On the bar visibility number of complete bipartite graphs

Weiting Cao* Douglas B. West† Yan Yang‡

May 6, 2019

Abstract

A *t-bar visibility representation* of a graph assigns each vertex up to t horizontal bars in the plane so that two vertices are adjacent if and only if some bar for one vertex can see some bar for the other via an unobstructed vertical channel of positive width. The least t such that G has a t-bar visibility representation is the *bar visibility number* of G, denoted by $b(G)$. For the complete bipartite graph $K_{m,n}$, the lower bound $b(K_{m,n}) \geq \lceil \frac{mn+4}{2m+2n} \rceil$ from Euler's Formula is well known. We prove that equality holds.

Keywords: bar visibility number; bar visibility graph; planar graph; thickness; complete bipartite graph.
MSC Codes: 05C62, 05C10

1 Introduction

In computational geometry, graphs are used to model visibility relations in the plane. For example, we may say that two vertices of a polygon "see" each other if the segment joining them lies inside the polygon. In the *visibility graph* on the vertex set, vertices are adjacent if they see each other. More complicated notions of visibility have been defined for families of rectangles and other geometric objects. Dozens of papers have been written concerning construction and recognition of visibility graphs and applications to search problems and motion planning. For a textbook on algorithms for visibility problems, see Ghosh [8].

*University of Illinois, Urbana, USA: caoweiting@gmail.com.
†Zhejiang Normal University, Jinhua, China, and University of Illinois, Urbana, USA: dwest@math.uiuc.edu. Supported by NNSF of China under Grant NSFC-11871439.
‡Tianjin University, Tianjin, China: yanyang@tju.edu.cn. Supported by NNSF of China under Grant NSFC-11401430.

We consider visibility among horizontal segments in the plane. A graph G is a *bar visibility graph* if each vertex can be assigned a horizontal line segment in the plane (called a *bar*) so that vertices are adjacent if and only if the corresponding bars can see each other along an unobstructed vertical channel with positive width. The assignment of bars is a *bar visibility representation* of G. The condition on positive width allows bars $[(a, y), (x, y)]$ and $[(x, z), (c, z)]$ to block visibility at x without seeing each other.

Tomassia and Tollis [13] and Wismath [16] found a simple characterization of bar visibility graphs. Hutchinson [11] later gave another simple proof for the 2-connected case.

Theorem 1.1 ([13, 16]). *A graph G has a bar visibility representation if and only if for some planar embedding of G all cut-vertices appear on the boundary of one face.*

Theorem 1.1 is quite restrictive. Nevertheless, assigning multiple bars to vertices permits representations of all graphs and leads to a complexity parameter measuring how many bars are needed per vertex, introduced by Chang, Hutchinson, Jacobson, Lehel, and West [5].

Definition 1.2 ([5]). A *t-bar visibility representation* of a graph assigns to each vertex at most t horizontal bars in the plane so that vertices are adjacent if and only if some bar assigned to one sees some bar assigned to the other via an unobstructed vertical channel of positive width. The *bar visibility number* of a graph G, denoted by $b(G)$, is the least integer t such that G has a t-bar visibility representation.

Results in [5] include the determination of visibility number for planar graphs (always at most 2), plus $b(K_n) = \lceil n/6 \rceil$ for $n \geq 7$, the determination of $b(K_{m,n})$ within 1, and $b(G) \leq \lceil n/6 \rceil + 2$ for every n-vertex graph G. Results on the visibility numbers for hypercubes [15] and an analogue for directed graphs [1] have also been obtained. For complete bipartite graphs, the result was as follows.

Theorem 1.3 ([5]). *$r \leq b(K_{m,n}) \leq r + 1$, where $r = \left\lceil \frac{mn+4}{2m+2n} \right\rceil$.*

To prove the lower bound, consider a t-bar representation, add edges to encode visibilities that produce edges of $K_{m,n}$, and then shrink bars to single points. This produces a bipartite plane graph H with at most $t(m+n)$ vertices and at least mn edges. Hence $mn \leq 2t(m+n)-4$ by Euler's Formula, so $b(K_{m,n}) \geq r$. Equality requires most faces in H to have length 4.

In this paper, we prove $b(K_{m,n}) = r$. Section 2 contains a short proof valid for $K_{n,n}$. For this case, it suffices to decompose the graph into r bar visibility graphs, where a *decomposition* of G is a set of edge-disjoint subgraphs whose union is G. The subgraphs can then be repesented with disjoint projections on the horizontal axis. In Section 3, we present a different approach that solves the problem for all complete bipartite graphs.

Our results are related to earlier work. A *t-split* of a graph G is a graph H in which each vertex is replaced by a set of at most t independent vertices in such a way that u and v are adjacent in G if and only if some vertex in the set representing u is adjacent in H to some vertex in the set representing v. The graph G used to prove the lower bound for Lemma 1.3 is an example of a t-split of $K_{m,n}$. As defined by Eppstein et al. [7], the *planar split thickness* (or simply *split thickness*) of a graph G, which we denote by $\sigma(G)$, is the minimum t such that G has a t-split that is a planar graph. As explained above, always $\sigma(G) \le b(G)$. If G has a $\sigma(G)$-split that is 2-connected, then $\sigma(G) = b(G)$.

This connection was noted earlier in the thesis of the first author [4], where planar split thickness was given the unfortunate name "split number", creating confusion with another concept. The *splitting number* of a graph is the minimum number of successive splits of one vertex into two (with each incident edge being inherited by one of the two new vertices) needed to produce a planar graph.

The notion of t-split originated with Heawood [9], who proved that K_{12} has a 2-split. Later, Ringel and Jackson [12] proved in effect that K_n has a $\lceil n/6 \rceil$-split. A short proof of this by Wessel [14] was used in [5] to prove $b(K_n) = \lceil n/6 \rceil$.

The results in [7] that concern complete bipartite graphs determine those that are 2-splittable. They are the same as those having bar visibility number at most 2. Their lower bounds on $\sigma(K_{m,n})$ use the same counting argument from Euler's Formula that yields the lower bounds for $b(K_{m,n})$ (see [5]).

In [7], the authors close the paper by asking whether graphs embeddable on the surface of genus k are $(k+1)$-splittable, as an open question. This follows from a recent result about the *thickness* $\theta(G)$ of a graph G, defined to be the minimum number of planar graphs needed to decompose G. A decomposition into k planar graphs is a k-split, so $\sigma(G) \le \theta(G)$; this motivates the term "split thickness". Xu and Zha [17] proved that $\theta(G) \le k + 1$ when G embeds on the surface of genus k, thereby providing a positive answer to the question in [7].

2 The bar visibility number of $K_{n,n}$

As noted above, thickness provides an upper bound on the split thickness, and the split thickness usually equals the bar visibility number. Beineke, Harary, and Moon [3] determined $\theta(K_{m,n})$ for most m and n.

Lemma 2.1 ([3])**.** $\theta(K_{n,n}) = \left\lceil \frac{n+2}{4} \right\rceil$.

When $\theta(K_{n,n})$ is the desired value for $b(K_{n,n})$, we aim to decompose $K_{n,n}$ into that number of bar visibility graphs. The difficult case is when $b(K_{n,n}) < \theta(K_{n,n})$.

Theorem 2.2. $b(K_{n,n}) = \left\lceil \frac{n+1}{4} \right\rceil$, *except for* $b(K_{3,3}) = 2$.

Proof. It is immediate that $K_{1,1}$ and $K_{2,2}$ are bar visibility graphs. Since $K_{3,3}$ is not planar, $b(K_{3,3}) \geq 2$; equality holds because $K_{3,3}$ decomposes into a 6-cycle and a matching of size 3, both of which are bar visibility graphs. Hence we may assume $n \geq 4$.

Let $r = \lceil (n+1)/4 \rceil$. When $\theta(K_{n,n}) = r$, we will decompose $K_{n,n}$ into r bar visibility graphs. This will leave the case where $n \equiv 3 \mod 4$, in which case $r < \theta(K_{n,n})$ and $K_{n,n}$ cannot decompose into r bar visibility graphs. Let U and V be the parts of $K_{n,n}$, with $U = \{u_1, \ldots, u_n\}$ and $V = \{v_1, \ldots, v_n\}$. Let $p = \lfloor n/4 \rfloor$.

For $n \equiv 0 \mod 4$, Chen and Yin [6] provided a decomposition of $K_{n,n}$ into $p+1$ planar subgraphs $\{G_1, \ldots, G_{p+1}\}$. Let $[p] = \{1, \ldots, p\}$. For $1 \leq j \leq p$, let $U_1^j = \bigcup_{i \in [p]-\{j\}} \{u_{4i-3}, u_{4i-2}\}$, let $U_2^j = \bigcup_{i \in [p]-\{j\}} \{u_{4i-1}, u_{4i}\}$, let $V_1^j = \bigcup_{i \in [p]-\{j\}} \{u_{4i-3}, u_{4i-1}\}$, and let $V_1^j = \bigcup_{i \in [p]-\{j\}} \{u_{4i-2}, u_{4i}\}$. Figure 1 shows the subgraph G_j, for $1 \leq j \leq p$. Being a 2-connected planar graph, it is a bar visibility graph. The subgraph induced by the eight special vertices $u_{4j-3}, \ldots, u_{4j}$ and $v_{4j-3}, \ldots, v_{4j}$ is $K_{4,4}$ minus the edges of the form $u_i v_i$. The remaining graph G_{p+1} is the matching consisting of $u_i v_i$ for $1 \leq i \leq 4p$. Again this is a bar visibility graph.

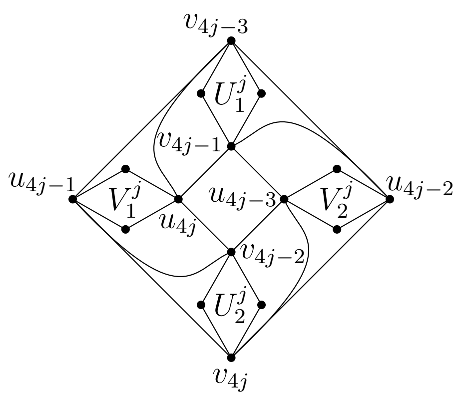

Figure 1: The graph G_j in a planar decomposition of $K_{4p,4p}$.

For $n = 4p+1$, we add two vertices u_{4p+1} and v_{4p+1}, with u_{4p+1} adjacent to V and v_{4p+1} adjacent to U. The edges incident to u_{4p+1} and v_{4p+1} can be added to the graph G_{p+1} of the previous case, as shown in Figure 2. Again this graph is planar and 2-connected, so again we have a decomposition $\tilde{G}_1, \ldots, \tilde{G}_{p+1}$ into $p+1$ bar visibility graphs.

For $n = 4p+2$, we modify the decomposition given for $K_{4p,4p}$ to accommodate the edges incident to $\{u_{4p+1}, u_{4p+2}, v_{4p+1}, v_{4p+2}\}$. First form $\widehat{G}_{p+1}$ by adding to the matching G_{p+1} the edges joining u_{4p+1} to $\bigcup_{i \in [p]} \{v_{4i-2}, v_{4i}\}$, joining u_{4p+2} to $\bigcup_{i \in [p]} \{v_{4i-3}, v_{4i-1}\}$, joining v_{4p+1} to $\bigcup_{i \in [p]} \{v_{4i-2}, v_{4i-3}\}$, and joining v_{4p+2} to $\bigcup_{i \in [p]} \{v_{4i}, v_{4i-1}\}$, plus the edges of the 4-cycle $[u_{4p+1}, v_{4p+1}, u_{4p+2}, v_{4p+2}]$, as shown in Figure 3. To include the remaining edges involving the four added vertices, for $1 \leq j \leq p$ obtain $\widehat{G}_j$ from G_j by adding u_{4p+i} to U_i^j and v_{4p+i} to V_i^j, for $i \in \{1, 2\}$. Each of these four vertices gains the two neighbors in G_j that are

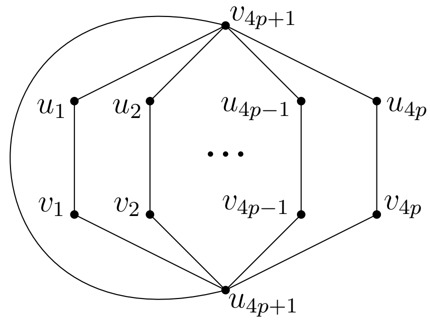

Figure 2: The subgraph $\widetilde{G}_{p+1}$ in the planar decomposition of $K_{4p+1,4p+1}$

shared by the vertices of the set to which it was added. Over the resulting $\widehat{G}_1, \ldots, \widehat{G}_p$, it gains precisely the neighbors in the other part that it does not have in $\widehat{G}_{p+1}$. We again have r 2-connected planar graphs decomposing $K_{n,n}$.

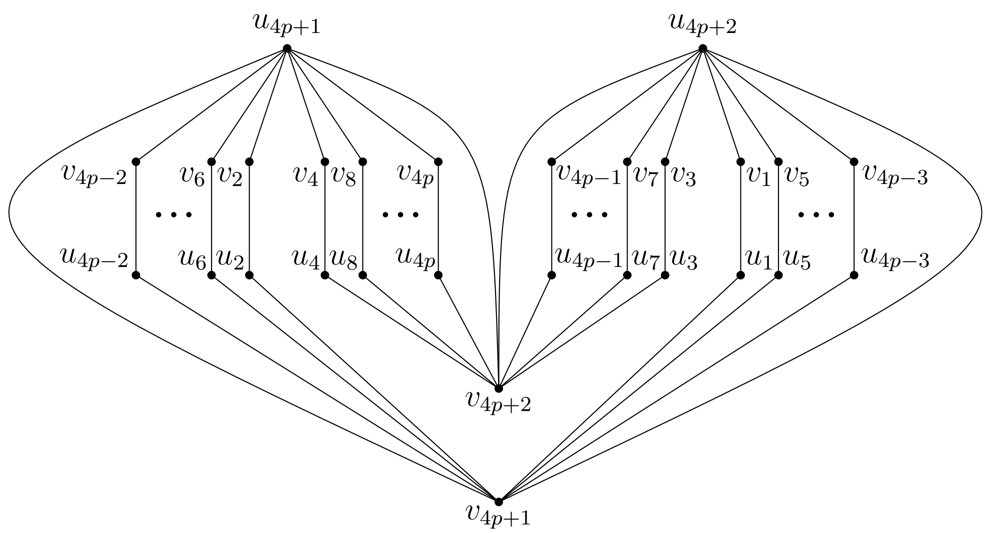

Figure 3: The subgraph $\widehat{G}_{p+1}$ in the planar decomposition of $K_{4p+2,4p+2}$

The remaining case is $n = 4p + 3$. A graph G is *thickness t-minimal* if $\theta(G) = t$ and every proper subgraph of G has thickness less than t. When $n = 4p+3$, the graph $K_{4p+3,4p+3}$ is a thickness $(p + 2)$-minimal graph. Hobbs and Grossman [10] and Bouwer and Broere [2] independently gave two different decompositions of $K_{4p+3,4p+3}$ into planar subgraphs $H_1, \ldots, H_{p+2}$. In each case, each H_i for $1 \le i \le p + 1$ is a 2-connected maximal planar bipartite graph (hence a bar visibility graph), and the graph H_{p+2} contains only one edge. Let this edge be u_iv_j (it is u_1v_1 in [10] and $u_{4p+3}v_{4p-1}$ in [2]).

The bar visibility representation algorithm of [13] uses "s, t-numberings", allowing one to choose any vertex of a bar visibility graph to be the unique lowest or highest bar in the representation. Since we have reduced to the case $n \geq 4$, we have $p+1 \geq 2$. Choose a representation of H_1 in which u_i is the lowest bar and a representation of H_2 in which v_j is the highest bar. Place the representation of H_1 above the representation of H_2 to incorporate the edge $u_i v_j$ without using an extra bar for u_i or v_j.

We must also show that the bars for u_i in H_1 and v_j in H_2 can prevent unwanted visibilities between bars for vertices above and below them. Since the graph is bipartite, we may assume that bars for the two parts occur on horizontal lines with those for U having odd vertical coordinates and those for V having even coordinates. In addition, the bars on one horizontal line can extend to meet at endpoints to block visibility between higher and lower bars for the other part (using both the requirement of positive width for visibility and the fact that we are representing the complete bipartite graph). The bars can extend so that on each horizontal line the leftmost occupied point is the same and the rightmost occupied point is the same. Now the two representations can combine as described above. □

3 General approach to $b(K_{m,n})$

Henceforth let $f(m,n) = \left\lceil \frac{mn+4}{2m+2n} \right\rceil$. Our proof of $b(K_{m,n}) \leq f(m,n)$ for $m,n \in \mathbb{N}$ is independent of the shorter proof for $m=n$ given in the previous section, which relied on thickness results from earlier papers. This proof is self-contained.

As mentioned in the introduction, it suffices to produce a 2-connected r-split of $K_{m,n}$, where $r = f(m,n)$; this is our aim. We will consider various cases depending on parity. In this section we present the common aspects of the constructions. We may assume $m \geq n$. Let the two parts of $K_{m,n}$ be X and Y with $X = \{x_1, \ldots, x_m\}$ and $Y = \{y_1, \ldots, y_n\}$.

When n is even and $m > \frac{1}{2}(n^2 - 2n - 4)$, or when n is odd and $m > n^2 - n - 4$, we compute $r = \left\lceil \frac{n}{2} \right\rceil$. In this case let G_i be the subgraph induced by $X \cup \{y_{2i-1}, y_{2i}\}$, except that $G_{(n+1)/2}$ is the subgraph induced by $X \cup \{y_n\}$ when n is odd. Since $K_{m,2}$ and $K_{m,1}$ are bar visibility graphs, this decomposes $K_{m,n}$ into r bar visibility graphs. Note that $3 > 3^2 - 3 - 4$, so when $n = 3$ we have already considered all cases, and henceforth we may assume $n \geq 4$.

We have also considered all cases with $r = \left\lceil \frac{n}{2} \right\rceil$, so henceforth we may assume $r \leq \left\lfloor \frac{n-1}{2} \right\rfloor$. Let $s = \left\lceil \frac{n}{2} \right\rceil - r$. For fixed n, the value of $\frac{mn+4}{2m+2n}$ increases with m. Since $m \geq n$, we have $r \geq \left\lceil \frac{n^2+1}{4n} \right\rceil \geq \left\lceil \frac{n+1}{4} \right\rceil$. Thus $s \leq r$. The case $s = r$ requires $s = r = \frac{n+1}{4}$ and hence $n \equiv 3 \mod 4$. For $m \in \{n, n+1, n+2\}$, the values of $\frac{mn+4}{2m+2n}$ are $\frac{n+1}{4}$, $\frac{n+0.5+15/(4n+2)}{4}$, and $\frac{n+1+3/(n+1)}{4}$, respectively. The last exceeds $\frac{n+1}{4}$. Thus the case $s = r$ occurs if and only if $n \equiv 3 \mod 4$ and $m \in \{n, n+1\}$. Otherwise, $s < n/4 < r$.

We will construct a 2-connected planar graph G that is an r-split of $K_{m,n}$. In G, each vertex will have a label in $X \cup Y$, with each label used at most r times. When no vertices labeled x_i and y_j are yet adjacent, we say that x_i *misses* y_j; otherwise x_i *hits* y_j. We place vertices in the coordinate plane, with vertices labeled by X on the horizontal axis and vertices labeled by Y on the vertical axis. Edges will join only the two axes, so no unwanted edges are formed. To facilitate understanding, we first exhibit in Figure 4 the graph G that we produce when $(m, n) = (8, 7)$. For clarity, we record only the subscripts of the labels on the vertices; the labels are from X on the horizontal axis and from Y on the vertical axis.

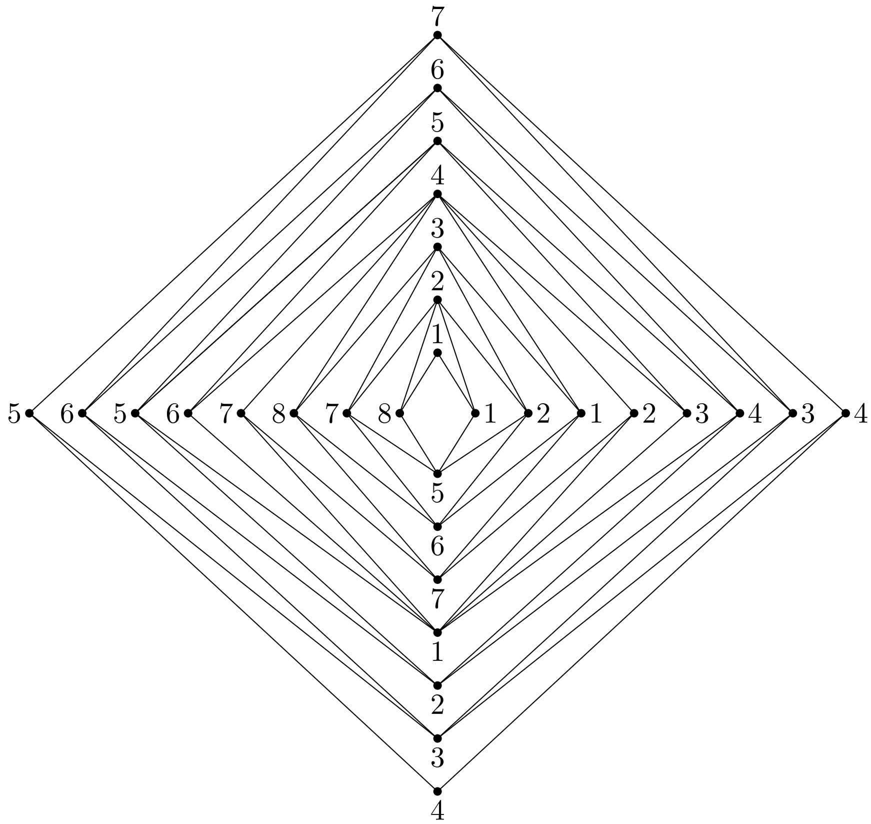

Figure 4: A bar visibility graph that is a 2-split of $K_{8,7}$

The plan: We first construct subgraphs separately in each half-plane bounded by the vertical axis. Combining these two subgraphs along the vertical axis will yield a 2-connected plane graph $\widehat{G}$ with $rn + sm$ vertices such that labels in X occur s times and labels in Y occur r times. Ideally, each $x_i \in X$ will hit $n - 2(r - s)$ different vertices of Y, and the vertices of Y that x_i misses will form $r - s$ pairs such that each pair lies on a face of length 4. We will then insert a copy of x_i in each such face, adjacent to the missed vertices of Y, so that x_i now hits all n vertices of Y. This brings the usage of each label to r vertices, and the result will be a 2-connected r-split of $K_{m,n}$. Because the parity of $n - 2(r - s)$ depends on the parity of n, we will need to use different building blocks for even n and odd n.

We begin by addressing the matter of 2-connectedness.

Lemma 3.1. *Let $\widehat{G}$ be a plane graph whose vertex set is comprised of sets A^+, A^-, B^+, and B^- placed along the positive and negative horizontal axes and the positive and negative vertices axes, respectively (as in Figure 4). If the four subgraphs induced by $B^+ \cup A^+$, $A^+ \cup B^-$, $B^- \cup A^-$, and $A^- \cup B^+$ are connected (and there are no other edges), then $\widehat{G}$ is 2-connected.*

Proof. We show that any two vertices u and v in $\widehat{G}$ are connected by two internally disjoint paths. Consider the four subgraphs combined as in Figure 4; the graph in each quadrant is connected. Suppose first that u and v are not on the same half-axis. Choose edges uu' and vv' such that u' and v' are on half-axes different from each other and from u and v. Now the four vertices are on distinct half-axes, and in every case we have chosen our two edges from the subgraphs in opposite quadrants. Because the remaining two quadrants are connected, we can choose a path in each to connect the vertices we have chosen on its half-axes. Now our four chosen vertices lie on a cycle, which contains the two desired u, v-paths.

If u and v lie on the same half-axis, then we can choose uu' and vv' so that u' and v' are on the same neighboring half-axis. Now the four vertices are in the boundary of the same quadrant. We find a u, v-path in one neighboring quadrant and a u', v'-path in the other neighboring quadrant. Again the four vertices lie on a cycle. □

To facilitate computations, we want to reduce to the critical values of m. The next lemma shows that examples for smaller m will cause no difficulty.

Lemma 3.2. *Let H be a 2-connected plane bipartite graph. If each part in H has at least three vertices, and $v \in V(H)$, then edges can be added to $H - v$ to obtain a 2-connected plane bipartite graph with the same vertex bipartition as $H - v$.*

Proof. Let X and Y be the parts of the bipartition of H; we may assume $v \in X$. If $H - v$ is 2-connected, then nothing need be done. Otherwise, $H - v$ has a cut-vertex w. Since $\{v, w\}$ is a separating set of H, in the embedding of H these two vertices must lie on the same face. Since $|X| \geq 3$, some component C of $H - \{v, w\}$ contains a vertex of X on its outside face in the embedding; let x be such a vertex. Each other component C' contains a neighbor of v (in Y), which must lie on the outside face of C' in the embedding. Make this vertex in each such component C' adjacent to x. The resulting graph is planar, 2-connected, and has the same bipartition as $H - v$. □

Lemma 3.3. *For $n, r \in \mathbb{N}$ with $r < n/2$, the largest m such that $f(m, n) = r$ is $\left\lfloor \frac{2rn-4}{n-2r} \right\rfloor$.*

Proof. If $\left\lceil \frac{mn+4}{2m+2n} \right\rceil = r$, then $r - 1 < \frac{mn+4}{2m+2n} \leq r$, equivalent to $\frac{2n(r-1)+4}{n-2(r-1)} < m \leq \frac{2nr-4}{n-2r}$. □

Lemma 3.4. *If $K_{m,n}$ has a 2-connected planar r-split whenever $r < n/2$ and $m = \lfloor \frac{2rn-4}{n-2r} \rfloor$, then $b(K_{m,n}) = \lceil \frac{mn+4}{2m+2n} \rceil$ for all $m, n \in \mathbb{N}$.*

Proof. The given r-split G yields the claim for $m = \lceil \frac{mn+4}{2m+2n} \rceil$, where the parts of $K_{m,n}$ are X and Y with $|X| = m$. By Lemma 3.3, we do not need r-split for this n with larger m. For smaller m, Lemma 3.2 allows us iteratively to delete the copies in G of one vertex of X, restoring 2-connectedness after each vertex deletion. We have an r-split of the resulting complete bipartite graph. Hence $b(K_{m,n}) \leq r$ for all m where $b(K_{m,n}) \leq r$ is desired. □

Although the details of the construction differ for even and odd n, the main idea is the same, so we can introduce some common notation.

Definition 3.5. For ease of illustration, we squeeze each half-plane into a strip, drawing its three axis rays along horizontal lines (see Figure 5). The vertices receiving labels in Y are the first $\lceil rn/2 \rceil$ integer points on the positive vertical axis and the first $\lfloor rn/2 \rfloor$ integer points on the negative vertical axis, called B^+ and B^-, respectively. Starting from the origin, label B^+ in order using $y_1, \ldots, y_n$, through increasing indices cyclically modulo n. Similarly label B^-, but start with $y_{\lceil n/2 \rceil + 1}$ and again continue increasing through indices modulo n (see Figure 5). The last labels on B^+ and B^- are $\{y_n, y_{\lceil n/2 \rceil}\}$, with y_n ending B^+ if r is even and B^- if r is odd. Each label y_i is used exactly r times. The vertices with labels in X are placed at integer points on the horizontal axis, with A^+ and A^- respectively denoting the sets of positive and negative points used. Let $A = A^+ \cup A^-$ and $B = B^+ \cup B^-$.

4 The case of even n

As seen in Figure 5, most vertices in A will have two consecutive neighbors in B^+ and in B^-; vertices in the middle row for the horizontal axis receive two neighbors above and below. For now ignore the edges added there for x_{13} and x_{14}. The main part of the construction consists of special building blocks that enable most vertices in X to hit $n - 2(r - s)$ labels in Y using s vertices on the horizontal axis. In Figure 5 these use four vertices in A and five vertices in B^+ above and B^- below. Throughout this section, n is even.

Definition 4.1. An *opposite pair* is a pair of labels in Y whose subscripts differ by $n/2$; that is, having the form $y_i, y_{i+n/2}$, where the computation in subscripts is viewed modulo n. The labels of vertices in B^+ and B^- at the same distance from the origin form an opposite pair.

An *i-brick* is a graph induced by $2s$ consecutive vertices in A^+ or A^- (with alternating labels x_{2i-1} and x_{2i}) and $2s + 1$ consecutive vertices in each of B^+ and B^- (see Figure 5). The vertices used from B form opposite pairs. The edges of the brick join the jth vertex among its vertices from A to the jth and $(j + 1)$th opposite pairs among its vertices from B.

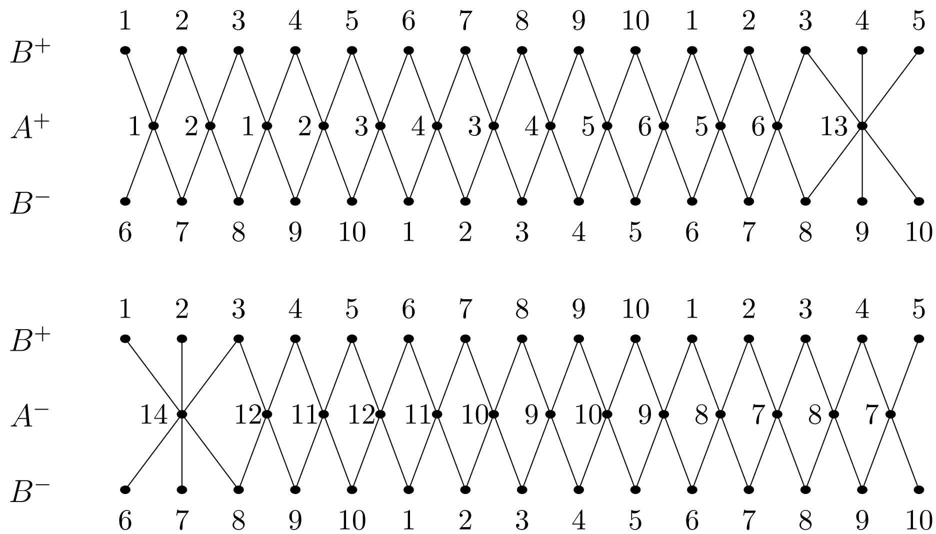

Figure 5: Pattern for even n, shown for $K_{14,10}$ with $(r, s) = (3, 2)$.

Lemma 4.2. *When n is even, the labels from Y that lie on a 4-face in an i-brick form an opposite pair. Each label for a vertex of X in a brick hits two intervals of $2s$ cyclically consecutive labels in Y, forming $2s$ distinct cyclically consecutive opposite pairs. The labels missed by such a vertex of X thus also come in opposite pairs.*

Proof. The claims follow immediately from the construction in Definition 4.1, because vertices in corresponding positions in the lists B^+ and B^- form opposite pairs. These pairs are distinct because $2s < n/2$. □

Theorem 4.3. *If n is even, then $b(K_{m,n}) = \lceil \frac{mn+4}{2m+2n} \rceil = r$.*

Proof. By Lemma 3.4, we may assume $m = \lfloor \frac{2rn-4}{n-2r} \rfloor$. We have also reduced to $r < n/2$. With $s = n/2 - r$, we have $s < n/4 < r$. Let $q = \frac{rn/2-1}{2s}$ and $t = \lfloor q \rfloor$. We have $m = \lfloor 4q \rfloor$, so $m = 4t + j$ for some $j \in \{0, 1, 2, 3\}$, where j depends on which fourth of $[0, 1)$ contains q.

We first put i-bricks into A^+, for $1 \le i \le t$ (from the left in Figure 5). The last vertex of B^+ in the i-brick is also the first vertex in the $(i + 1)$-brick (similarly for B^-). Thus these bricks use $1 + 2st$ vertices from B^+ (and B^-). Since $|B^+| = rn/2 = 1 + 2sq$, there is room for these bricks. Similarly, working inward from the outer face (from the right in Figure 5), we put i-bricks into A^- for $t + 1 \le i \le 2t$. Counting the last vertex of the t-brick, the number of vertices remaining visible to unused vertices of A^+ at the right end of B^+ and B^- is $(rn/2) - 2st$, which equals $1 + 2s(q - t)$. Similarly, this many vertices are visible to unused vertices of A^- at the left end.

Since all opposite pairs remain available on the faces with A^+ involving the first $n/2$ vertices in B^+ and B^-, we have now satisfied the vertices $x_1, \ldots, x_{4t}$. Each such label has been used s times and hit $4s$ labels in Y. Since $n - 4s = 2(r - s)$ and the $2(r - s)$ missed labels occur in opposite pairs, we can add $r - s$ vertices with this label in the appropriate faces to hit the remaining $2(r - s)$ labels in Y.

Since $m = 4t + j$, there remain j vertices to process in X, where $j \leq 3$ (none if $j = 0$; the example in Figure 5 has $j = 2$). Note that the opposite pair in B^+ and B^- that is seen from the left (inner) end of A^- cyclically follows the opposite pair seen from the right (outer) end of A^+. Thus if a label in X sees consecutive pairs in $B^+ \cup B^-$ using vertices in A^+, or in A^-, or from the end of A^+ and beginning of A^-, then the labels in Y hit by that vertex will be distinct as long as the number of pairs is at most $n/2$.

When j is odd, x_m will receive one vertex at the end of A^+ and one at the beginning of A^-. When $j \geq 2$, we assign one vertex in A^+ to x_{m-j+1} and one vertex in A^- to x_{m-j+2}. If this gives p vertices to a vertex x_i and each of the remaining $r - p$ vertices for x_i will see one opposite pair by putting it into a 4-face, then the specified p vertices for x_i need to hit $n - 2(r - p)$ labels in Y. This value equals $2s + 2p$, so x_i needs to hit $s + p$ consecutive opposite pairs.

For $p \leq 2$, ensuring that the labels hit are distinct requires $s+2 \leq n/2$. Since $s \leq (n-2)/4$ when n is even, it suffices to have $(n - 2)/4 \leq n/2 - 2$, which is equivalent to $n \geq 6$ (and when $n = 4$ we cannot have $s < n/4$).

It remains only to show that $B^+ \cup B^-$ has enough such pairs available. For $j \in \{1, 2, 3\}$, we need in total to hit $s + 2$, $2s + 2$, or $3s + 4$ consecutive pairs, respectively. We have observed that there are $1 + 2s(q - t)$ opposite pairs visible both from the right end of A^+ and the left end of A^-. Since $q - t \geq j/4$, the $2 + 4s(q - t)$ pairs are at least $s + 2$, $2s + 2$, and $3s + 2$ for $j \in \{1, 2, 3\}$, respectively. However, when $j = 3$ we are using two vertices in each of A^+ and A^-, meaning that the last pair seen by one vertex can also be the first pair seen by the other. This means that in total the vertices can see $3s + 4$ pairs instead of $3s + 2$, which is the number needed.

Finally, we must ensure that the graph $\widehat{G}$ produced before adding the excess labels in faces is 2-connected (it is an elementary exercise that adding vertices of degree 2 to a 2-connected graph preserves 2-connectedness). By Lemma 3.1, it suffices to show that the four subgraphs induced by $B^+ \cup A^+$, $A^+ \cup B^-$, $B^- \cup A^-$, and $A^- \cup B^+$ are connected. After adding the vertices with labels $x_{m-j+1}, \ldots, x_m$ in the last step, we may have left some vertices of B^+ or B^- unhit. Add edges joining A and B in each of the four induced subgraphs to make them connected, while remaining planar and bipartite with the same bipartition. Because we are seeking a representation of a complete bipartite graph, extra visibilities between the parts do not cause a problem. □

5 The case of odd n

Our general aim was to have s vertices for x_i hit $n-2(r-s)$ labels in Y. When n is odd, this amount is odd, so we change the definition of bricks. They will still use $2s$ vertices from A, but now they will use one less vertex each in B^+ and B^-. Indeed, the bricks we used before are too big to fit onto B^+ and B^-. Throughout this section, n is odd.

Definition 5.1. A *skew pair* is a pair of labels in Y whose subscripts differ by $\lfloor n/2 \rfloor$; that is, having the form $y_i, y_{i+(n-1)/2}$, where the computation in subscripts is viewed modulo n.

For odd n, an *i-brick* is a graph induced by $2s$ consecutive vertices in A^+ or A^- (alternating labels x_{2i-1} and x_{2i}) and $2s$ consecutive vertices in B^+ and B^- (see Figure 6). Bricks using A^+ start from the left end (near the origin). Those using A^- start from the right (not near the origin). Vertices from B^+ and B^- in a brick have the same distances from the start.

Let $\{\widehat{B}, \tilde{B}\} = \{B^+, B^-\}$. Measured from the start, the edges of a brick join the jth copy of x_{2i-1} to the $(2j-1)$th and $2j$th vertices of the brick in $\widehat{B}$ and the $(2j-2)$th and $(2j-1)$th vertices of the brick in $\tilde{B}$, except that the first copy of x_{2i-1} hits in $\tilde{B}$ only the first vertex. Similarly, the jth copy of x_{2i} hits the $2j$th and $(2j+1)$th vertices of the brick in $\widehat{B}$ and the $(2j-1)$ and $2j$th vertices of the brick in $\tilde{B}$, except that the sth copy of x_{2i} hits in $\widehat{B}$ only the last vertex. For bricks from the left (using A^+), set $\widehat{B} = B^+$ and $\tilde{B} = B^-$. For bricks from the right (using A^-), let $\widehat{B}$ be the member of $\{B^+, B^-\}$ whose last label is y_n, and let $\tilde{B}$ be the member whose last label is $y_{(n+1)/2}$.

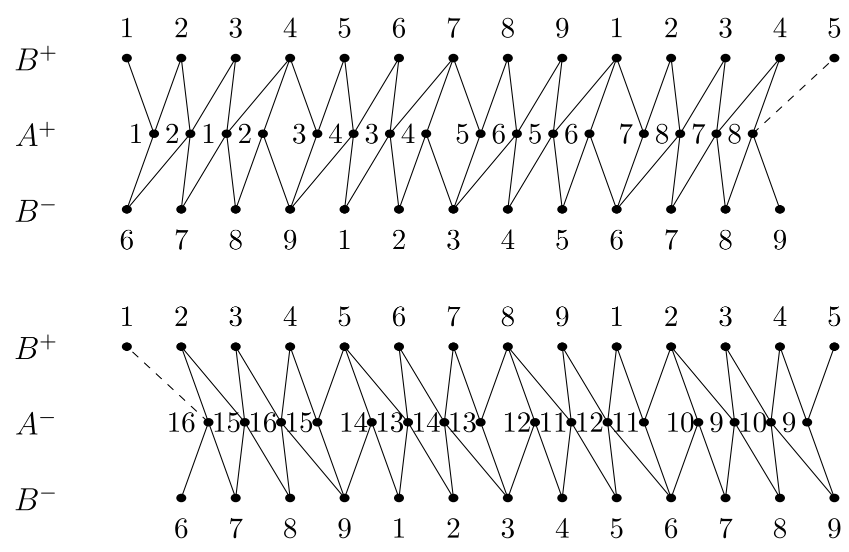

Figure 6: Pattern for odd n, shown for $K_{16,9}$ with $(r,s) = (3,2)$.

Lemma 5.2. *When n is odd, the labels from Y that lie on a 4-face in an i-brick form a skew pair. Each label for a vertex of X in a brick hits two intervals of $2s-2$ cyclically consecutive labels in Y, forming $2s-2$ cyclically consecutive skew pairs, plus one more label at the end of one of those intervals. These labels are distinct. The labels missed by such a vertex of X also come in skew pairs.*

Proof. The labels on a 4-face in a brick are the jth vertex of $\widehat{B}$ and $(j+1)$th of $\tilde{B}$ from the start. Since B^+ starts with y_1 and B^- starts with $y_{(n+3)/2}$, for bricks using A^+ the two labels on the face are y_{j+1} and $y_{j+(n+1)/2}$, which form a skew pair taking subscripts modulo n. For bricks using A^-, starting from the other end, $\widehat{B}$ starts with y_n and $\tilde{B}$ starts with $y_{(n+1)/2}$. (As specified in Definition 3.5, y_n ends B^+ if r is even and B^- if r is odd.) On a 4-face in such a brick we have y_{n-j} and $y_{(n+1)/2-j-1}$. Since $(n-1)/2-j \equiv n-j+(n-1)/2 \mod n$, again we have a skew pair.

Since corresponding positions in B^+ and B^- are labeled by skew pairs, the $4s$ labels in Y occurring in a brick are distinct unless $2s=(n+1)/2$, which can occur when $n \equiv 3 \mod 4$ and $s=r=(n+1)/4$. In this case, the first label from B^+ is the same as the last label from B^- in a brick. However, as constructed in Definition 5.1, the first label from B^+ is hit only by x_{2i-1}, and the last label from B^- is hit only by x_{2i} in the brick, so each label from X still hits $4s-1$ distinct labels in Y.

These $4s-1$ distinct labels group into $2s-1$ cyclically consecutive skew pairs plus one more label. The two intervals of labels hit by the pairs leave two intervals of labels missed, and the lengths of the intervals of missed labels are $(n+1)/2-s+1$ and $(n-1)/2-s+1$. The extra label hit by x_i is at the end of the longer interval. No matter which end of the longer interval it shortens, the remaining missed labels match up as skew pairs. □

The approach to the construction is the same as for even n in Theorem 4.3, but the technical details are different.

Theorem 5.3. *If n is odd, then $b(K_{m,n}) = \left\lceil \frac{mn+4}{2m+2n} \right\rceil = r$.*

Proof. By Lemma 3.4, we may assume $m = \left\lfloor \frac{2rn-4}{n-2r} \right\rfloor$. We have reduced to $\frac{n+1}{4} \le r < \frac{n}{2}$. With $s=\frac{n+1}{2}-r$, we have $s \le \frac{n+1}{4} \le r$. Also $n-2r=2s-1$. Let $q=\frac{rn/2-1}{2s-1}$ and $t=\lfloor q \rfloor$. We have $m=\lfloor 4q \rfloor$, so $m=4t+j$ for some $j \in \{0,1,2,3\}$, depending on where in $[0,1)$ is q.

We put i-bricks into A^+ for $1 \le i \le t$ (from the left in Figure 6). The last vertex of B^+ in the i-brick is the first vertex in the $(i+1)$-brick (similarly for B^-). Thus these bricks use $1+(2s-1)t$ vertices from B^+ (and B^-). Since B^+ and B^- each have at least $(rn-1)/2$ vertices, and $(rn-1)/2=(1/2)+(2s-1)q \ge (1/2)+(2s-1)t$, there is room for these bricks. Similarly, working inward from the outer face (from the right in Figure 6), we put i-bricks into A^- for $t+1 \le i \le 2t$. Counting the last vertex of the t-brick, the number of vertices remaining visible to unused vertices of A^+ at the right end of B^+ and B^- together is $rn - 2(2s-1)t$, which equals $2 + 2(2s-1)(q-t)$. Similarly, this many vertices are visible to unused vertices of A^- at the left end.

Since all skew pairs remain available on the faces with A^+ involving the first n vertices in B^+ and B^-, we have now satisfied the vertices $x_1, \dots, x_{4t}$. Each such label has been used s times and hit $4s-1$ labels in Y. Since $n+1-4s = 2(r-s)$ and the $2(r-s)$ missed labels occur in skew pairs, we can add $r-s$ vertices with this label in the appropriate faces to hit the remaining $2(r-s)$ labels in Y.

Since $m = 4t + j$, there remain j vertices to process in X, where $j \le 3$ (none if $j = 0$; as in the example in Figures 4 and 6). The labels that end B^+ and B^- and may be visible at the end of A^+ are $\{y_n, y_{(n+1)/2}\}$, a skew pair. The labels that begin B^+ and B^- are $\{y_1, y_{(n+3)/2}\}$, the next skew pair. Thus if a label in X sees consecutive skew pairs in $B^+ \cup B^-$ using vertices in A^+, or in A^-, or from the end of A^+ and beginning of A^-, then the labels in Y hit by that vertex will be distinct as long as the number of pairs is at most $n/2$.

When j is odd, x_m will receive one vertex at the end of A^+ and one at the beginning of A^-. When $j \ge 2$, we assign one vertex in A^+ to x_{m-j+1} and one vertex in A^- to x_{m-j+2}. If this assigns p vertices to a vertex x_i and each of the remaining $r-p$ vertices for x_i will see one skew pair by putting it into a 4-face, then the specified p vertices for x_i need to hit $n - 2(r-p)$ labels in Y. This value equals $2s - 1 + 2p$, so it suffices for x_i to hit $s + p - 1$ consecutive skew pairs and one label from the next pair.

For $p \le 2$, ensuring that the labels hit are distinct requires $s + 1 \le (n-1)/2$. Since $s \le (n+1)/4$, it suffices to have $(n+1)/4 \le (n-3)/2$, which is equivalent to $n \ge 7$. Since we have reduced to $n \ge 4$, and $s \le (n-1)/4$ when $n \equiv 1 \mod 4$, all cases are covered.

It remains only to show that $B^+ \cup B^-$ has enough vertices available. For $j \in \{1, 2, 3\}$, we need in total to hit $2s+3$, $4s+2$, or $6s+5$ labels, respectively. We have observed that there are in total $2 + 2(2s-1)(q-t)$ vertices of B visible both from the right end of A^+ and the left end of A^-. Since $q - t \ge j/4$, the total number of vertices is at least $4 + (2s-1)j$, which is enough when $j \le 2$. When $j = 3$ we are using two vertices in each of A^+ and A^-, meaning that the last pair seen by one vertex can also be the first pair seen by the other. This provides four additional visibilities to reach the needed $6s + 5$.

Finally, we must ensure that the graph $\widehat{G}$ produced before adding the excess labels in faces is 2-connected. Here the argument applying Lemma 3.1 to the subgraphs induced by $B^+ \cup A^+$, $A^+ \cup B^-$, $B^- \cup A^-$, and $A^- \cup B^+$ is the same as in Theorem 4.3. $\square$

References

[1] M. Axenovich, A. Beveridge, J.P. Hutchinson and D.B. West, Visibility number of directed graphs, *SIAM J. Discrete Math.*, **27** (2013), 1429–1449.

[2] I.Z. Brouwer and I. Broere, Note on t-minimal complete bipartite graphs, *Canad. Math. Bull.*, **11** (1968), 729–732.

[3] L.W. Beineke, F. Harary and J.W. Moon, On the thickness of the complete bipartite graph, *Proc. Cambridge Philos. Soc.*, **60** (1964), 1–5.

[4] W. Cao, *Some Problems in Structural Graph Theory*, Ph.D. Thesis, University of Illinois (2006).

[5] Y.W. Chang, J.P. Hutchinson, M.S. Jacobson, J. Lehel and D.B. West, The bar visibility number of a graph, *SIAM J. Discrete Math.*, **18** (2004), 462–471.

[6] Y.C. Chen and X.L. Yin, The thickness of the Cartesian product of two graphs, *Canad. Math. Bull.*, **59(4)** (2016), 705–720.

[7] D. Eppstein, P. Kindermann, S. Kobourov, G. Liotta, A. Lubiw, A. Maignan, D. Mondal, H. Vosoughpour, S. Whitesides, and S. Wismath, On the planar split thickness of graphs. *Algorithmica*, **80** (2018), 977–994.

[8] S.K. Ghosh, *Visibility algorithms in the plane*, (Cambridge University Press, 2007).

[9] P.J. Heawood, Map colour theorem, *Quart. J. Math*, **24** (1890), 332–338.

[10] A.M. Hobbs and J.W. Grossman, A class of thickness-minimal graphs, *J. Res. Nat. Bur. Standards Sect. B*, **72B(2)** (1968), 145–153.

[11] J.P. Hutchinson, A note on rectilinear and polar visibility graphs, *Discrete Appl. Math.* **148** (2005), 263–272.

[12] G. Ringel and B. Jackson, Solution of Heawood's empire problem in the plane, *J. Reine Angew. Math.* **347** (1984), 146–153.

[13] R. Tamassia and I.G. Tollis, A unified approach to visibility representations of planar graphs, *Discrete Comput. Geom.*, **1** (1986), 321–341.

[14] W. Wessel, A short solution of Heawood's empire problem in the plane. *Discrete Math.* 191 (1998), 241–245.

[15] D.B. West and J.I. Wise, Bar visibility numbers for hypercubes and outerplanar digraphs, *Graphs and Combin.*, **33** (2017), 221–231.

[16] S. Wismath, Characterizing bar line-of-sight graphs, in Proceedings of 1st ACM Symposium on Comput. Geom., Baltimore, MD, 1985, 147–152.

[17] B. Xu, and X. Zha, Thickness and outerthickness for embedded graphs, *Discrete Math.*, **341** (2018), 1688–1695.