Bar visibility numbers for hypercubes and outerplanar digraphs

Douglas B. West* and Jennifer I. Wise†

July 19, 2016

Abstract

A t-bar visibility representation of a graph G assigns each vertex a union of at most t horizontal segments ("bars") in the plane so that vertices u and v are adjacent if and only if some bar assigned to u and some bar assigned to v are joined by an unobstructed vertical line of sight having positive width. For an oriented graph G, having an edge uv requires visibility from a bar for u upward to a bar for v. The visibility number of a graph or digraph G, written $b(G)$, is the least t such that G has a t-bar visibility representation.

For the n-dimensional hypercube Q_n, Euler's Formula yields $b(Q_n) \geq \lceil \frac{n+1}{4} \rceil$. To prove equality, we decompose Q_{4k-1} explicitly into k spanning subgraphs whose components have the form $C_4 \square P_{2^l}$. When G is an outerplanar oriented graph, always $b(G) \leq 2$; we give a forbidden substructure characterization for those with $b(G) = 1$.

1 Introduction

A *bar visibility representation* of a graph G assigns the vertices distinct horizontal line segments ("bars") in the plane such that $uv \in E(G)$ if and only if there is an unobstructed vertical line of sight (having positive width) joining the bar assigned to u and the bar assigned to v. A graph is a *bar visibility graph* if it has a bar visibility representation.

Inspired by a similar concept for interval graphs, Chang et al. [3] introduced t-bar visibility representations, where each vertex of G is represented by at most t bars, and $uv \in E(G)$ if and only if there is an unobstructed vertical line of sight (having positive width) joining some bar for u to some bar for v.

The model above uses undirected lines of sight. Axenovich et al. [2] introduced an analogous model for oriented graphs. In a t-bar visibility representation of an oriented graph G, there is an edge uv if and only some bar for u sees upward to some bar for v.

*Zhejiang Normal University and University of Illinois, dwest@math.uiuc.edu. Research supported by Recruitment Program of Foreign Experts, 1000 Talent Plan, State Admin. of Foreign Experts Affairs, China.

†Virginia Polytechnic Institute and State University, jiwise@vt.edu. Research partially supported by NSF grant DMS 08-38434, "EMSW21-MCTP: Research Experience for Graduate Students".

In both models, the *bar visibility number* $b(G)$ is the least t such that G has a t-bar visibility representation. Requiring positive width for lines of sight is important. It permits us to use closed bars so that bars with endpoints having a common x-coordinate cannot see each other but can block vertical visibility between them.

Early work used the name "visibility representation" for a slightly different model using zero-width lines of sight [1, 6, 7, 9]. The model we study here was introduced by Melnikov [8] under the name ϵ-visibility graphs, but it is now commonly called bar visibility representation. Wismath [10] and Tamassia and Tollis [9] independently gave a simple characterization of bar visibility graphs under this requirement. Hutchinson [4] later gave a simpler proof.

Theorem 1.1. *[4, 9, 10] A graph G is a bar visibility graph if and only if G has an embedding in the plane in which all cut-vertices lie on a single face.*

Let Q_n be the n-dimensional hypercube, defined by $V(Q_n) = \{0,1\}^n$ and $xy \in E(Q_n)$ if and only if x and y differ in exactly one coordinate. Given a t-bar representation of Q_n with minimum t, let $\hat{Q}_n$ be the planar graph formed by letting the bars be vertices and the lines of sight corresponding to $E(Q_n)$ be edges. Note that $\hat{Q}_n$ is triangle-free, since any triangle in $\hat{Q}_n$ would yield a triangle in Q_n. For a triangle-free planar graph G, Euler's Formula yields $|E(G)| \le 2|V(G)| - 4$. Note also that $t \ge |V(\hat{Q}_n)|/2^n$. Thus

$$b(Q_n) = t \ge \left\lceil \frac{|E(\hat{Q}_n)| + 4}{2 \cdot 2^n} \right\rceil = \left\lceil \frac{n \cdot 2^{n-1} + 4}{2 \cdot 2^n} \right\rceil = \left\lceil \frac{n+1}{4} \right\rceil.$$

Axenovich et al. [2] asked whether this trivial lower bound suffices, yielding $b(Q_n) = \lceil \frac{n+1}{4} \rceil$. Kleinert [5] decomposed Q_n into $\lceil \frac{n+1}{4} \rceil$ planar graphs. In light of Theorem 1.1, proving that his graphs are 2-connected answers the question in the affirmative.

Theorem 1.2. $b(Q_n) = \lceil \frac{n+1}{4} \rceil$ *for $n \in \mathbb{N}$.*

In Section 2 we obtain an explicit such decomposition; the decomposition is of independent interest. Let C_n denote the n-vertex cycle and P_n denote the n-vertex path. The cartesian product $G \square H$ of graphs G and H is the graph with vertex set $V(G) \times V(H)$ in which two vertices are adjacent if in one coordinate they are equal and in the other they are adjacent. For the case $n = 4k - 1$, our decomposition consists of spanning subgraphs $G_1, \ldots, G_k$ such that the components of G_i for $1 \le i < k$ are isomorphic to $C_4 \square P_{2^{i+1}}$, and $G_k \cong G_{k-1}$.

The product $C_4 \square P_{2^l}$ can be drawn using 2^l concentric 4-cycles joined by matchings, as shown in Figure 1. The labels will be used later to explain how the component is assembled as a subgraph of Q_n. Clearly $C_4 \square P_{2^l}$ is planar and 2-connected. Our graphs are isomorphic to those used by Kleinert, but our presentation is more intuitive and somewhat simpler.

Note that $b(D) \ge b(G)$ whenever D is an orientation of G. Our construction showing $b(Q_n) = \lceil \frac{n+1}{4} \rceil$ yields an orientation of Q_n where equality holds, answering another question of [2]. One can ask how large $b(D)-b(Q_n)$ can be when D is an orientation of Q_n. Although $b(Q_2)=b(Q_3)=1$, these graphs have orientations with bar visibility number 2.

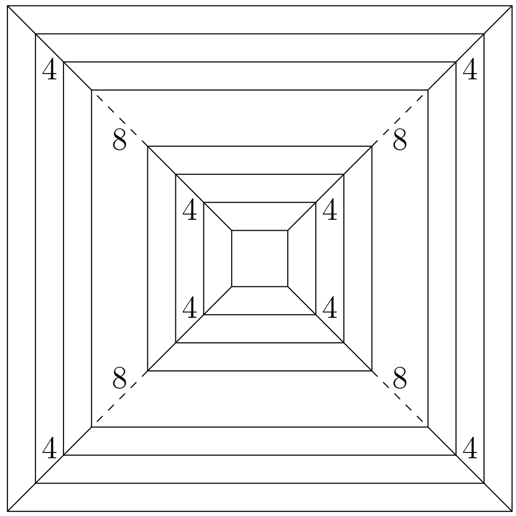

Figure 1

Oriented graphs with bar visibility number 1 are *bar-visibility digraphs*. These have a characterization related to bar-visibility graphs that uses the following notions.

Definition 1.3. In a digraph, a vertex with indegree 0 or outdegree 0 is a *source* or a *sink*, respectively. A *consistent cycle* is an orientation of a cycle having no source or sink. Given a digraph G, the *auxiliary digraph* G' is formed by adding two vertices s and t, a directed edge sv for every source vertex v in G, a directed edge wt for every sink vertex w, and the directed edge st.

Shrinking bars to vertices converts a bar visibility representation of a digraph G into a planar embedding of G; hence every bar visibility digraph is planar. Thomassia and Tollis [9] and Wismath [11] characterized the planar digraphs that are bar visibility digraphs.

Theorem 1.4. [9, 11] *A planar digraph G is a bar-visibility digraph if and only if G has no consistent cycle and its auxiliary digraph G' is planar.*

Axenovich et al. [2] proved $b(G)\le 2$ for every oriented outerplanar graph G. In Section 3, we give a forbidden substructure characterization for the oriented outerplanar graphs with $b(G)=1$, based on Theorem 1.4. Theorem 1.4 also computes $b(G)$ when G is an outerplanar digraph, via a check for acyclicity and a planarity test of the auxiliary digraph. In essense, our result analyzes how nonplanar subgraphs can arise in the auxiliary digraph of an outerplanar digraph.

2 Decomposition of Q_n

We first reduce the problem of computing $b(Q_n)$ to that of decomposing Q_n when $n \equiv 3 \bmod 4$, where there is no slack in the counting argument using Euler's Formula.

Lemma 2.1. *To prove Theorem 1.2, it suffices to decompose Q_{4k-1} into k planar subgraphs $G_1, \dots, G_k$ whose components are 2-connected, for $k \in \mathbb{N}$.*

Proof. Let $n = 4k - 1 + s$, where $1 \leq s \leq 3$. Given such a decomposition of Q_{4k-1}, decompose Q_n into 2^s copies of Q_{4k-1} and 2^{4k-1} copies of Q_s. Construct a $(k+1)$-bar visibility representation of Q_{4k-1+s} by placing 2^s copies of the representation of Q_{4k-1} in disjoint vertical strips and 2^{4k-1} copies of the bar visibility representation of Q_s in other disjoint vertical strips. Since $b(Q_s) = 1$ for $s \in \{1, 2, 3\}$, the extra bar allowed per vertex suffices for these representations. □

Our decomposition of Q_{4k-1} is described by allocating edges to subgraphs, based on the coordinates where the endpoints of edges differ and the constant values in other coordinates.

Definition 2.2. An edge $e \in E(Q_{4k-1})$ is of *type* r if its endpoints differ in coordinate r. Let E_r be the set of edges of type r. The edges of type $4j$ for $1 \leq j \leq k-1$ are *special edges*. For an edge e of type r and $s \neq r$, let e_s denote the common value that the endpoints of e have in coordinate s. For $e \in E_r$ with $r \leq 4k-4$, let e' denote the edge in Q_{4k-5} obtained by deleting the last four coordinates of the endpoints of e; the edge e' is the *truncation* of e, and e is an *extension* of e'.

We will decompose Q_{4k-1} into spanning subgraphs $G_1, \dots, G_k$. Example 2.3 shows the allocation of edges to subgraphs for $k \leq 5$. We will prove several properties inductively.

For $1 \leq i < k$, the spanning subgraph G_i has 2^n vertices in $2^{4(k-1)-i}$ components, each isomorphic to $C_4 \square P_{2^{i+1}}$; also $G_k \cong G_{k-1}$. The subgraph G_i contains $E_{4(k-i)+1} \cup E_{4(k-i)+2} \cup E_{4(k-i)+3}$. The resulting 3-dimensional subcubes are linked into larger components using special edges. For example, for $k > 2$ the components of G_2 are copies of $C_4 \square P_8$. As shown in Figure 1 for $k = 3$, these arise by combining four copies of Q_3 using eight edges of type $4k-8$ and four edges of type $4k-4$ (when $k = 3$ these are types 4 and 8, respectively).

To discuss $G_1, \dots, G_k$ for Q_{4k-1}, let $G'_1, \dots, G'_{k-1}$ be the specified decomposition of Q_{4k-5}. The key idea, illustrated in Example 2.3, is that for $i > 1$ the subgraph G_i for Q_{4k-1} is obtained from G'_{i-1}. Indeed, G_i begins with 16 copies of G'_{i-1}, extending its vertices by fixed choices in the four new coordinates. These copies will be linked in pairs using the new special type $4k-4$. It thus follows inductively that edges of all types other than $4k-4$ are used exactly once in the decomposition. We will need to prove this also for type $4k-4$ and allocate its edges to combine components in pairs.

Example 2.3. Decompositions of Q_3, Q_7, Q_{11}, Q_{15}, and Q_{19}:

$k{=}1$	
G_1	E_1 E_2 E_3

→ ↘

$k{=}2$		$e \in E_4$
G_1	E_5 E_6 E_7	$e_1 = e_5$
G_2	E_1 E_2 E_3	$e_1 \neq e_5$

↘ ↘

$k{=}3$		$e \in E_4$	$e \in E_8$
G_1	E_9 E_{10} E_{11}		$e_4 = e_9$
G_2	E_5 E_6 E_7	$e_1 = e_5$	$e_1 \neq e_5$ $e_4 \neq e_9$
G_3	E_1 E_2 E_3	$e_1 \neq e_5$	$e_1 = e_5$ $e_4 \neq e_9$

$k{=}4$		$e \in E_4$	$e \in E_8$	$e \in E_{12}$
G_1	E_{13} E_{14} E_{15}			$e_8 = e_{13}$
G_2	E_9 E_{10} E_{11}		$e_4 = e_9$	$e_4 \neq e_9$ $e_8 \neq e_{13}$
G_3	E_5 E_6 E_7	$e_1 = e_5$	$e_1 \neq e_5$ $e_4 \neq e_9$	$e_1 \neq e_5$ $e_4 = e_9$ $e_8 \neq e_{13}$
G_4	E_1 E_2 E_3	$e_1 \neq e_5$	$e_1 = e_5$ $e_4 \neq e_9$	$e_1 = e_5$ $e_4 = e_9$ $e_8 \neq e_{13}$

↘ ↘ ↘ ↘

$k{=}5$		$e \in E_4$	$e \in E_8$	$e \in E_{12}$	E_{16}
G_1	E_{17} E_{18} E_{19}				$e_{12} = e_{17}$
G_2	E_{13} E_{14} E_{15}			$e_8 = e_{13}$	$e_8 \neq e_{13}$ $e_{12} \neq e_{17}$
G_3	E_9 E_{10} E_{11}		$e_4 = e_9$	$e_4 \neq e_9$ $e_8 \neq e_{13}$	$e_4 \neq e_9$ $e_8 = e_{13}$ $e_{12} \neq e_{17}$
G_4	E_5 E_6 E_7	$e_1 = e_5$	$e_1 \neq e_5$ $e_4 \neq e_9$	$e_1 \neq e_5$ $e_4 = e_9$ $e_8 \neq e_{13}$	$e_1 \neq e_5$ $e_4 = e_9$ $e_8 = e_{13}$ $e_{12} \neq e_{17}$
G_5	E_1 E_2 E_3	$e_1 \neq e_5$	$e_1 = e_5$ $e_4 \neq e_9$	$e_1 = e_5$ $e_4 = e_9$ $e_8 \neq e_{13}$	$e_1 = e_5$ $e_4 = e_9$ $e_8 = e_{13}$ $e_{12} \neq e_{17}$

The constraints used to allocate E_{4k-4} to subgraphs will ensure the desired properties. Before allocating type $4k-4$ edges, G_1 consists of 2^{4k-4} copies of Q_3, which is $C_4 \square P_{2^1}$. To combine into copies of $C_4 \square P_{2^2}$, we need to use four edges 2^{4k-5} times, for a total of 2^{4k-3} edges. Since there are 2^{4k-2} edges of each type, we use half the edges of E_{4k-4}, which is accomplished by imposing one constraint on the coordinates.

With each increase in i, the size of the components doubles, and the number of components halves. Hence the number of special edges needed also halves. To obtain the desired number of edges of E_{4k-4}, we add one constraint at each step, and the constraints are satisfied by half of the remaining edges. The last step has the same number of constraints as the step before it, using the remaining edges of E_{4k-4}. Hence we allocate each edge of E_{4k-4} once, but we still must show that this produces the claimed subgraphs.

The inductive specification facilitates proof.

Construction 2.4. *Decomposition of* Q_{4k-1} We define spanning subgraphs $G_1, \ldots, G_k$ by specifying the edge sets, letting $F_i = E(G_i)$. For $k = 1$, let $F_1 = E_1 \cup E_2 \cup E_3 = E(Q_3)$. For $k > 1$, let $G'_1, \ldots, G'_{k-1}$ be the decomposition of Q_{4k-5}, with $F'_i = E(G'_i)$.

1. For $e \in E_r$ with $r < 4k-4$, put $e \in F_i$ if $e' \in F'_{i-1}$. Also put $E_{4k-1} \cup E_{4k-2} \cup E_{4k-3} \in F_1$.
2. For $k = 2$ and $e \in E_4$, put e in F_1 if $e_1 = e_5$, in F_2 if $e_1 \neq e_5$.
3. For $k > 2$ and $e \in E_{4k-4}$ with $e_{4k-8} = e_{4k-3}$, put $e \in F_1$.
4. For $k > 2$ and $e \in E_{4k-4}$ with $e_{4k-8} \neq e_{4k-3}$,
   (a) If $e_{4j-4} = e_{4j+1}$ for $i' < j \leq k-2$ and $e_{4i'-4} \neq e_{4i'+1}$, then put $e \in F_{k-i'}$ (here $i' \geq 2$).
   (b) If $e_{4j-4} = e_{4j+1}$ for $2 \leq j \leq k-2$, then put $e \in F_{k-|e_1-e_5|}$.

Theorem 2.5. *Construction 2.4 decomposes* Q_{4k-1} *into spanning subgraphs* $G_1, \ldots, G_k$ *such that each component of* G_i *is isomorphic to* $C_4 \square P_{2^{i+1}}$ *for* $i < k$ *and to* $C_4 \square P_{2^k}$ *for* $i = k$.

Proof. The proof is by induction on k. We first check that $G_1, \ldots, G_k$ is a decomposition. For $k = 1$, $Q_3 = G_1$. For $k \geq 2$, let $G'_1, \ldots, G'_{k-1}$ be the specified decomposition of Q_{4k-5}. Rule (1) allocates all types other than E_{4k-4}, putting edges into G_i for $1 \leq i \leq k$ forming a spanning subgraph whose components are isomorphic to Q_3 for $i = 1$ and to G'_{i-1} for $i > 1$.

To allocate E_{4k-4}, in Rules 3 and 4a of Construction 2.4 we impose i constraints on the edges to be used in F_i for $i < k-1$. Furthermore, the constraints for edges put in F_i are not satisfied by those put in $F_1, \ldots, F_{i-1}$. In Rule 4b, we allocate all the remaining edges, half to F_{k-1} and half to F_k. Thus each edge of E_{4k-4} is allocated exactly once, and $G_1, \ldots, G_k$ is a decomposition of Q_{4k-1}.

We must show that for $1 \leq i < k$, the edges of E_{4k-4} in G_i combine copies of $C_4 \square P_{2^i}$ into copies of $C_4 \square P_{2^{i+1}}$, and that in G_k they combine copies of $C_4 \square P_{2^{k-1}}$ into copies of $C_4 \square P_{2^k}$.

Let $H_i = C_4 \square P_{2^i}$. The vertices in the two copies of H_i that will be linked by four edges of E_{4k-4} are obtained by adding four coordinates to extend the names of the vertices in H_i. These two extensions will differ only in coordinate $4k-4$, which is one of the four new coordinates. This allows us to link them using edges of E_{4k-4}.

In H_i, we call the four copies of P_{2^i} the *diagonals*. A copy of H_i can be embedded in the plane with a specified end of the diagonals on the unbounded face or on the central face. The two copies we want to link have extensions differing only in coordinate $4k-4$. We embed them with one inside the other, so that the outer face of the inner copy has extensions of the same vertices whose extensions are on the inner face of the outer copy. Hence the vertices incident to the region bounded by the two embeddings are matched in Q_{4k-1} by edges of E_{4k-4}. Adding those edges creates copies of H_{i+1} in which the diagonals have twice as many vertices as in H_i, and the central edge of each diagonal is in E_{4k-4}.

It remains to show that the edges of type $4k-4$ we use to link these pairs are in fact the ones we have specified for G_i. Each pairing uses the following discussion. When each diagonal has exactly one edge of type r, the vertices on the outer face and those on the central face in an embedding of H_i have opposite values in coordinate r. If also the copy of H_i has no edge of type s, then an embedding of H_i has coordinates r and s agreeing in the vertices of one extreme face and disagreeing in the vertices of the other extreme face. Furthermore, either property can be chosen for the outer face.

For $k = 2$, the components of the graph G'_1 are copies of H_1 (which is Q_3) using types 1, 2, and 3. To form G_2, draw the copies of G'_1 with type 1 edges on the diagonal. Because coordinate 5 is constant on the two copies and the diagonal edge in each is type 1, we may embed the inside and outside copies of H_1 so that $e_1 \neq e_5$ for the endpoints of each edge e of type 4 used to link the two copies across the region between them.

We also combine copies of Q_3 to form G_1 for $k = 2$. The copies of Q_3 use edges of types 5, 6, and 7, with type 5 along the diagonal. We link two copies whose extensions differ only in coordinate 4. Since coordinate 1 is constant on the two copies of Q_3 being linked and the diagonal edges are type 5, we may choose to embed the inside and outside copies of Q_3 so that $e_1 = e_5$ for the endpoints of each edge e of type 4 used to link the two copies. Hence we have the claimed allocation for $k = 2$, as specified by Rule 2.

For $k \geq 3$, we form G_1 in almost the same manner as for $k = 2$. We take copies of Q_3 using edges of types $4k-1$, $4k-2$, and $4k-3$, with type $4k-3$ on the diagonals. Since coordinate $4k-8$ is constant on the copies of Q_3 being linked and the diagonal edges are type $4k-3$, we may embed the copies of Q_3 so that coordinates $4k-8$ and $4k-3$ are equal at vertices on the outer face of the inner copy of Q_3 and on the inner face of the outer copy of Q_3 linked to it. Thus we link the copies by edges of E_{4k-4} satisfying $e_{4k-8} = e_{4k-3}$, as specified in Rule 3. (The case $k = 2$ differs from this in using e_1, since $1 \neq 4 \cdot 2 - 8$.)

In forming G_1 we have used every edge $e \in E_{4k-4}$ such that $e_{4k-8} = e_{4k-3}$; there are 2^{4k-3} of them in 2^{4k-5} components. To allocate the edges satisfying $e_{4k-8} \neq e_{4k-3}$, we want to agree with Rule 4. By the induction hypothesis, for $2 \leq i \leq k$ the central edge of the diagonal in each component of G'_{i-1} has type $4k-8$. No other edges of G'_{i-1} have type $4k-8$, and coordinate $4k-3$ is constant on the vertices in components being paired. Hence by the usual discussion we embed the paired components so that the outer vertices of the inner copy and the inner vertices of the outer copy are matched via edges in E_{4k-4} satisfying $e_{4k-8} \neq e_{4k-3}$, agreeing with Rule 4.

Finally, to show that each edge of E_{4k-4} is used only once, we check that the remaining specified constraints on which subgraph contains which edges of E_{4k-4} are satisfied. We are checking Rule 4, involving only the edges of E_{4k-4} satisfying $e_{4k-8} \neq e_{4k-3}$. The condition in Rule 4b is vacuous when $k = 3$, so Rule 4b always puts edges into F_{k-1} and F_k, while Rule 4a puts edges into $F_2, \ldots, F_{k-2}$ for $k \geq 4$. For easier understanding, we recommend that the reader compare the arguments with Example 2.3.

A component of G_i is formed by combining extensions of two copies of G'_{i-1}. By the inductive construction, the central edge on each diagonal in G'_{i-1} is type $4k-8$. When $i = 2$, the edges of type $4k-8$ in G'_1 satisfy $e_{4k-12} = e_{4k-7}$. The other diagonal edges have type $4k-7$, and one of them is traversed to reach the edge of type $4k-4$. Thus the edges of type $4k-4$ in G_2 satisfy $e_{4k-12} \neq e_{4k-7}$, as specified by Rule 4a with $i' = k-2$.

To understand the constraints on the other edges, it is helpful to describe the list of types along each diagonal in a component of G_i; let this list be $L_k(i)$. By construction, $L_k(1) = 4k-3, 4k-4, 4k-3$. For $i \geq 2$, we have a recursive concatenation: $L_k(i) = L_{k-1}(i-1), 4k-4, L_{k-1}(i-1)$. Note that the special types on the diagonal are types $4k-4$ down to $4i'$, where $i' = k-i$. The key point follows inductively: for $2 < i \leq k-2$ and $i' < j \leq k-2$, between the central edge e of type $4k-4$ and the nearest edge of type $4j$ in either direction, there is exactly one edge of type $4j-4$. Note also that type $4i'$ is the lowest special type in G_i, and the special edges of type $4i'$ nearest to the central edge e of type $4k-4$ on the diagonal are separated from e only by an edge of type $4i'+1$.

By Rule 1, the edges of special type $4j$ in G_i were created when forming the $(k-2-j)$th subgraph in the decomposition of $Q_{4(j+1)-1}$. If $j = i'$, then the $(k-2-j)$th subgraph is the first. By Rule 3, the endpoints of these edges have the same value in coordinates $4i'-4$ and $4i'+1$. Since we follow an edge of type $4i'+1$ to reach the new edge e of type $4k-4$, we have $e_{4i'-4} \neq e_{4i'+1}$, as specified in Rule 4a. If $i' < j \leq k-2$, then the edges of type $4j$ are not introduced into the first subgraph in decomposing $Q_{4(j+1)-1}$, so their endpoints have different values in coordinates $4j-4$ and $4j+1$. Since we follow one edge of type $4j-4$ and no edges of type $4j+1$ in reaching the new edge e of type $4k-4$ on the diagonal, we have $e_{4j-4} = e_{4j+1}$, again as specified in Rule 4a.

Finally, we consider the edges of type $4k-4$ placed in F_{k-1} and F_k. In these graphs all special types occur. For $2 \leq j \leq k-2$, the reasoning is as above: the edges of type $4j$ are not introduced into a subgraph other than the first when decomposing $Q_{4(j+1)-1}$, so their endpoints have different values in coordinates $4j-4$ and $4j+1$. We follow one edge of type $4j-4$ and none of type $4j+1$ to reach the new edge e, and hence $e_{4j-4} = e_{4j+1}$.

To determine the last constraint on edge e for G_{k-1} and G_k, we consider the edges of type 4 nearest to e along the diagonal. Inductively, these edges were originally created in copies of the first or second subgraph when decomposing Q_7 (that is, $k=2$). In G_{k-1}, the types around e along the diagonal are $5, 4, 5, 4k-4, 5, 4, 5$ in order. In G_k, they are $1, 4, 1, 4k-4, 1, 4, 1$. The endpoints of the edges of type 4 have coordinates 1 and 5 the same in G_{k-1}, different in G_k. Since the edge e is separated from these edges by one edge of type 1 or 5, we have $e_1 \neq e_5$ in G_{k-1} and $e_1 = e_5$ in G_k, as specified in Rule 4b. $\square$

It is easy to describe an explicit bar visibility representation for $C_4 \square P_{2^i}$. We indicate the resulting representation for Q_{11} in Figure 2. Vertical cuts in horizontal segments indicate shared endpoints of bars.

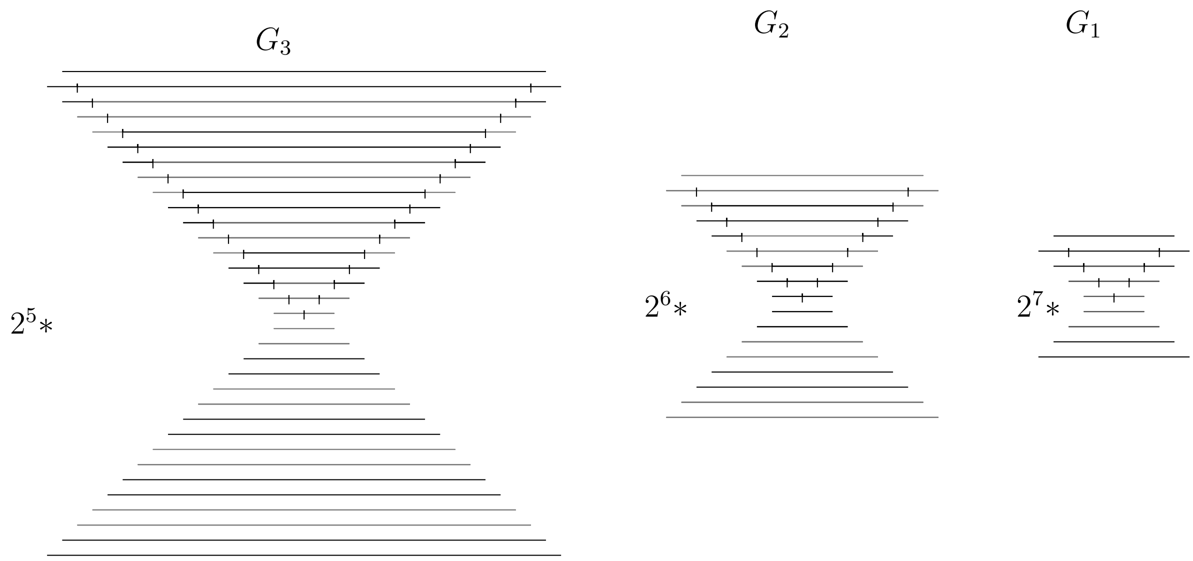

Figure 2: 3-Bar visibility representation for Q_{11}

3 Bar visibility outerplanar oriented graphs

We define several substructures that must be forbidden from bar-visibility digraphs. An *oriented cycle* is an orientation of a cycle, not necessarily a consistent cycle. A *claw* is a copy of the star $K_{1,3}$.

Definition 3.1. A *flower* in a digraph D consists of an oriented cycle C such that from three distinct vertices on C there are paths in the underlying graph to a sink and a source of D, and all six paths are disjoint except for the initial three vertices on C.

A *gear* in D consists of an oriented cycle C such that from four distinct vertices on C (in order) there are paths in the underlying graph to a source, a sink, a source, and a sink of D. These paths may have length 0.

A *tripod* in a digraph D consists of a claw such that the underlying graph has from each leaf of the claw a path to a source in D and a path to a sink in D; the two paths from one leaf need not be disjoint, but the paths from one leaf are disjoint from the paths from the other leaves (and avoid the center).

Axenovich et al. [2] also characterized the oriented trees T that are bar visibility graphs.

Theorem 3.2. [2] *An oriented tree is a bar-visibility graph if and only if it contains three internally disjoint inconsistent paths from a single vertex.*

In a tree, the forbidden condition in Theorem 3.2 is equivalent to the existence of a tripod.

Theorem 3.3. *An oriented tree is a bar-visibility graph if and only if it contains no tripod.*

Proof. We show that an oriented tree T contains three internally disjoint inconsistent paths from a single vertex if and only if it contains a tripod.

Given a tripod in T whose claw has center w, let v be a leaf of the claw. If the edge is oriented from v to w, then appending the tripod path from v to a sink yields an inconsistent path starting at w. If it is oriented from w to v, then appending the tripod path from v to a source yields an inconsistent path starting at w. Doing this for each leaf of the claw yields the three desired paths.

Given three inconsistent paths from w, from each path we obtain one leaf in the claw for a tripod. By symmetry, suppose that such a path begins with the directed edge wv. Following a consistent path from v eventually reaches a sink in T, since T has no cycle. The given inconsistent path from w contains a first edge yx oriented toward w. Following a consistent path from y (in reverse) eventually reaches a source in T. Hence we obtain two paths from v in the underlying graph (which may share an initial portion) to a sink and a source in T. □

Axenovich et al. [2] also noted that a bar visibility digraph cannot contain a consistent cycle. Our main result is that also forbidding the configurations of Definition 3.1 characterizes outerplanar digraphs that are bar visibility digraphs.

Theorem 3.4. *If G is an outerplanar digraph, then $b(G) = 1$ if and only if G contains no flower, gear, tripod, or consistent cycle.*

Proof. We first prove necessity.

If G contains a consistent cycle, then a bar visibility representation of G must put every bar in the cycle above the previous bar, which is impossible.

If G contains a flower, then in the auxiliary digraph G' of Theorem 1.4 the three given vertices on the cycle C plus s and t are the branch vertices of a K_5-subdivision. The paths in the subdivison are three on C, the edge st, and the six paths from the specified vertices C to a source or sink and then to s or t.

If G contains a gear, then in G' the four vertices on the cycle C plus s and t are the branch vertices of a $K_{3,3}$-subdivision. The branch vertices for one part of $K_{3,3}$ consist of s and the two vertices on the cycle whose specified paths lead to sinks; the other part consists of t and the two specified vertices whose paths lead to sources. We add four paths on C, the four paths to leading to sinks or sources and then to s or t, and the edge st itself.

If G contains a tripod, then in G' there is a $K_{3,3}$-subdivision whose branch vertices are s, t, and the center of the claw in one part, and the last common vertex on the specified paths to a source and a sink from each of the three leaves in the other part. The tripod provides internally disjoint paths from each vertex of one part to each vertex of the other.

For sufficiency, we suppose that G is not a bar visibility digraph and show that G contains a forbidden substructure. We may assume that G has no consistent cycle. Thus the auxiliary digraph G' must be nonplanar, containing in its underlying graph a subdivision of $K_{3,3}$ or K_5. Since G is outerplanar, G contains no subdivision of $K_{2,3}$ or K_4. Thus s and t must both be branch vertices in the subdivision of $K_{3,3}$ or K_5.

Suppose first that G' contains a $K_{3,3}$-subdivision H. If s and t are branch vertices in the same part, then let w be the third branch vertex in that part, and let Y be the set of branch vertices in the other part. The three edges incident to w in H form a claw that extends on to Y. From each vertex of Y, in H there is a path to a source followed by an edge from s and a path to a sink followed by an edge to t. Thus H contains a tripod (the two paths from a neighbor of w to a source and a sink run together until they reach Y).

If s and t are branch vertices in opposite parts, then let s_1 and s_2 be the other two branch vertices in the same part as s, and let t_1 and t_2 be the other two branch vertices in the same part as t. Note that s_1, t_1, s_2, and t_2 lie in order on a cycle C. In H there must be paths from s_1 and s_2 to t, and paths to t reach a sink just before t. Similarly, in H there are paths from t_1 and t_2 to s, reaching a source just before s. Hence we obtain a gear in G.

Finally, if G' contains a K_5-subdivision H, then let Z be the set of the three branch vertices other than s and t. Note that H contains a cycle C through Z. Leaving C, in H there are paths to s and t from each vertex of Z. These paths reach s and t via a source or sink vertex of G, respectively. Hence we obtain a flower in G. □

References

[1] Andreae, T.; Some results on visibility graphs. Discrete Appl. Math., 40 (1992), 5-17.

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

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

[4] Hutchinson J. P.; A note on rectilinear and polar visibility graphs. Discrete Applied Math. 148, 3 (2005), 263–272.

[5] Kleinert, M.; Die Dicke des n-dimensionalen Wrfel-Graphen. (German) J. Combinatorial Theory, 3 (1967), 10–15.

[6] Luccio F.; Mazzone S.; Wong C. K.; Visibility graphs. Prog. Naz. Teoria degli Algoritmi, Report 9 (University of Pisa, 1983).

[7] Luccio F.; Mazzone S.; Wong C. K.; A note on visibility graphs. Discrete Math., 64 (1987), 209–219.

[8] Melnikov L. S.; Problem at the 6th Hungarian Colloquium on Combinatorics, 1981.

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

[10] Wismath S. K.; Characterizing bar line-of-sight graphs. In Proc. ACM Symposium on Computational Geometry, Baltimore, MD (ACM, 1985), 147–152.

[11] Wismath S. K.; Bar-representable visibility graphs and a related network ow problem. PhD thesis, U. British Columbia, 1989.