pair. (d) The lowest pair is type-1. (e) The lowest pair is type-2

in the remaining graph. If the edge $v_1 v_2$ contains some of the isolated vertices, then it has a sub-segment joining two vertices of different colours. If these vertices lie on the same side of h then the graphs were not edge maximal after all. If they are on different sides, then they are a bichromatic visible pair. Thus we may assume that there are no isolated vertices in G .

Let l be the line containing a visible pair (v_1, v_2) , then the *height* of the pair is the point at which l intersects h . Call the pair *type-1* if v_1 and v_2 both have a neighbour strictly under the line l (Fig. 3(a)). Call the pair *type-2* if there are edges $v_1 w_1$ in G_1 and $v_2 w_2$ in G_2 such that the line g containing $v_1 w_1$ intersects $v_2 w_2$ (call this point x), w_2 lies strictly under g , and the closed triangle $v_1 v_2 x$ contains no other vertex (Fig. 3(b)). Here x may equal v_2 , in which case $g = l$. A visible pair is also type-2 in the equivalent case with the subscripts interchanged.

A particular visible pair may be neither type-1 nor type-2, but we may assume there exists a type-1 or type-2 pair. To see this, consider the highest visible pair (v_1, v_2) and assume it is neither type-1 nor type-2 (see Fig. 3(c)). Note that $v_1 v_2$ is an edge of the convex hull of G . Since all of G lies on or below the line l containing $v_1 v_2$, both vertices must have degree 1 and their neighbours w_1 and w_2 must lie on l . For $i = 1, 2$, let x_i be a vertex of G_i not on l that minimizes the angle $\angle v_i w_i x_i$. Since $V(G)$ is not collinear, at least one of x_i exists. By symmetry, we may assume that either only x_1 exists, or both x_1 and x_2 exist and $\text{dist}(x_1, l) \leq \text{dist}(x_2, l)$. In either case, (x_1, v_2) and (x_1, w_2) are visible pairs and at least one of them is bichromatic.

So now assuming there exists a type-1 or type-2 visible pair, let (u_1, u_2) be the lowest such pair:

Case (i) The pair (u_1, u_2) is type-1 (see Fig. 3(d)). Let $u_1 w_1$ be the first edge of G_1 incident to u_1 in a clockwise direction, starting at $u_1 u_2$. Let $u_2 w_2$ be the first edge of G_2 incident to u_2 in a counterclockwise direction, starting at $u_2 u_1$. Let x be the point on the segment $u_1 w_1$ closest to w_1 such that the open triangle $u_1 u_2 x$ is disjoint from G . Similarly, let y be the point on the segment $u_2 w_2$ closest to w_2 such that the open triangle $u_1 u_2 y$ is disjoint from G .

Without loss of generality, the intersection of $u_1 y$ and $u_2 x$ is to the left of h , or on h . Therefore the segment $x u_2$ is disjoint from G_2 . Let $v \in V(G_1)$ be the vertex on $x u_2$ closest to u_2 . Thus (v, u_2) is a visible pair of height less than (u_1, u_2) . We may assume that $v \neq w_1$, otherwise (v, u_2) would be bichromatic. The point w_2 is under the line $v u_2$ and v has no neighbour above the line $v u_2$. Hence v either has

a neighbour under the line vu_2 and (v, u_2) is type-1, or v has a neighbour on the line vu_2 and (v, u_2) is type-2. This contradicts the assumption that (u_1, u_2) was the lowest pair of either type.

Case (ii) The pair (u_1, u_2) is type-2 with neighbours w_1 and w_2 , such that the line u_2w_2 intersects the edge u_1w_1 at some point x (see Fig. 3(e)). Let y be the point on the segment u_1w_1 closest to w_1 such that the open triangle u_1u_2y is disjoint from G . Note y is below x by the definition of type-2.

First assume that G_2 intersects yu_2 . Let v_2 be the closest vertex to u_2 on yu_2 . Thus (u_1, v_2) is a visible pair of height less than (u_1, u_2) . Let z be a neighbour of v_2 . If z is under the line u_1v_2 then (u_1, v_2) is type-1 since w_1 is also under this line. Note z cannot lie above the line yu_2 since u_1 and u_2 see each other and the open triangle u_1u_2y is empty. Furthermore, if z lies on yu_2 then $z = u_2$ and (u_1, v_2) is bichromatic. Thus if z is not under the line u_1v_2 , the line v_2z must intersect the edge u_1w_1 at a point above y , so (u_1, v_2) is a type-2 pair. Hence the pair (u_1, v_2) is type-1 or type-2, a contradiction.

Now assume that yu_2 does not intersect G_2 , and therefore does intersect G_1 , and let $v_1 \in V(G_1)$ be the vertex on yu_2 closest to u_2 . Thus (v_1, u_2) is a visible pair of height less than (u_1, u_2) . We may assume that $v_1 \neq w_1$, otherwise (v_1, u_2) would be bichromatic. Since w_2 is under the line v_1u_2 , if v_1 has a neighbour under the line v_1u_2 then (v_1, u_2) is a type-1 pair. Otherwise the only neighbour of v_1 is on the line v_1u_2 which makes (v_1, u_2) a type-2 pair. Hence the pair (v_1, u_2) is type-1 or type-2, a contradiction. \square

4 Vertex-Connectivity

As is common practice, we will often refer to vertex-connectivity simply as connectivity. Connectivity of visibility graphs is not as straightforward as edge-connectivity since there are visibility graphs with connectivity strictly less than the minimum degree (see Fig. 1). Our aim in this section is to show that the connectivity of a visibility graph is at least half the minimum degree (Theorem 9). This follows from Theorem 8, which says that bivisibility graphs contain large non-crossing subgraphs. In the proof of Theorem 8 we will need a version of the Ham Sandwich Theorem for point sets in the plane, and also Lemma 7.

Theorem 6 (Ham Sandwich. See [7]) *Let A and B be finite sets of points in the plane. Then there exists a line h such that each closed half-plane determined by h contains at least half of the points in A and at least half of the points in B .*

Lemma 7 *Let A be a set of points lying on a line l . Let B be a set of points, none of them lying on l . Let $|A| \geq |B|$. Then there is a non-crossing spanning tree in the bivisibility graph of A and B .*

Proof We proceed by induction on $|B|$. If $|B| = 1$ then the point in B sees every point in A , and we are done. Now assume $1 < |B| \leq |A|$.

First suppose that all of B lies to one side of l and consider the convex hull C of $A \cup B$. An end point a of A is a corner of C and there is a point b of B visible

to it in the boundary of C . There exists a line h that separates $\{a, b\}$ from the rest of $A \cup B$. Applying induction and Lemma 5 we find a non-crossing spanning tree among $A \cup B \setminus \{a, b\}$ and an edge across h to the edge ab , giving a non-crossing spanning tree of $\mathcal{B}(A, B)$.

Now suppose that there are points of B on either side of l . Then we may apply the inductive hypothesis on each side to obtain two spanning trees. Their union is connected, and thus contains a spanning tree. \square

Theorem 8 *Let A and B be disjoint sets of points in the plane with $|A| = |B| = n$ such that $A \cup B$ is not collinear. Then the bivisibility graph $\mathcal{B}(A, B)$ contains a non-crossing subgraph with at least $n + 1$ edges.*

Proof We proceed by induction on n . The statement holds for $n = 1$, since no valid configuration exists. For $n = 2$, any triangulation of $A \cup B$ contains at least five edges. At most one edge has both endpoints in A , and similarly for B . Removing these edges, we obtain a non-crossing subgraph of $\mathcal{B}(A, B)$ with at least three edges. Now assume $n > 2$.

Case (i) First suppose that there exists a line l that contains at least n points of $A \cup B$. Let $A_0 := A \cap l$, $B_0 := B \cap l$, $A_1 := A \setminus l$ and $B_1 := B \setminus l$. Without loss of generality, $|A_0| \geq |B_0|$.

If $|A_0| > |B_0|$ then $|A_0| + |B_1| > |B_0| + |B_1| = n$. Since $|A_0| + |B_0| \geq n = |B_1| + |B_0|$ we have $|A_0| \geq |B_1|$, so we may apply Lemma 7 to A_0 and B_1 . We obtain a non-crossing subgraph of $\mathcal{B}(A, B)$ with $|A_0| + |B_1| - 1 \geq n$ edges, and by adding an edge along l if needed, we are done.

Now assume $|A_0| = |B_0|$. We apply Lemma 7 to A_0 and B_1 , obtaining a non-crossing subgraph with $n - 1$ edges, to which we may add one edge along l . We still need one more edge. Suppose first that one open half-plane determined by l contains points of both A_1 and B_1 . Let a and b be the furthest points of A_1 and B_1 from l in this half-plane. Since $|A_0| = |B_0|$ we may assume that a is at least as far from l as b . Then we may add an edge along the segment ab , because none of the edges from A_0 to B_1 cross it. It remains to consider the case where l separates A_1 from B_1 . Then applying Lemma 7 on each side of l we find a non-crossing subgraph with $2n - 1$ edges: $|A_0| + |B_1| - 1$ on one side, $|B_0| + |A_1| - 1$ on the other side, and one more along l .

Case (ii) Now assume that no line contains n points in $A \cup B$. By Theorem 6 there exists a line h such that each of the closed half-planes determined by h contains at least $\frac{n}{2}$ points from each of A and B . Assume that h is horizontal. Let A^+ be the points of A that lie above h along with any that lie on h that we choose to assign to A^+ . Define A^- , B^+ and B^- in a similar fashion. Now assign the points on h to these sets so that each has exactly $\lceil \frac{n}{2} \rceil$ points. In particular, assign the required number of *leftmost* points of $h \cap A$ to A^+ and *rightmost* points of $h \cap A$ to A^- . Do the same for $h \cap B$ with left and right interchanged. If n is even then $A^+ \cup A^-$ and $B^+ \cup B^-$ are partitions of A and B . If n is odd then $|A^+ \cap A^-| = |B^+ \cap B^-| = 1$.

Since there is no line containing n points of $A \cup B$, the inductive hypothesis may be applied on either side of h . Thus there is a non-crossing subgraph with $\lceil \frac{n}{2} \rceil + 1$ edges on each side. The union of these subgraphs has at least $n + 2$ edges, but some

edges along h may overlap. Due to the way the points on h were assigned, one of the subgraphs has at most one edge along h . (If n is odd, this is the edge between the two points that get assigned to both sides.) Deleting this edge from the union yields a non-crossing subgraph of $\mathcal{B}(A, B)$ with at least $n + 1$ edges. \square

Theorem 9 *Every non-collinear visibility graph with minimum degree δ has connectivity at least $\frac{\delta}{2} + 1$.*

Proof Suppose $\{A, B, C\}$ is a partition of the vertex set of a non-collinear visibility graph such that C separates A and B , and $|A| \leq |B|$. By considering a point in A we see that $\delta \leq |A| + |C| - 1$. By removing points from B until $|A| = |B|$ whilst ensuring that $A \cup B$ is not collinear, we may apply Theorem 8 and Observation 1 to get $|C| \geq |A| + 1$. Combining these inequalities yields $|C| \geq \frac{\delta}{2} + 1$. \square

5 Vertex-Connectivity with Bounded Collinearities

For the visibility graphs of point sets with n points and at most ℓ collinear, connectivity is at least $\frac{n-1}{\ell-1}$, just as for edge-connectivity. Bivisibility graphs will play a central role in the proof of this result. For point sets A and B an *AB-line* is a line containing points from both sets.


Theorem 10 *Let $A \cup B$ be a non-trivial partition of a set of n points with at most ℓ on any AB -line. Then the bivisibility graph $\mathcal{B}(A, B)$ contains a non-crossing forest with at least $\frac{n-1}{\ell-1}$ edges. In particular, if $\ell = 2$ then the forest is a spanning tree.*

Proof The idea of the proof is to cover the points of $A \cup B$ with a large set of disjoint line segments each containing an edge of $G := \mathcal{B}(A, B)$. Start with a point $v \in A$. Consider all open ended rays starting at v and containing a point of B . Each such ray contains at least one edge of G and at most $\ell - 1$ points of $(A \cup B) \setminus v$. For each ray r , choose a point $w \in B \cap r$. Draw all maximal line segments with an open end at w and a closed end at a point of A in the interior of the sector clockwise from r . Figure 4 shows an example. If one sector S has central angle larger than π then some points of A may not be covered. In this case we bisect S , and draw segments from each of its bounding rays into the corresponding half of S (assign points on the bisecting line to one sector arbitrarily). Like the rays, these line segments all contain at least one edge of G and at most $\ell - 1$ points of $(A \cup B) \setminus \{v, w\}$. Together with the rays, they are pairwise disjoint and cover all of $(A \cup B) \setminus v$. Hence the edges of G contained in them form a non-crossing forest with at least $\frac{n-1}{\ell-1}$ edges. Note that if $\ell = 2$ we have a forest with $n - 1$ edges, hence a spanning tree. \square

Note that the $\ell = 2$ case of Theorem 10 is well known [4].

Corollary 11 *Let G be the visibility graph of a set of n points with at most ℓ collinear. Then G has connectivity at least $\frac{n-1}{\ell-1}$, which is best possible.*

Fig. 4 Covering $A \cup B$ with rays and segments (a), each of which contains an edge of the bivisibility graph (b)



Proof Let $\{A, B, C\}$ be a partition of $V(G)$ such that C separates A and B . Consider the bivisibility graph of $A \cup B$. Applying Observation 1 and Theorem 10 (with $n' = n - |C|$ and $\ell' = \ell - 1$) yields $|C| \geq \frac{n-|C|-1}{\ell-2}$, which implies $|C| \geq \frac{n-1}{\ell-1}$. As in the case of edge-connectivity, the example in Fig. 2(a) shows that this bound is best possible. \square

The following theorem yields an improvement of Theorem 9 for visibility graphs with at most four collinear vertices. Lemma 5 is an important tool in the proof.

Theorem 12 *Let A and B be disjoint point sets in the plane with $|A| = |B| = n$ such that $A \cup B$ has at most three points on any AB -line. Then the bivisibility graph $\mathcal{B}(A, B)$ contains a non-crossing spanning tree.*

Proof We proceed by induction on n . The statement is true for $n = 1$. Apply Theorem 6 to find a line h such that each closed half-plane defined by h has at least $\frac{n}{2}$ points from each of A and B . Assume that h is horizontal. The idea of the proof is to apply induction on each side of h to get two spanning trees, and then find an edge joining them together. In most cases the joining edge will be found by applying Lemma 5.


We will construct a set A^+ containing the points of A that lie above h along with any that lie on h that we choose to assign to A^+ . We will also construct A^-, B^+ and B^- in a similar fashion. By the properties of h , there exists an assignment¹ of each point in $h \cap (A \cup B)$ to one of these sets such that $|A^+| = |B^+| = \lceil \frac{n}{2} \rceil$ and $|A^-| = |B^-| = \lfloor \frac{n}{2} \rfloor$.

Consider the sequence s_h of signs (+ or -) given by the chosen assignment of points on h from left to right. If s_h is all the same sign, or alternates only once from one sign to the other, then it is possible to perturb h to h' so that $A^+ \cup B^+$ lies strictly above h' and $A^- \cup B^-$ lies strictly below h' . Thus we may apply induction on each side to obtain non-crossing spanning trees in $\mathcal{B}(A^+, B^+)$ and $\mathcal{B}(A^-, B^-)$. Then apply Lemma 5 to find an edge between these two spanning trees, creating a non-crossing spanning tree of $\mathcal{B}(A, B)$.

Otherwise, s_h alternates at least twice (so there are at least three points on h). This need never happen if there are only points from one set on h , since the points required

¹We need only consider one of the sets, say A . Say there are x points above h , y points on h and z points below h . Then $x + y \geq \lceil n/2 \rceil \geq \lfloor n/2 \rfloor \geq z$ so we can ensure $|A^+| = \lceil n/2 \rceil$. A^- is the complement and therefore has $\lfloor n/2 \rfloor$ points.

Fig. 5 The only case in which h may not be perturbed to separate the points assigned *above* h from those assigned *below*



above h can be taken from the left and those required below h from the right. Without loss of generality, the only remaining case to consider is that h contains one point from A and two from B . If the two points from B are consecutive on h , then without loss of generality $s_h = (+, -, +)$ and the points of B are on the left. In this case the signs of the points from B may be swapped so s_h becomes $(-, +, +)$. If the point from A lies between the other two points, it is possible that s_h must alternate twice. In this case, use induction to find spanning trees in $\mathcal{B}(A^+, B^+)$ and $\mathcal{B}(A^-, B^-)$. These spanning trees have no edges along h , so we may add an edge along h to connect them, as shown in Fig. 5. □

Theorem 13 *Let G be a visibility graph with minimum degree δ and at most four collinear vertices. Then G has connectivity at least $\frac{2\delta+1}{3}$.*

Proof Let $\{A, B, C\}$ be a partition of $V(G)$ such that C separates A and B and $|A| \leq |B|$. By considering a point in A we can see that $\delta \leq |A| + |C| - 1$. If necessary remove points from B so that $|A| = |B|$. Applying Theorem 12 and Observation 1 yields $|C| \geq 2|A| - 1$. Combining these inequalities yields $|C| \geq \frac{2\delta+1}{3}$. □

It turns out that Theorem 13 is best possible. There are visibility graphs with at most three collinear vertices and connectivity $\frac{2\delta+1}{3}$. The construction was discovered by Roger Alperin, Joe Buhler, Adam Chalcraft and Joel Rosenberg in response to a problem posed by Noam Elkies. Elkies communicated their solution to Todd Trimble who published it on his blog [15]. Here we provide a brief description of the construction, but skip over most background details. Note that the original problem and construction were not described in terms of visibility graphs, so we have translated them into our terminology.

The construction uses real points on an elliptic curve. For our purposes a *real elliptic curve* \mathcal{C} is a curve in the real projective plane (which we model as the Euclidean plane with an extra ‘line at infinity’) defined by an equation of the form $y^2 = x^3 + \alpha x + \beta$ (see Fig. 6(a)). The constants α and β are chosen so that the discriminant $\Delta = -16(4\alpha^3 + 27\beta^2)$ is non-zero, which ensures that the curve is non-singular. We define a group operation ‘+’ on the points of \mathcal{C} by declaring that $a + b + c = 0$ if the line through a and b also intersects \mathcal{C} at c , that is, if a, b and c are collinear. The identity element 0 corresponds to the point at infinity in the $\pm y$ -direction, so that for instance $a + b + 0 = 0$ if the line through a and b is parallel to the y -axis. Furthermore, $a + a + b = 0$ if the tangent line at a also intersects \mathcal{C} at b . It can be shown that this operation defines an Abelian group structure on the points of \mathcal{C} .




Fig. 6 (a) The elliptic curve $y^2 = x^3 - x$. (b) The black points separate the white points from the grey points

We will use two facts about real elliptic curves and the group structure on them. First, no line intersects an elliptic curve in more than three points. Second, the group acts continuously: adding a point e which is close to 0 to another point a results in a point close to a (in terms of distance along \mathcal{C}).

Proposition 14 (Alperin, Buhler, Chalcraft and Rosenberg) *For infinitely many integers δ , there is a visibility graph with at most three vertices collinear, minimum degree δ , and connectivity $\frac{2\delta+1}{3}$.*

Proof Begin by choosing three non-zero collinear points a, b and c on a real elliptic curve \mathcal{C} , such that c lies between a and b (see Fig. 6(b)). Choose a point e very close to 0. Define $A := \{a + ie : 0 \leq i \leq m - 1\}$, $B := \{b + je : 0 \leq j \leq m - 1\}$ and $C := \{-(a + b + ke) : 0 \leq k \leq 2m - 2\}$. Let G be the visibility graph of $A \cup B \cup C$. Since the points are all on \mathcal{C} , G has at most three vertices collinear. Observe that the points $a + ie$ and $b + je$ are collinear with the point $-(a + b + (i + j)e)$. Since e was chosen to be very close to 0, by continuity the set A is contained in a small neighbourhood of a , and similarly for B and C . Therefore, the point from C is the middle point in each collinear triple, and so C is a vertex cut in G , separating A and B .

By choosing a, b and c away from any points of inflection, we can guarantee that there are no further collinear triples among the sets A, B or C . Thus a point in A sees all other points in $A \cup C$, a point in B sees all other points in $B \cup C$, and a point in C sees all other points. Therefore the minimum degree of G is $\delta = 3m - 2$, attained by the vertices in $A \cup B$. Hence (also using Theorem 13) the connectivity of G is $|C| = 2m - 1 = \frac{2\delta+1}{3}$. □

We close with the following conjecture.

Conjecture 15 *Every visibility graph with minimum degree δ has connectivity at least $\frac{2\delta+1}{3}$.*

Acknowledgements David Wood is supported by a QEII Research Fellowship from the Australian Research Council. Michael Payne is supported by an Australian Postgraduate Award from the Australian Government. Work by Pavel Valtr is supported by the Ministry of Education of the Czech Republic under project No. 1M0545.

This work began during the Special Semester on Discrete and Computational Geometry in 2010 at the École Polytechnique Fédérale de Lausanne. We are grateful to the Centre Interfacultaire Bernoulli for hosting and the Swiss National Science Foundation for supporting this event.

References

1. Abel, Z., Ballinger, B., Bose, P., Collette, S., Dujmović, V., Hurtado, F., Kominers, S.D., Langerman, S., Pór, A., Wood, D.R.: Every large point set contains many collinear points or an empty pentagon. *Graphs Comb.* **27**, 47–60 (2011)
2. Dankelmann, P., Volkmann, L.: Degree sequence conditions for maximally edge-connected graphs depending on the clique number. *Discrete Math.* **211**, 217–223 (2000)
3. Dumitrescu, A., Pach, J., Tóth, G.: A note on blocking visibility between points. *Geombinatorics* **19**, 67–73 (2009)
4. Kaneko, A., Kano, M.: Discrete geometry on red and blue points in the plane—a survey. In: *Discrete and Computational Geometry. Algorithms Combin.*, vol. 25, pp. 551–570. Springer, Berlin (2003)
5. Kára, J., Pór, A., Wood, D.R.: On the chromatic number of the visibility graph of a set of points in the plane. *Discrete Comput. Geom.* **34**, 497–506 (2005)
6. König, D.: Über Graphen und ihre Anwendung auf Determinantentheorie und Mengenlehre. *Math. Ann.* **77**, 453–465 (1916)
7. Matoušek, J.: Using the Borsuk–Ulam theorem. In: *Universitext*. Springer, Berlin (2003)
8. Matoušek, J.: Blocking visibility for points in general position. *Discrete Comput. Geom.* **42**, 219–223 (2009)
9. Payne, M., Pór, A., Valtr, P., Wood, D.R.: On the connectivity of visibility graphs (2011). <http://arxiv.org/abs/1106.3622>
10. Pfender, F.: Visibility graphs of point sets in the plane. *Discrete Comput. Geom.* **39**, 455–459 (2008)
11. Plesník, J.: Critical graphs of given diameter. *Acta Fac. Rerum Nat. Univ. Comen. Math.* **30**, 71–93 (1975)
12. Plesník, J., Známa, Š.: On equality of edge-connectivity and minimum degree of a graph. *Arch. Math. (Brno)* **25**, 19–25 (1989)
13. Pór, A., Wood, D.R.: On visibility and blockers. *J. Comput. Geom.* **1**, 29–40 (2010)
14. Székely, L.A.: Crossing numbers and hard Erdős problems in discrete geometry. *Comb. Probab. Comput.* **6**, 353–358 (1997)
15. Trimble, T.: Solution to problem of the week No. 13—highly coincidental. Accessed on 30/3/2011. <http://topologicalmusings.wordpress.com/2009/07/22/solution-to-pow-13-highly-coincidental/>
16. Volkmann, L.: Degree sequence conditions for equal edge-connectivity and minimum degree, depending on the clique number. *J. Graph Theory* **42**, 234–245 (2003)