A SUBLINEAR BOUND ON THE STACK NUMBER OF UPWARD PLANAR GRAPHS

MACIEJ MIKOŁAJCZAK

Abstract. The stack number of a directed graph G is the smallest number k for which there exist a k-colouring of the edges of G and a topological ordering of the vertices of G such that no two edges of the same colour cross in this ordering, i.e. have alternating endpoints along this ordering. A long-standing conjecture states that the stack number of upward planar graphs is bounded by a constant. Nevertheless, no non-trivial bounds on the stack number of upward planar graphs were established for many years. In 2023 Jungeblut, Merker, and Ueckerdt proved an upper bound of $\widetilde{\mathcal{O}}(n^{2/3})$ for n-vertex upward planar graphs. Their proof, while seemingly simple, turns out to be surprisingly subtle. In this thesis we present this proof in an effort to make it more transparent and accessible.

Undergraduate thesis

Advisor: dr hab. Piotr Micek

Theoretical Computer Science Department

Faculty of Mathematics and Computer Science

Jagiellonian University

June 19, 2026

Kraków, Poland

E-mail address: `pub@mikolajczak.dev`.

CONTENTS

1. Introduction

A *directed graph* is a pair $G = (V, E)$, where V is a set whose elements are called *vertices* and $E \subseteq V \times V$ is a set of *directed edges*. A *cycle* in a directed graph $G = (V, E)$ is a sequence of vertices $u_1, \ldots, u_k$ such that $(u_i, u_{i+1}) \in E$ for $i \in \{1, \ldots, k-1\}$ and $(u_k, u_1) \in E$. G is called *acyclic* if there is no cycle in G. A *topological ordering* of a directed graph G is a linear order $u_1, \ldots, u_n$ on the vertices of G such that if (u_i, u_j) is an edge of G, then $i < j$. It is known that a directed graph has a topological ordering if and only if it is acyclic.

A *planar drawing* of a directed graph $G = (V, E)$ is a representation where every vertex $v \in V$ corresponds to a point in the Euclidean plane, and every edge $(u, v) \in E$ is a curve connecting the points representing u and v. Furthermore, different edges have different sets of endpoints and the interior of an edge contains no vertex and no point of any other edge. A directed graph that admits such a drawing is called *planar*. An *upward planar* drawing of a directed graph G is a planar drawing such that every edge is a strictly y-monotone curve (increasing in the vertical direction, see the left of Figure 1). A directed graph that admits such a drawing is called *upward planar*. Notice that an upward planar graph is necessarily acyclic.

In a *stack embedding* (often referred to in the literature as a book embedding) of a directed acyclic graph $G = (V, E)$, the vertex set V is assigned a topological ordering $<$ (known as the *spine ordering*), and the edge set E is partitioned into so-called *stacks* (or pages) in such a way that no stack contains two edges (u_1, v_1), (u_2, v_2) that cross with respect to $<$. Specifically, edges cross if their endpoints are pairwise distinct and alternate along the topological ordering, meaning either $u_1 < u_2 < v_1 < v_2$ or $u_2 < u_1 < v_2 < v_1$ (see the right of Figure 1). The *stack number* (page number), denoted $\mathrm{sn}(G)$, is the minimum number of stacks required to admit such an embedding. Stack embeddings of graphs have found applications in contexts such as VLSI and fault-tolerant computing [6], sorting networks [21], and complexity theory [11].

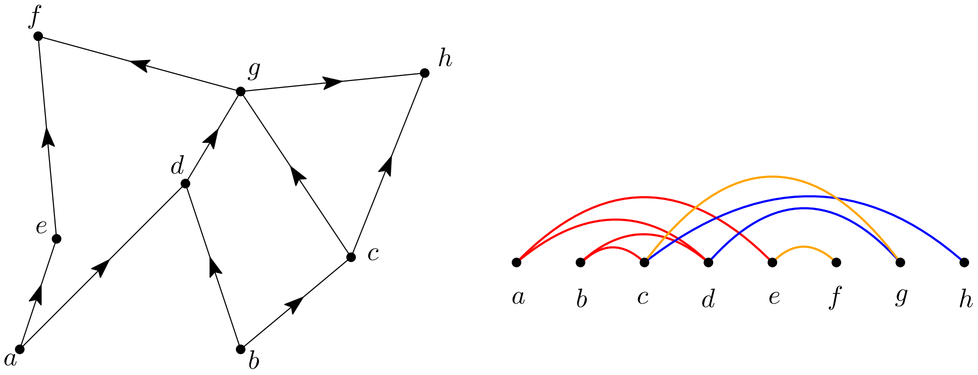

Figure 1. Left: An upward planar graph. Right: Its stack embedding into three stacks.

The concept of the stack number of undirected graphs (where the spine may be any arbitrary vertex ordering) was first introduced by Bernhart and Kainen in 1979 (by the name of book thickness) [4], building upon earlier notions suggested by Ollmann [20]. They conjectured that the stack number of planar graphs was unbounded. This, however, was quickly disproved [5], and it was subsequently established that any planar graph can be embedded into 4 stacks [22], a bound that has been proved to be best possible [23].

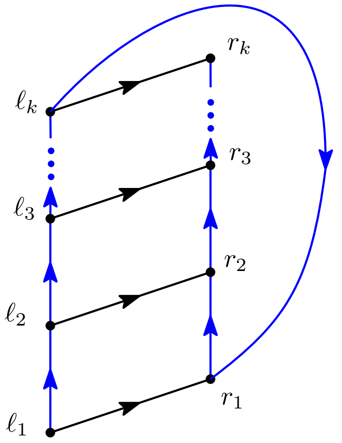

FIGURE 2. Example of a planar directed graph with large stack number. This graph has only one topological ordering (denoted in blue) and edges (ℓ_i, r_i) pairwise cross in this ordering.

In the case of directed graphs it is easy to construct families of graphs (even planar ones) that have unbounded stack number (see Figure 2). This naturally raises the question of what stronger structural conditions are required to guarantee a bounded stack number. The notion of stack embeddings of directed graphs was first considered by Nowakowski and Parker in 1989 [19]. They also introduced the stack number of a poset P by considering its cover graph and restricting the spine ordering strictly to a topological ordering of that cover graph, which is equivalent to a linear extension of P. This led them to pose a question:

Do posets with a planar diagram (equivalently, upward planar and transitively reduced graphs) have a bounded stack number?

Significant effort has been directed towards this question, exploring both the realm of posets [1, 2, 13] and directed graphs [12, 14, 18]. Notably, the stack number of outerplanar directed graphs has been shown to be bounded by a constant [16]. Despite these results, there has been little progress in bounding the stack number of upward planar graphs. Until recently, the best asymptotic bound on the stack number of n-vertex upward planar graphs was a trivial $\mathcal{O}(n)$ bound achieved simply by putting each edge in its own stack. Sublinear bounds were known only for very specific classes of graphs, such as upward planar triangulations with $o(n/\log n)$ diameter [10].

In 2023 Jungeblut, Merker, and Ueckerdt proved an upper bound of $\mathcal{O}(n^{2/3}\log^{2/3} n)$ [15]. To achieve this, they proposed an argument that bounds the stack number of the considered graphs in terms of their width, then in terms of their height, and then combines both approaches in order to prove a bound in terms of the number of vertices. This result, which remains the current state of the art, is the subject of this thesis.

2. PRELIMINARIES

Let $G = (V, E)$ be a directed acyclic graph. We write $V(G)$ and $E(G)$ to denote the set of vertices (respectively edges) of G. Let X be a subset of vertices of G. We denote the subgraph of G induced by X as $G[X]$. We shall refer to subsets of edges of directed acyclic graphs as *layers*. In the context of a fixed topological ordering, a layer is a *stack* if all edges in it are pairwise non-crossing with respect to this topological ordering.

A *path* P in a directed acyclic graph G is a sequence of vertices $(u_1, \ldots, u_k)$ of G such that $(u_i, u_{i+1}) \in E(G)$ for each $i \in \{1, \ldots, k-1\}$. We say the number $k-1$ is the *length* of P. We write $V(P)$ to denote the set of vertices of a path P. Notice that these vertices are pairwise distinct – if a vertex appeared twice in a path, then we would be able to close a cycle in the acyclic graph G. Moreover, P induces a linear order on $V(P)$. We make use of this fact by using phrases such as "the smallest vertex in P satisfying some condition" or "the vertex u precedes v along P".

If $P = (u_1, \ldots, u_k)$ is a path in a directed acyclic graph G, then for u_i, u_j such that u_i precedes u_j we define $u_i P u_j = (u_i, \ldots, u_j)$. This is the unique subpath of P originating in u_i and terminating in u_j. To compactly describe the concatenation of multiple subpaths and vertices, we list them sequentially by their shared endpoints. Formally, if P is a path containing the vertices u and v (where u precedes v along P), and Q is a path containing the vertices v and w (where v precedes w along Q), then $uPvQw$ denotes the unique path formed by concatenating the subpaths $(u, \ldots, v)$ of P and $(v, \ldots, w)$ of Q into a path $(u, \ldots, v, \ldots, w)$. The introduced notation naturally extends to any finite sequence of paths and intermediate vertices.

Let $G = (V, E)$ be a directed graph. A *subdivision* of an edge $e = (u, v)$ is the operation that replaces e in G by a path of length 2, that is, creates a graph $\widetilde{G} = (V \cup \{w\}, (E \setminus \{e\}) \cup \{(u, w), (w, v)\})$, where $w \notin V$ is a new vertex.

A *partial order*, or *poset*, on a set X is a pair $P = (X, \leqslant)$, where $\leqslant$ is a binary relation that is reflexive, antisymmetric, and transitive. Two elements $x, y \in X$ are called *comparable* if $x \leqslant y$ or $y \leqslant x$; otherwise, they are *incomparable*. We define the height $\mathrm{h}(P)$ and width $\mathrm{w}(P)$ of a poset P as the largest number of pairwise comparable (respectively incomparable) elements in P. We call a set of pairwise comparable elements of a poset a *chain*.

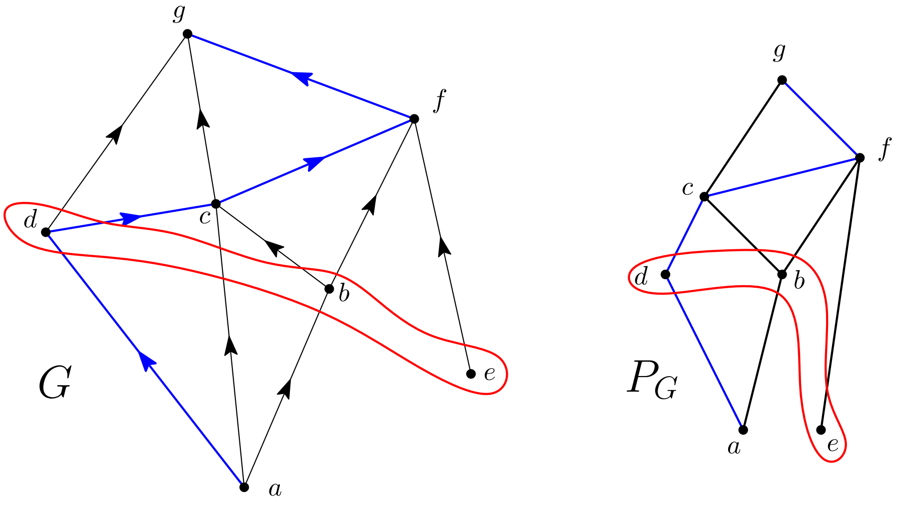

FIGURE 3. Left: An upward planar graph G of height 5 and width 3. Right: The reachability poset P_G of G.

We say that a vertex v of a directed acyclic graph G is *reachable* from a vertex u if there exists a path $(w_1, \ldots, w_k)$ in G such that $w_1 = u$ and $w_k = v$. We denote this by $u \prec_G v$ (often omitting the index when the context is clear) and write $u \preccurlyeq v$ to indicate that either $u \prec v$ or $u = v$. This reachability relation naturally defines a partial order on the vertex set V. Thus, we can define the reachability poset $P_G = (V, \preccurlyeq)$, consisting of the vertices of G partially ordered by their directed reachability. We say that two vertices u and v are *comparable* in G, respectively *incomparable* in G, if they are comparable, respectively incomparable, in P_G. We define the *height* $\mathrm{h}(G)$ as $\mathrm{h}(P_G)$ and *width* $\mathrm{w}(G)$ as $\mathrm{w}(P_G)$, see Figure 3. For a subset X of vertices of G its *height* $\mathrm{h}_G(X)$ and *width* $\mathrm{w}_G(X)$ are defined as the maximum number of vertices in X that are pairwise comparable (respectively incomparable) in G (including comparabilities that are witnessed by paths containing vertices outside of X). We omit the index if the context is clear.

In the context of a fixed drawing of a graph G, we denote the x-coordinate and y-coordinate of a vertex $v \in V(G)$ respectively by $x(v)$ and $y(v)$. We extend this notation to any object that is a point in the drawing.

Let $G = (V, E)$ be a directed acyclic graph. We call a vertex $s \in V$ a *source* if $s \preccurlyeq v$ for all $v \in V$. We call a vertex $t \in V$ a *sink* if $v \preccurlyeq t$ for all $v \in V$. A planar directed acyclic graph G is called an st-graph if it contains a unique source s and a unique sink t. An st-path in G is a directed path from s to t in G. It was shown by Kelly [17, Corollary 1] that any upward planar graph G can be augmented into an st-graph $\overline{G}$ on the same vertex set by adding some directed edges. Because adding edges to a graph can only increase its stack number (i.e., $\mathrm{sn}(G) \leqslant \mathrm{sn}(\overline{G})$), establishing an upper bound on the stack number of the st-graph $\overline{G}$ automatically yields an upper bound for the original upward planar graph G.

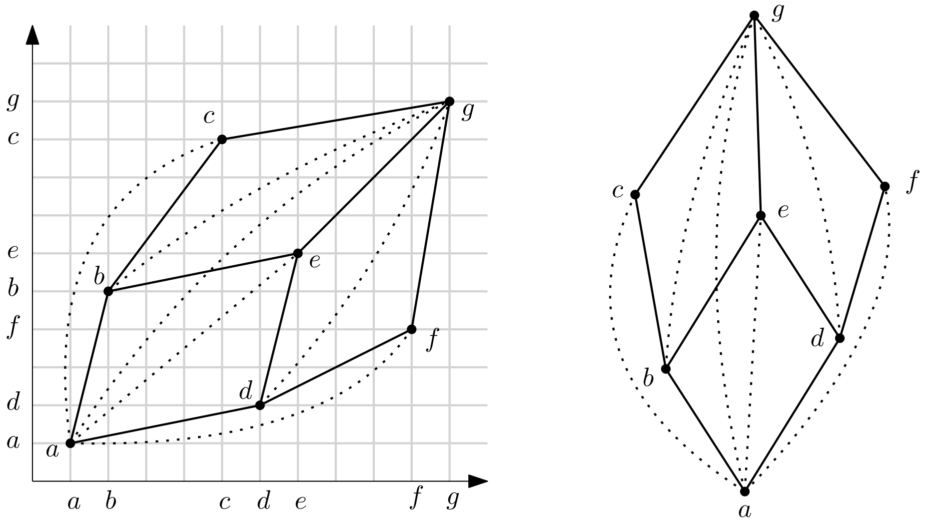

Figure 4. An st-graph together with its dominance drawing (left) and a more "upward-looking" drawing (right). The transitive edges are drawn dotted for better readability. We omit the arrows indicating direction, as all edges are pointed upward.

It is also known [3] that any st-graph admits a *dominance drawing*: a planar drawing such that between any two vertices u and v, there is a path from u to v if and only if $x(u) \leqslant x(y)$ and $y(u) \leqslant y(v)$ (see Figure 4). Moreover, this drawing can be realised with edges that are monotone polylines – chains of straight-line segments that are monotone both in the x- and y-coordinate. If there are two distinct vertices u, v such that $x(u) = x(v)$ in the dominance drawing, then we can slightly move one of them horizontally in such a way that the obtained drawing is still a dominance drawing and the x-coordinate of the moved vertex is different from the x-coordinate of all the other vertices. Performing a similar operation for all the vertices in both the x- and y-coordinates, we obtain a dominance drawing such that no two vertices have the same x-coordinate or y-coordinate. When we refer to a dominance drawing, we will implicitly assume it has this property. As such a dominance drawing is also an upward planar drawing, any st-graph is also an upward planar graph. In light of these observations, st-graphs can be thought of as the maximal upward planar graphs, representing the most challenging instances for which to bound the stack number in terms of the number of vertices. Consequently, we will be able to restrict ourselves to st-graphs in most of our proofs.

A concept closely related to the stack number is the twist number. For a given topological ordering $<$ of a directed acyclic graph G, a *k-twist* is a set of k edges that pairwise cross with respect to $<$ (see Figure 5). The *twist number* $\mathrm{tn}(G)$ of a directed acyclic graph is defined as the minimum k for which there exists a topological ordering of G containing no $(k+1)$-twist. Clearly, $\mathrm{tn}(G) \leqslant \mathrm{sn}(G)$, as in any topological ordering of G one can find a $\mathrm{tn}(G)$-twist, whose edges must be assigned to pairwise distinct stacks.

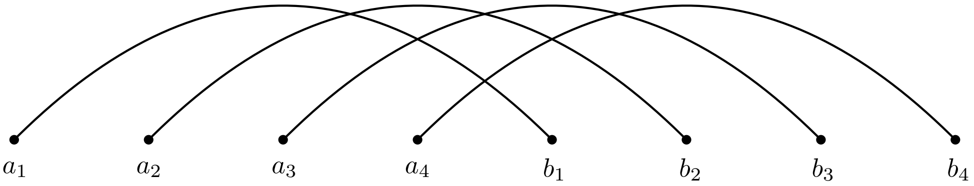

FIGURE 5. Edges (a_i, b_i) for $i = 1, \ldots, 4$ forming a 4-twist.

For a directed acyclic graph G and a fixed topological ordering $<$ of G one might consider the *crossing graph* $C_{G,<}$, which is an undirected graph on the edges of G with two vertices of $C_{G,<}$ connected by an edge when the corresponding directed edges of G cross with respect to $<$. A clique of size k in $C_{G,<}$ corresponds to a k-twist in G with respect to $<$. What is more, having already decided to take $<$ as the spine ordering, finding a stack embedding of G is equivalent to finding a proper colouring of $C_{G,<}$. The crossing graph $C_{G,<}$ belongs to an undirected graph family known as circle graphs. A *circle graph* is an undirected graph that can be represented as the intersection graph of a set of chords inside a circle. We can easily verify that $C_{G,<}$ is a circle graph: consider a circle and arrange the vertices of G along its boundary in the order dictated by $<$. Each directed edge $(u, v) \in E(G)$ can then be drawn as a straight-line chord connecting the boundary points u and v. Two chords intersect exactly when their endpoints alternate along the perimeter, which matches the definition of crossing with respect to $<$ (see Figure 6). Hence, the intersection graph of these chords is exactly $C_{G,<}$.

It is known that circle graphs are χ-bounded. In particular, the following asymptotically tight bound holds.

Lemma 1 (Davies [7])**.** If H is a circle graph with clique number at most ω, then the chromatic number of H is at most $2\omega \log(\omega) + 2\omega \log(\log(\omega)) + 10\omega \leqslant 14\omega \log(\omega)$.[1]

Together with the previous comments, this implies the following observation.

---

[1] All logarithms in this thesis are of base 2.

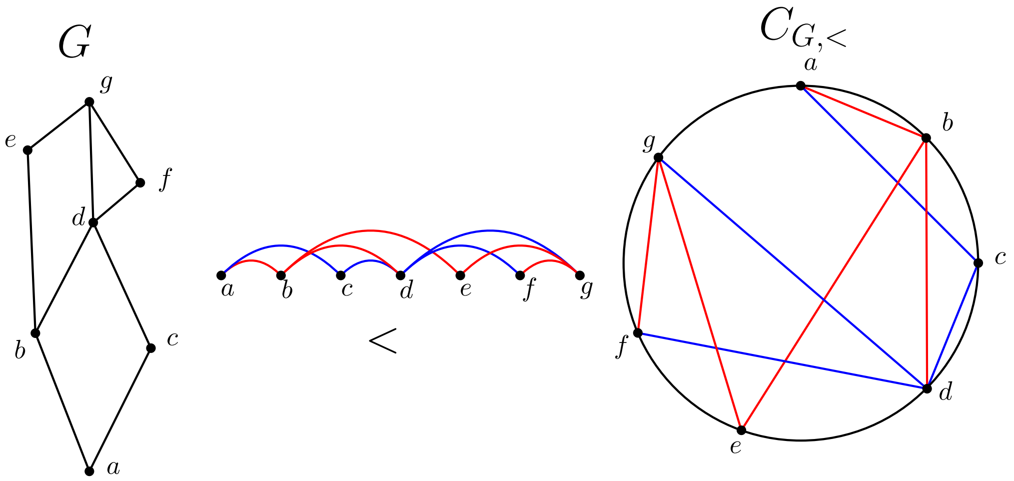

FIGURE 6. An upward planar graph G with its topological ordering $<$ and crossing graph $C_{G,<}$. A proper vertex colouring of $C_{G,<}$ corresponds to a stack embedding of G.

Observation 2. Let G be a directed acyclic graph. Then

$$\operatorname{sn}(G) \leqslant 14 \operatorname{tn}(G) \log(\operatorname{tn}(G)).$$

Moreover, this bound can be certified by a stack embedding using any vertex ordering with maximum twist size $\operatorname{tn}(G)$.

It is often more convenient to bound the twist number of a graph G instead of its stack number. Observation 2 implies that the two notions are functionally equivalent and it is actually enough to only consider bounds on the twist number.

Now we are ready to state the results from the work of Jungeblut, Merker, and Ueckerdt [15] that form the core of this thesis. We start with two technical lemmas.

Lemma 3. Let G be an st-graph and let X be a subset of its vertices of height at most h. Then G admits a topological ordering $<$ such that the size of every twist with respect to $<$ consisting of edges with at least one endpoint in X is at most $4h$.

Lemma 4. Let G be an st-graph and let X be a subset of its vertices of width at most w. There exists an st-graph G' such that

(1) $V(G) \subseteq V(G')$ and for every two vertices $u, v \in V(G)$ with $u \prec_G v$ we have $u \prec_{G'} v$. In other words, the poset $P_{G'}$ extends the poset P_G.

(2) the edges of $E(G'[X]) \cup (E(G) \setminus E(G'))$ can be partitioned into at most $14w$ layers that are stacks with respect to any topological ordering of G'.

We will later combine these results in order to prove the following theorem.

Theorem 5. For every upward planar graph G on n vertices, the stack number of G is $\mathcal{O}(n^{2/3} \log^{2/3} n)$.

3. Bounded height regime

In this section, we present the proof of Lemma 3. Recall that the height $\mathrm{h}(X)$ of a subset X of the vertices of a directed acyclic graph G is the maximal number of vertices in X that are pairwise comparable in G, and $\mathrm{h}(G)$ is the height of $V(G)$. Note that Lemma 3 does not easily imply a similar result for any upward planar graph, as in the process of augmenting an upward planar graph into an *st*-graph we add edges, and thus create new comparabilities. The height of the obtained *st*-graph can get much larger than the height of the starting upward planar graph.

In our proof, we will need the following classical result.

Lemma 6 (Erdős, Szekeres [9])**.** Given $r, s \geqslant 2$, any sequence of distinct real numbers with length at least $(r-1)(s-1)+1$ contains a monotonically increasing subsequence of length r or a monotonically decreasing subsequence of length s.

Proof of Lemma 3. Recall that as an *st*-graph, G admits a dominance drawing. Let $<_x$ and $<_y$ denote the orderings of vertices of G by their x-coordinate (respectively y-coordinate) in this drawing. For any $u, v \in V(G)$, we have $u \prec_G v \iff u <_x v$ and $u <_y v$. In particular, this implies that $<_x$ and $<_y$ are topological orderings of G. We will prove that $<_x$ satisfies the condition from the statement. It will be clear from the proof that $<_y$ would be an equally good choice.

Consider the largest twist (with respect to $<_x$) consisting of edges with at least one endpoint in X. These are edges $(a_1, b_1), \dots, (a_k, b_k)$ such that $a_i \in X$ or $b_i \in X$ for all $i \in \{1, \dots, k\}$ and $a_1 <_x \cdots <_x a_k <_x b_1 <_x \cdots <_x b_k$. Assume for the sake of contradiction that $k > 4h$. By the pigeon-hole principle, at least $\frac{k}{2}$ of the vertices $a_1, \dots, a_k$ are in X or at least $\frac{k}{2}$ of the vertices $b_1, \dots, b_k$ are in X. We assume the former, as it will be clear from our proof that the other case is symmetric. Thus, we have $a_{j_1}, \dots, a_{j_{2h}} \in X$ for some indices $j_1, \dots, j_{2h}$. By Lemma 6 applied to the sequence $(y(a_{j_1}), \dots, y(a_{j_{2h}}))$ one of the following is true:

(1) there exists a sequence $i_1 < \cdots < i_{h+1}$ of indices such that $a_{i_1} <_y \cdots <_y a_{i_{h+1}}$.

(2) there exists a sequence $i_1 < i_2 < i_3$ of indices such that $a_{i_1} >_y a_{i_2} >_y a_{i_3}$.

In the first case, we obtain $a_{i_1} <_x \cdots <_x a_{i_{h+1}}$ and $a_{i_1} <_y \cdots <_y a_{i_{h+1}}$. This means that $a_{i_1} \prec \cdots \prec a_{i_{h+1}}$ and these vertices are pairwise comparable. As they all belong to X, this means that $\mathrm{h}(X) \geqslant h+1$ – a contradiction with the lemma's assumptions.

Now assume the other case, that is, there exist vertices $a_{i_1}, a_{i_2}, a_{i_3}$ such that $a_{i_1} <_x a_{i_2} <_x a_{i_3}$ and $a_{i_1} >_y a_{i_2} >_y a_{i_3}$ (see the left of Figure 7). We know that $a_{i_2} <_x a_{i_3} <_x b_{i_1} <_x b_{i_2}$ and $a_{i_3} <_y a_{i_2} <_y a_{i_1} <_y b_{i_1}$. This implies that $a_{i_3} \prec b_{i_1}$. Hence, there exists a path P in G from a_{i_3} to b_{i_1} that is realised in the drawing as a monotone poly-line. Now consider the edge $e = (a_{i_2}, b_{i_2})$. We have $a_{i_2} <_x a_{i_3}$ and $a_{i_3} <_y a_{i_2}$. This means the monotone poly-line representing the edge e starts to the left of P and above it. $b_{i_1} <_x b_{i_2}$ implies e ends to the right of P. If e was above P at all times, then it would intersect the edge (a_{i_1}, b_{i_1}) in the drawing (consult the left of Figure 7). Thus, there exists a point on the poly-line e that is below the point on P with the same x-coordinate. As poly-lines are continuous, this implies that (a_{i_2}, b_{i_2}) and P must intersect in the drawing. This is a contradiction with the planarity of the drawing. Thus our initial assumption $k > 4h$ yields a contradiction. This completes the proof. □

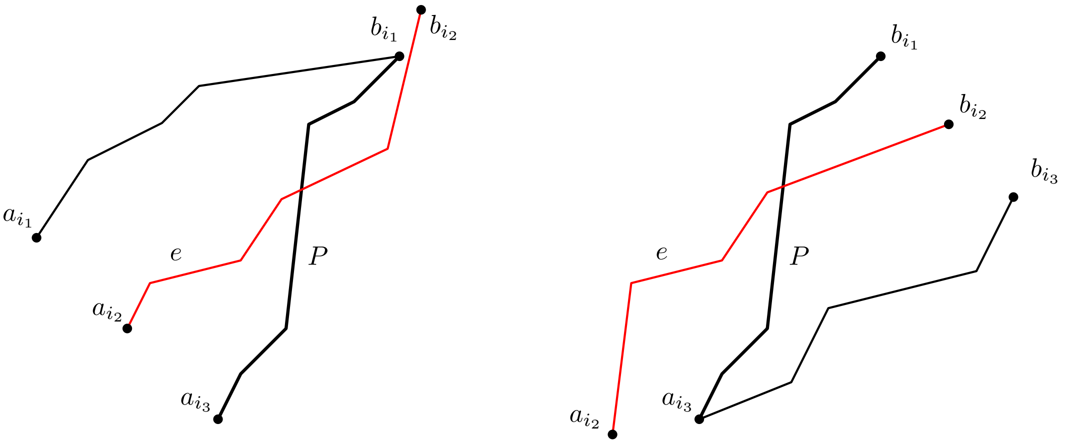

FIGURE 7. Left: The situation yielding the final contradiction in our proof, where we have $a_{i_1} >_y a_{i_2} >_y a_{i_3}$. Right: The symmetric case with $b_{i_1} >_y b_{i_2} >_y b_{i_3}$.

4. BOUNDED WIDTH REGIME

In this section, we present the proof of Lemma 4. We begin by introducing some topological notions that will be of use in the proof.

Lemma 7 (Dilworth [8]). Any poset P of width at most w can be decomposed into w chains.

Observation 8. Let G be an st-graph and let $X \subseteq V(G)$. All vertices of X can be covered by a family $\mathcal{P}$ of st-paths such that $|\mathcal{P}| = \mathrm{w}(X)$.

Proof. Consider the directed acyclic graph H with vertex set X and an edge from $u \in X$ to $v \in X$ if and only if $u \prec_G v$. We have $\mathrm{w}(P_H) = \mathrm{w}(X)$, where P_H is the reachability poset of H. By Lemma 7, P_H can be decomposed into $\mathrm{w}(X)$ chains $C_1, \dots, C_{\mathrm{w}(X)}$. A chain C_j is a set of vertices $u_1, \dots, u_k \in X$ such that $u_1 \prec_G \cdots \prec_G u_k$. For $i \in \{1, \dots, k-1\}$, let P_i be a path in G witnessing $u_i \prec_G u_{i+1}$. The path $P_j = u_1 P_1 u_2 \dots u_{k-1} P_{k-1} u_k$ contains all the vertices of C_j and can be easily extended to an st-path P'_j. The family $\mathcal{P} = \left\{ P'_j : j \in \{1, \dots, \mathrm{w}(X)\} \right\}$ satisfies all the requirements of the statement. ☐

Given an upward planar drawing of an st-graph G, we can define an ordering on the outgoing edges of any vertex $v \in V(G)$. Consider a ray originating at v that points directly to the left. Rotating this ray clockwise sweeps the upper half-plane (see the left of Figure 8). We call the order in which the outgoing edges of v appear in such a sweep a *left-to-right ordering* of these edges. If e, f are two outgoing edges of v, then we say e is to the left of f if e precedes f in the left-to-right ordering. In such a case, we say f is to the right of e. Similarly, we can define a left-to-right ordering of the incoming edges of v by considering the same ray and sweeping counterclockwise (see the right of Figure 8). We naturally extend the introduced notation to this ordering as well.

Let P and Q be two distinct st-paths in an st-graph G. The set of their common vertices $V(P) \cap V(Q)$ can be ordered sequentially along the paths as $s = u_1, u_2, \dots, u_k = t$.

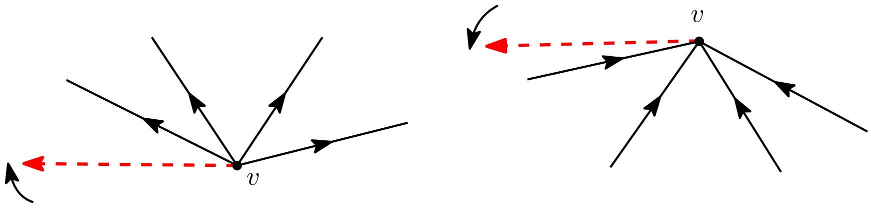

FIGURE 8. Left: The outgoing edges of v and the ray sweeping them. Right: The same for the incoming edges.

For one of these vertices u_j (where $u_j \neq t$), we say that it is a *divergence vertex* if the immediate next vertices of P and Q after u_j are different, that is, the paths branch apart at this vertex. In such a case, we call u_{j+1} a *convergence vertex* – this is the next common vertex of the paths P, Q after u_j. In other words, the paths converge at this vertex. Let $(u_{j_1}, u_{j_1+1}), \ldots, (u_{j_\ell}, u_{j_\ell+1})$ be the sequentially ordered list of pairs of divergence and convergence vertices. Note that there might be $u_{j_i+1} = u_{j_{i+1}}$, that is, a convergence vertex might also be a divergence vertex (see Figure 9). This list is non-empty, as the paths are distinct and must diverge at some point.

Let u_{j_i} be a divergence vertex and let u_{j_i+1} be its corresponding convergence vertex. If we consider an upward planar drawing of G, then the subpaths $u_{j_i}Pu_{j_i+1}$ and $u_{j_i}Qu_{j_i+1}$ enclose a region in the drawing. We denote this region by L_i and call it a *lens*. We call u_{j_i} the minimum of L_i and $u_{j_{i+1}}$ the maximum of L_i. We refer to the vertices of $u_{j_i}Pu_{j_i+1}$ and $u_{j_i}Qu_{j_i+1}$ as the *boundary* of the lens. The sequence $L_1, \ldots, L_\ell$ forms the *lens sequence* of the paths P and Q. We say that the lens L_i is P-oriented if the outgoing edge of P at u_{j_i} is to the left of the outgoing edge of Q in the left-to-right ordering. If this order is reversed, we say that L_i is Q-oriented.

We say that the paths P and Q *cross* if there exist two lenses L, L' in their lens sequence such that L is P-oriented and L' is Q-oriented (see Figure 9). We say that a family of st-paths $\mathcal{P}$ is *non-crossing* if no two paths $P, Q \in \mathcal{P}$ cross.

Observation 9. Let G be an st-graph and let $\mathcal{P}$ be a family of st-paths in G. There exists a non-crossing family of st-paths $\mathcal{P}'$ such that $\bigcup_{P\in\mathcal{P}} V(P) = \bigcup_{P\in\mathcal{P}'} V(P)$ and $|\mathcal{P}'| \leqslant |\mathcal{P}|$.

Proof. Let $k = |\mathcal{P}|$. For every edge $e \in E(G)$, its multiplicity $m(e)$ is defined as the number of paths in $\mathcal{P}$ that traverse e. For a vertex $v \in V(G)$, we define $d^+(v)$ as the sum of multiplicities of the outgoing edges of v. We define $d^-(v)$ as the sum of multiplicities of the incoming edges of v. We have $d^+(s) = d^-(t) = k$, $d^-(s) = d^+(t) = 0$, and $d^+(v) = d^-(v)$ for any $v \in V(G) \setminus \{s, t\}$.

We construct the new family $\mathcal{P}' = \{P'_1, \ldots, P'_k\}$ iteratively. To construct a path, we begin at the source s and move along the edges of the graph until we reach t. We choose the edges according to the following procedure: if the current vertex is v, then we examine the outgoing edges of v in their left-to-right ordering. We traverse the first edge with a positive multiplicity. Note that if we reached a vertex v, then $d^-(v) > 0$. Thus, if $v \neq t$, then $d^+(v) > 0$ and there exists an edge that we can move along. This proves that this procedure constructs an st-path. We take this path to be P'_1. We then decrease by one the multiplicity of every edge that P'_1 traverses. We repeat this procedure, producing paths $P'_2, P'_3, \ldots$ until $d^+(s) = 0$.

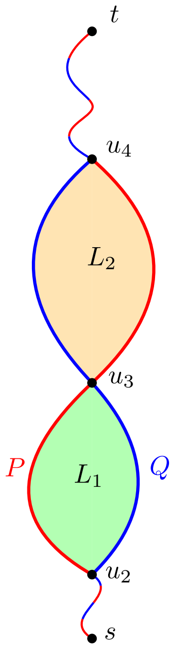

FIGURE 9. Two crossing st-paths P and Q. u_2 is a divergence point, u_3 is both a divergence point and a convergence point, and u_4 is a convergence point. The lens sequence consists of two lenses L_1 and L_2. L_1 is P-oriented and L_2 is Q-oriented.

Note that the described procedure maintains the invariant $d^+(v) = d^-(v)$ for any $v \in V(G) \setminus \{s,t\}$. This implies that after the procedure stops, we have $d^-(v) = 0$ for all $v \in V(G)$. Indeed, if $d^-(v_0) > 0$ for a vertex v_0, then there exists an edge (v_1, v_0) of positive multiplicity. If $v_1 \neq s$, then $d^-(v_1) = d^+(v_1) > 0$. Repeating this argument, we obtain a sequence of distinct vertices $v_0, v_1, \ldots$. This sequence has to be finite. This means this process terminates and we have $v_\ell = s$ for some ℓ. But then $d^+(v_\ell) > 0$, a contradiction.

Let us denote $X = \bigcup_{P \in \mathcal{P}} V(P)$. When we begin our construction, we have $d^+(s) = k$ and $d^-(v) > 0$ for any vertex $v \in X \setminus \{s\}$. When it ends, we have $d^+(s) = 0$ and $d^-(v) = 0$ for any vertex $v \in V(G)$. This means we constructed exactly k paths (though some of them might be the same path) and $|\mathcal{P}'| \leqslant k$. Moreover, $X \subseteq \bigcup_{P \in \mathcal{P}'} V(P)$. It is evident that $\bigcup_{P \in \mathcal{P}''} V(P) \subseteq X$.

It remains to show that $\mathcal{P}'$ is non-crossing. Let P'_i, P'_j be two paths of $\mathcal{P}'$ such that $i < j$ (this means P'_i was constructed before P'_j). Let u be a divergence vertex of these paths. By our construction, the outgoing edge of P'_i at u is to the left of the outgoing edge of P'_j. This means the lens induced by u and its corresponding convergence vertex is P'_i-oriented. As u was an arbitrary divergence vertex, we get that P'_i and P'_j do not cross. $\square$

If P is a path in an st-graph G and y_0 is such that $y(s) \leqslant y_0 \leqslant y(t)$, then we denote by P_{y_0} the unique point on P such that $y(P_{y_0}) = y_0$ (assuming such a point exists). Note that if P is an st-path, then P_{y_0} necessarily exists, as P is a continuous curve between s and t.

Observation 10. Let G be an st-graph and let P be an st-path. Let Q be any path. Let y_0, y_1 be such that $y(s) \leqslant y_0, y_1 \leqslant y(t)$ and both Q_{y_0}, Q_{y_1} exist. If $x(Q_{y_0}) \leqslant x(P_{y_0})$ and $x(Q_{y_1}) \geqslant x(P_{y_1})$, then the paths P and Q intersect at a vertex v such that $y(v)$ is between y_0 and y_1.

Proof. The statement follows from the continuity of y-monotone curves P, Q and the planarity of the drawing. $\square$

Observation 11. Let G be an st-graph and let P, Q be two st-paths that do not cross. If $x(P_{y_0}) < x(Q_{y_0})$ for some y_0 such that $y(s) \leqslant y_0 \leqslant y(t)$, then $x(P_y) \leqslant x(Q_y)$ for all $y(s) \leqslant y \leqslant y(t)$.

Proof. Let v be the last divergence vertex of P and Q such that $y(v) < y_0$. Assume its corresponding lens L is Q-oriented, that is, the outgoing edge (v, u) of P is right of the outgoing edge (v, u') of Q in the left-to-right ordering. This implies $x(u) > x(u')$ and by Observation 10 the paths P and Q converge at a vertex w such that $y(w) < y_0$. This contradicts the choice of v. Thus, L is P-oriented.

Assume for a contradiction that there exists y_1 such that $y(s) \leqslant y_1 \leqslant y(t)$ and $x(P_{y_1}) > x(Q_{y_1})$. Let v' be the last divergence vertex of P and Q such that $y(v') \leqslant y_1$. Let L' be its corresponding lens. Similarly as before, we show that L' is Q-oriented. This is a contradiction, as P and Q do not cross. This completes the proof. $\square$

Let $\mathcal{P}'$ be a non-crossing family of st-paths. Observation 11 allows us to define a natural ordering of the paths in $\mathcal{P}'$, which we shall refer to as the left-to-right ordering of $\mathcal{P}'$. Let $P, Q \in \mathcal{P}'$ be st-paths. We say that P is left of Q when $x(P_y) \leqslant x(Q_y)$ for all y such that $y(s) \leqslant y \leqslant y(t)$ (see Figure 10). We also say that Q is right of P. It is evident from Observation 11 that this is indeed a linear ordering of the paths in $\mathcal{P}'$.

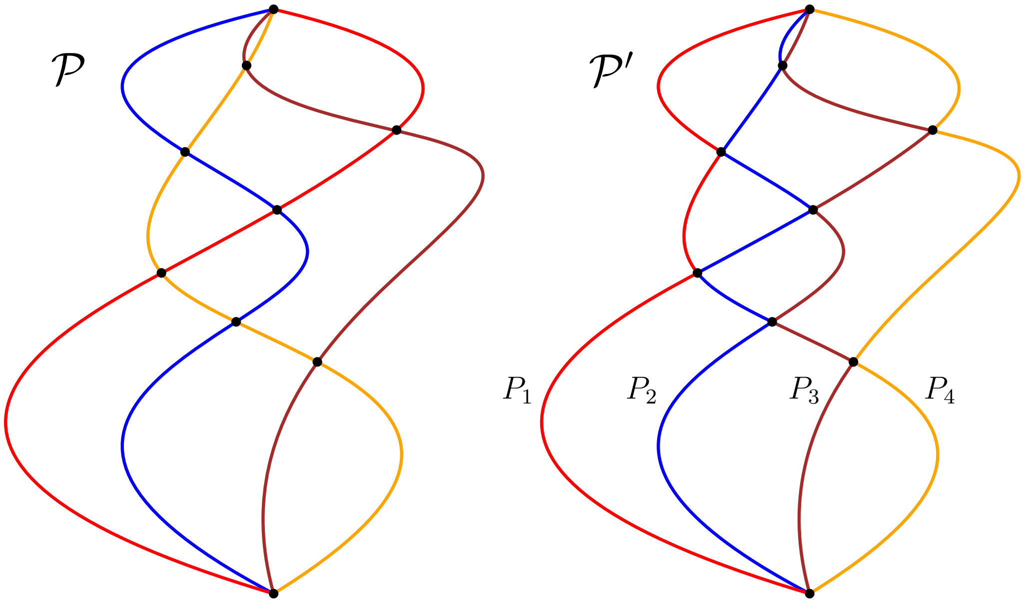

FIGURE 10. A family of st-paths $\mathcal{P}$ and a corresponding non-crossing family $\mathcal{P}' = \{P_1, P_2, P_3, P_4\}$.

For a fixed non-crossing family of st-paths $\mathcal{P}$, let us denote by $X = \bigcup_{P \in \mathcal{P}} V(P)$ the set of all vertices covered by $\mathcal{P}$. For an edge $e \in E(G[X])$, we say that e is an *intra-path edge* if there exists a path $P \in \mathcal{P}$ such that both endpoints of e belong to $V(P)$. If this is not the case, that is, e connects two different paths, then we call e an *inter-path edge*.

Observation 12 ([4])**.** Let G be an st-graph and let P be an st-path. The edges of $G[V(P)]$ can be partitioned into two layers E_1, E_2 that are stacks with respect to any topological ordering of G.

Proof. Notice that the reachability poset of a path P is a linear order. This means that there exists only one topological ordering of P (the natural ordering of vertices along P) and every topological ordering of G restricted to these vertices is this ordering.

Let $e = (u, v)$ be an edge of $G[V(P)]$. For y such that $y(u) \leqslant y \leqslant y(v)$ let e_y be the unique point on e such that $y(e_y) = y$. By Observation 10, either $x(e_y) \leqslant x(P_y)$ for all $y(u) \leqslant y \leqslant y(v)$, or $x(e_y) \geqslant x(P_y)$ for all $y(u) \leqslant y \leqslant y(v)$. If the first condition is true, then we place the edge in E_1. Otherwise we place it in E_2. Notice that the only edges that satisfy both conditions are the edges of P. We place them either in E_1 or in E_2, arbitrarily.

Two edges of E_i (for $i = 1, 2$) that are not edges of P cross with respect to the topological ordering of P if and only if they intersect in the drawing (see Figure 11). As the drawing is planar and the edges of P cannot cross any other edges in this ordering, E_i is a stack. □

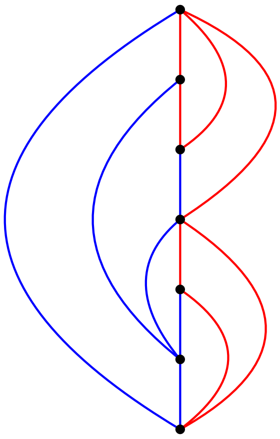

FIGURE 11. An st-path P with the edges of $G[V(P)]$ partitioned into stacks.

Observation 13. Let G be an st-graph and let $\mathcal{P} = \{P_1, \dots, P_k\}$ be a family of pairwise non-crossing st-paths indexed according to their left-to-right ordering. If $e = (u, v)$ is an inter-path edge, then $u \in V(P_i)$ and $v \in V(P_j)$ for some $i, j \in \{1, \dots, k\}$ with $|i - j| = 1$. Moreover, in an upward planar drawing of G, e is contained in some lens in the lens sequence of P_i and P_j with its endpoints on the boundary of this lens.

Proof. Let $e = (u, v)$ be an inter-path edge. Let $i, j \in \{1, \dots, k\}$ be such that $u \in V(P_i)$, $v \in V(P_j)$ and $|i - j|$ is minimal. Assume $|i - j| > 1$. This means there exists an st-path P_d such that P_i is left of P_d and P_d is left of P_j. Observation 10 implies that the edge e must intersect the path P_d in the drawing. By the planarity of the drawing, this intersection happens in an endpoint of e (see the left of Figure 12). This means $u \in V(P_d)$ or $v \in V(P_d)$. Both cases give a contradiction with the definition of i, j, as $i < d < j$. This implies $|i - j| \leqslant 1$. $|i - j| \geqslant 1$ follows from the fact that e is not an intra-path edge.

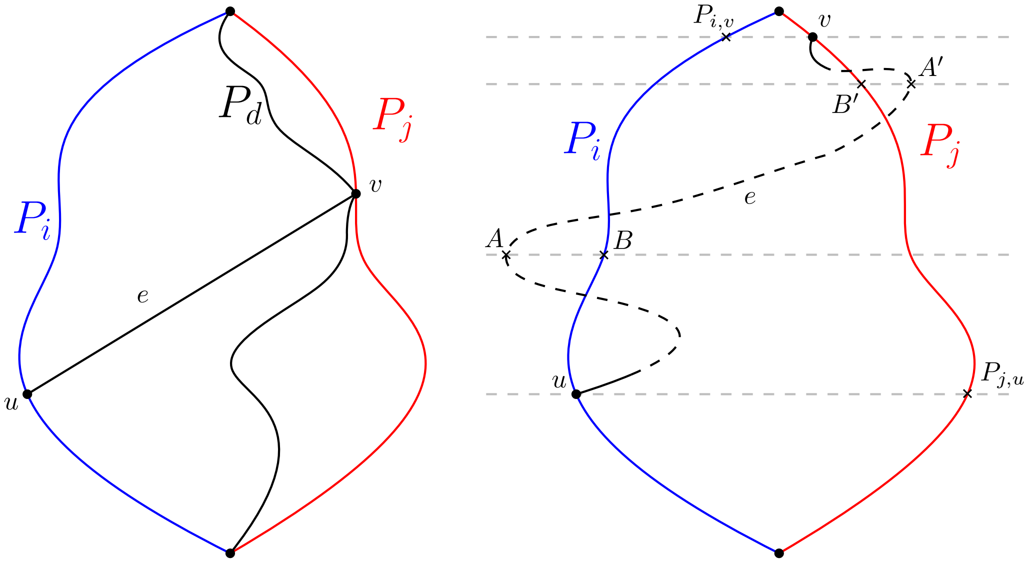

Figure 12. Left: The edge e intersecting both the paths P_d and P_j in its endpoint. Right: The edge e and the points A, B and A', B' yielding a contradiction.

Assume without loss of generality that $j = i + 1$, that is, P_i is left of P_j. Let $P_{j,u}$ be the point on P_j such that $y(P_{j,u}) = y(u)$. Let $P_{i,v}$ be the point on P_i such that $y(P_{i,v}) = y(v)$. We have $x(u) < x(P_{j,u})$ and $x(v) > x(P_{i,v})$ (the strict inequalities follow from the fact that e is not an intra-path edge). Assume there exists a point A on the edge e and a point B on P_i such that $y(A) = y(B)$ and $x(A) < x(B)$. Together with $x(v) > x(P_{i,v})$, applying Observation 10 implies that e intersects P_i at a point different than u (see the right of Figure 12). This is a contradiction. Assume there exists a point A' on the edge e and a point B' on P_j such that $y(A') = y(B')$ and $x(A') > x(B')$. Together with $x(u) < x(P_{j,u})$, applying Observation 10 implies that e intersects P_j at a point different than v. This is a contradiction. Thus, for every point C on the edge e its x-coordinate is between the x-coordinates of its corresponding points on P_i and P_j. We know that e cannot intersect the paths P_i, P_j at points other than the endpoints of e. This implies that u and v lie on the boundary of the same lens in the lens sequence of P_i and P_j and e is contained in this lens. ☐

Observation 14. Let G be an st-graph and let P, Q be two st-paths. Let L_1, L_2 be two lenses in the lens sequence of these paths. If the edge e_1 is contained in L_1 and the edge e_2 is contained in L_2, then these edges do not cross with respect to any topological ordering of G.

Proof. Assume without loss of generality that L_1 precedes L_2 in the lens sequence of P and Q. Let $e_1 = (v_1, w_1)$ be an edge contained in L_1. There exists a path P_1 connecting w_1 to the sink t. P_1 must intersect the boundary of L_1. Thus, we have $w_1 \preccurlyeq b_1$, where b_1 is the maximum of L_1. Let $e_2 = (v_2, w_2)$ be an edge contained in L_2. There exists a path P_2 connecting the source s to v_2. P_2 must intersect the boundary of L_2. Thus, we have $a_2 \preccurlyeq v_2$, where a_2 is the minimum of L_2. This implies $v_1 \prec w_1 \preccurlyeq b_1 \preccurlyeq a_2 \preccurlyeq v_2 \prec w_2$. These vertices appear in this order in any topological ordering of G. Thus, the edges (v_1, w_1) and (v_2, w_2) do not cross with respect to any such ordering. □

With these observations at hand, we are ready to prove Lemma 4. We restate it here for convenience.

Lemma 4. Let G be an st-graph and let X be a subset of its vertices of width at most w. There exists an st-graph G' such that

(1) $V(G) \subseteq V(G')$ and for every two vertices $u, v \in V(G)$ with $u \prec_G v$ we have $u \prec_{G'} v$. In other words, the poset $P_{G'}$ extends the poset P_G.

(2) the edges of $E(G'[X]) \cup (E(G) \setminus E(G'))$ can be partitioned into at most $14w$ layers that are stacks with respect to any topological ordering of G'.

Proof. By Observation 8 and Observation 9, there exists a non-crossing family $\mathcal{P}$ of st-paths such that $X \subseteq \bigcup_{P \in \mathcal{P}} V(P)$ and $k = |\mathcal{P}| \leqslant w$. Let $P_1, \ldots, P_k$ denote these paths, indexed according to their left-to-right ordering. Our strategy is to construct a family of layers that covers all edges between these paths (both the intra-path edges and the inter-path edges). Applying Observation 12 to each path individually, we obtain a family $\mathcal{I}$ of $2k$ layers such that all the intra-path edges of these paths belong to some layer and these layers are stacks with respect to any topological ordering of G. Now we deal with the inter-path edges, which require a more sophisticated argument.

We will incrementally construct the st-graph G' mentioned in the statement. We start by setting $G' = G$. In the course of the proof we will modify G' by subdividing its edges and adding new edges. We will do it in such a way that G' remains planar and acyclic. Thus, G' will be an st-graph. Recall that the definition of st-graphs does not require an upward planar drawing, but only a planar one. Thus, we will start from an upward planar drawing of G and produce from it a planar drawing of G'. It is easy to see that a graph obtained from G by subdividing edges and adding edges satisfies (1). In particular, any topological ordering of G' will be a topological ordering of G (when restricted to the vertices of G). Notice that subdividing an edge of $P \in \mathcal{P}$ requires us to add the new vertex to P in order for it to remain a path. Thus, it will be convenient for us to assume the convention that whenever we subdivide an edge of a path $P \in \mathcal{P}$, we add the newly created vertex to this path in between the vertices connected by the subdivided edge.

Let $P, Q \in \mathcal{P}$ with $P = P_i$ and $Q = P_{i+1}$ for some $i \in \{1, \ldots, k-1\}$. Let L be a lens in the lens sequence of P and Q. Let a and b be the minimum and maximum of L, respectively. Our goal is to partition the edges with endpoints on the boundary of L into 12 layers. We will then use our previous observations about lenses and inter-path edges to construct a family of layers satisfying (2). Let $\overrightarrow{E}$ (respectively, $\overleftarrow{E}$) denote the set of edges contained in L with both endpoints on the boundary of L such that each edge originates at a vertex of P (respectively, a vertex of Q) and terminates at a vertex of Q (respectively, a vertex of P). We mostly consider the edges $\overrightarrow{E}$, as $\overleftarrow{E}$ can be handled separately in a symmetric manner.

Let $v_1, \ldots, v_\ell$ be the sequence of vertices of P (indexed according to the natural ordering of P) such that for each v_i, there exists at least one edge of $\overrightarrow{E}$ originating in v_i. For

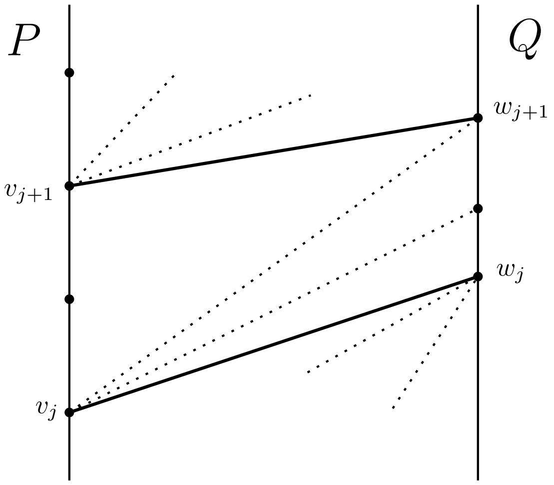

FIGURE 13. The edges (v_j, w_j) and (v_{j+1}, w_{j+1}). Other edges of $\overrightarrow{E}$ are draw dotted.

a fixed v_i, let w_i be the smallest vertex in Q such that $(v_i, w_i) \in \overrightarrow{E}$. Notice that if $(v_j, w) \in \overrightarrow{E}$ for some $w \in Q$ and $j < i$, then w precedes w_i along Q. Otherwise the edges (v_j, w) and (v_i, w_i) would intersect in the drawing (see Figure 13).

For each edge $e_i = (v_i, w_i)$, we perform the following procedure. We subdivide the edge of P going out of v_i and subdivide the edge of Q going into w_i. We obtain new vertices v'_i and w'_i on P and Q, respectively. Notice that the existence of e_i implies $v_i \neq b$ and $w_i \neq a$. This means the edges we subdivided were well-defined. Let $E_{v,i}$ be the set of outgoing edges of v_i contained in L that are to the left of e_i. Let $E_{w,i}$ be the set of incoming edges of w_i contained in L that are to the right of e_i. Now we subdivide the edge $e_i = (v_i, w_i)$, turning it into a path (v_i, u_i, w_i). Let $(\alpha_1, w_i), \ldots, (\alpha_d, w_i)$ be the edges of $E_{w,i}$ indexed according to the reversed left-to-right ordering. We subdivide each (α_j, w_i) into two edges $(\alpha_j, q_j), (q_j, w_i)$. We add the edges $(w'_i, q_1), (q_1, q_2), \ldots, (q_{d-1}, q_d), (q_d, u_i)$ to the graph. Similarly, let $(v_i, \beta_1), \ldots, (v_i, \beta_h)$ be the edges of $E_{v,i}$ indexed according to the reversed left-to-right ordering. We subdivide each (v_i, β_j) into two edges $(v_i, p_j), (p_j, \beta_j)$. We add the edges $(u_i, p_1), (p_1, p_2), \ldots, (p_{h-1}, p_h), (p_h, v'_i)$ to the graph. The new edges form a path from w'_i to v'_i. Let us denote it by T. In the drawing we can place T arbitrarily close to the path (v_j, u_j, w_j). This implies we can assume it does not intersect any edges other than the ones in $E_{v,i}$ and $E_{w,i}$. These edges were removed from the graph – we subdivided them in such a way that the obtained edges do not intersect with T (in points other than their common vertices). Thus, the obtained graph is planar. Notice that the obtained drawing need not be upward planar. In fact, the path T is almost certainly a decreasing curve, as in the drawing it looks like the path (v_i, u_i, w_i) going in reverse. The entire construction is presented on the left of Figure 14.

Now we prove that adding the path T to the graph does not create a cycle. We denote $q_0 = w'_i$, $p_0 = u_i$, and $p_{h+1} = v'_i$. We have $T = (q_0, \ldots, q_d, p_0, \ldots, p_{h+1})$. Assume there exists a cycle in the graph that contains an edge of T. Let C be the shortest such cycle. Let x and y be respectively the first and the last vertex of T contained in C. We consider a few cases.

- Assume $x = p_j$ and $y = p_{j'}$ for $0 \leqslant j < j' \leqslant h+1$. These vertices have only two incoming edges each: the edge from v_i and the previous edge of T. v_i cannot

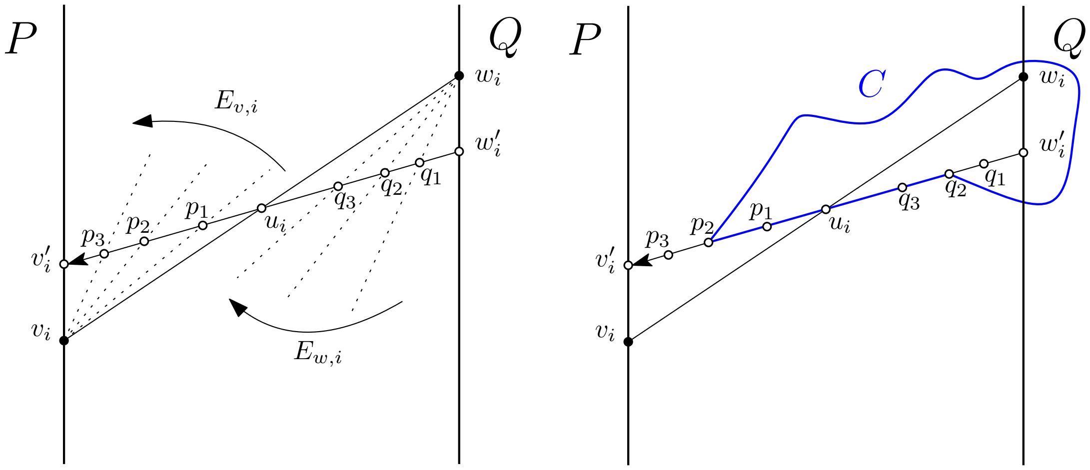

FIGURE 14. Left: The performed construction. Right: The cycle C yielding a contradiction.

appear in C twice by the choice of C. Thus, $v_i x T y$ is a path contained in C that can be replaced by a shorter path $v_i y$. This contradicts the choice of C.

- Assume $x = q_j$ and $y = q_{j'}$ for $0 \leqslant j < j' \leqslant d$. These vertices have only two outgoing edges each: the edge to w_i and the next edge of T. w_i cannot appear in C twice by the choice of C. Thus, $xTyw_i$ is a path contained in C that can be replaced by a shorter path xw_i. This contradicts the choice of C.
- Assume $x = q_j$ and $y = p_{j'}$ for $0 \leqslant j \leqslant d$ and $0 \leqslant j' \leqslant h+1$. Similarly as before, xTy is a path contained in C. Thus, the rest of C is a path connecting y to x that is disjoint from T. Let us denote it by T'.

  Let $D = (v_i, u_i, w_i)$. This path divides L into two closed regions such that y is contained in one of them and x is in the other. The path T' does not intersect D. Thus, T' leaves L and enters it again. When T' leaves or enters L, it intersects the boundary of L. If T' leaves L by intersecting Q in a vertex r above w_i and enters L by intersecting Q in a vertex g below w_i, then $rT'g$ and gQr form a cycle in the graph without T – a contradiction (see the right of Figure 14). If T' intersects P, we obtain an identical contradiction.

We proved the performed procedure keeps the desired properties of our graph intact. Now we partition edges into layers. Notice that each edge in $\overrightarrow{E}$ is contained in $E_{v,i}$ for some i (or is the edge e_i for some i). Thus, we only need to partition the edges we subdivided in order to cover both them and the edges of $\overrightarrow{E}$. Recall that the subdivided edges of P and Q are already covered by $\mathcal{I}$. Hence, we only need to consider the edges of $E_{v,i}$ and $E_{w,i}$.

Let $V_j = \{v_i : i \equiv j \pmod 3\}$ for $j = 0, 1, 2$. Let E_j be the layer containing the edges of $E'_{v,i} = E_{v,i} \cup \{e_i\}$ for all $v_i \in V_j$. We claim that E_j is a stack with respect to any topological ordering of G'. Obviously edges originating in the same vertex cannot cross in any ordering. Hence, we only need to show that the edges of $E'_{v,i}$ and $E'_{v,i'}$ are pairwise non-crossing for $i \equiv i' \pmod 3$. Let us fix v_i and $v_{i'}$ such that $i \equiv i' \pmod 3$ and $i < i'$. Fix $(v, w) = e \in E'_{v,i}$. Notice that the paths $v_i P v_{i+1} u_{i+1} w_{i+1}$ and $v_i u_i w_i Q w_{i+1}$ enclose a region R in the drawing (see Figure 15). By the definition of an st-graph, the sink t is reachable from w. Hence, the path witnessing this reachability must leave R and

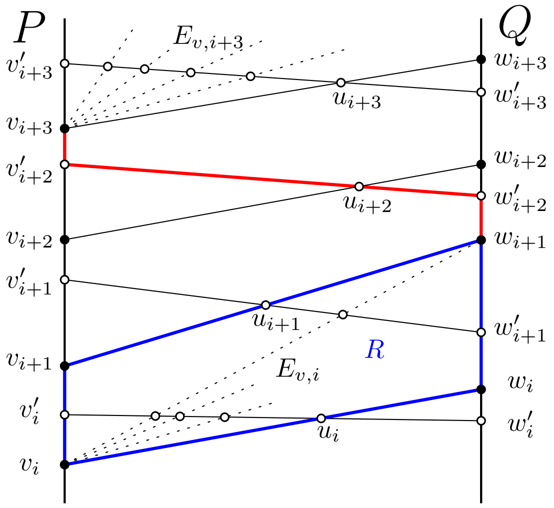

Figure 15. The region R (blue) ensuring every edge of $E_{v,i}$ is below w_{i+1} and the path (red) ensuring $w_{i+1} \prec v_{i+3}$.

thus contains a vertex of the boundary of R. This implies a reachability $w \prec_{G'} w_{i+1}$. By our construction, we have $w_{i+1} \prec_{G'} w'_{i+2} \prec_{G'} v'_{i+2} \prec_{G'} v_{i+3} \prec_{G'} v_{i'}$. This means $v, w \prec_{G'} v_{i'}$. Hence, both endpoints of e are before $v_{i'}$ in any topological ordering of G' and e does not cross any edge originating in $v_{i'}$. As e was chosen arbitrarily, this implies that the edges of $E'_{v,i}$ and $E'_{v,i'}$ are pairwise non-crossing with respect to any topological ordering of G'.

We deal with the edges of $E_{w,i}$ in a similar manner. Let $W_j = \{w_i : i \equiv j \pmod 3\}$ for $j = 0, 1, 2$. Let F_j be the layer containing the edges of $E_{w,i}$ for all $w_i \in W_j$. Let us fix w_i and $w_{i'}$ such that $i \equiv i' \pmod 3$ and $i > i'$. Fix $(v, w) = e \in E_{w,i}$. The source s is reachable from v, and thus $v_{i-1} \prec_{G'} v$. We have $w_{i'} \prec_{G'} w_{i-3} \prec_{G'} w'_{i-2} \prec_{G'} v'_{i-2} \prec_{G'} v_{i-1}$. This means $w_{i'} \prec_{G'} v, w$ and the edges of $E_{w,i}$ and $E_{w,i'}$ are pairwise non-crossing with respect to any topological ordering of G'.

The layers $E_0, E_1, E_2, F_0, F_1, F_2$ form a family $\mathcal{F}$ of 6 layers that cover $\overrightarrow{E}$ and all the edges we subdivided. Notice that during the described process we did not subdivide any edges of $\overleftarrow{E}$. Thus, we can perform a symmetric procedure for $\overleftarrow{E}$ and obtain a family $\mathcal{F}'$ of 6 layers that cover $\overleftarrow{E}$ and all the edges subdivided during the procedure. This way we covered all the edges with both endpoints on the boundary of L with a family $\mathcal{D}$ of 12 layers. Now let L' be a different lens in the lens sequence of P and Q. Performing the same procedure for L' produces an appropriate family of 12 layers $\mathcal{D}'$. By Observation 14, we can pair up the layers of $\mathcal{D}$ and $\mathcal{D}'$ and merge the paired layers into one layer while keeping the property that the edges in one layer are pairwise non-crossing with respect to any topological ordering of G'. This yields a family of 12 layers covering the edges with both endpoints either on the boundary of L or on the boundary of L'. Repeating this procedure multiple times, we obtain a family of 12 layers covering the edges with both endpoints on the boundary of some lens in the lens sequence of P and Q. Let $\mathcal{E}_1, \ldots, \mathcal{E}_{k-1}$ be families obtained this way for the consecutive paths in the left-to-right ordering of $\mathcal{P}$, that is, respectively for the paths P_1 and P_2, P_2 and P_3 and so on until P_{k-1} and P_k. By Observation 13, each inter-path edge is covered by some $\mathcal{E}_i$. Recall that each intra-path edge is covered by $\mathcal{I}$. By our construction, the edges we subdivided are covered by $\bigcup_{i=1}^{k-1} \mathcal{E}_i \cup \mathcal{I}$. This family has size $12\,(k-1)+2k \leqslant 14w$ and satisfies (2). This completes the proof. □

5. Main theorem

In this section we prove Theorem 5 by combining the approaches presented in previous sections. The proof's main idea is to use Lemma 4 to iteratively construct a sequence of supergraphs of G. Along with them we will construct a growing set of vertices of small width such that there are not too many pairwise comparable vertices in its complement.

Theorem 5. For every upward planar graph G on n vertices, the stack number of G is $\mathcal{O}(n^{2/3}\log^{2/3} n)$.

Proof. Let G be an upward planar graph on n vertices. We can add edges to G and obtain an st-graph $\overline{G}$ [17]. We have $\operatorname{sn}(G) \leqslant \operatorname{sn}(\overline{G})$ and $\overline{G}$ has the same number of vertices as G. Thus, proving $\operatorname{sn}(\overline{G}) = \mathcal{O}(n^{2/3}\log^{2/3} n)$ will yield the same bound for G and we can henceforth assume that G is an st-graph.

Using Lemma 4 we will construct a sequence of st-graphs $G_0, G_1, \ldots$, a sequence of sets of vertices $L_0, L_1, \ldots$, and a sequence of families of layers $\mathcal{E}_0, \mathcal{E}_1, \ldots$. We begin by setting $G_0 = G$, $L_0 = \emptyset$, and $\mathcal{E}_0 = \emptyset$. We shall assume G_i, L_i, $\mathcal{E}_i$ are already defined for some $i \geqslant 0$ and construct G_{i+1}, L_{i+1}, $\mathcal{E}_{i+1}$. We do it in such a way that

1. $V(G_i) \subseteq V(G_{i+1})$ and for every two vertices $u, v \in V(G_i)$ with $u \prec_{G_i} v$ we have $u \prec_{G_{i+1}} v$.
2. for $i \geqslant 1$ we have $L_i \subseteq V(G_{i-1})$ and $\mathrm{w}_{G_{i-1}}(L_i) \leqslant i$.
3. for $i \geqslant 1$, under the notation $E_\Delta(i) = E(G_i) \setminus E(G_{i-1})$, all the edges of $E(G_i[L_i]) \cup E_\Delta(i)$ belong to some layer in $\mathcal{E}_i$.
4. we have $|\mathcal{E}_i| \leqslant 14i$ and all layers of $\mathcal{E}_i$ are stacks with respect to any topological ordering of G_i.

Let $\ell = \frac{n^{2/3}}{\log^{1/3} n}$. Assume G_i, L_i, $\mathcal{E}_i$ are already defined. Assume there exists an st-path in G_i that contains at least ℓ vertices of the initial graph G that are not contained in L_i. Intuitively, this means that there exists a large set of vertices that are pairwise comparable in G_i and have not been "handled" previously. If such a path does not exist, we finish our construction. If it does, then let us denote it by P. We set $L_{i+1} = L_i \cup V(P)$. Let us verify (2). As $V(P) \subseteq V(G_i)$ and $L_i \subseteq V(G_{i-1}) \subseteq V(G_i)$ (or $L_i = \emptyset$ when $i = 0$), we have $L_{i+1} \subseteq V(G_i)$. If $i \geqslant 1$, then we know $\mathrm{w}_{G_{i-1}}(L_i) \leqslant i$. By (1), we get that every pair of vertices incomparable in G_i is incomparable in G_{i-1}. This implies $\mathrm{w}_{G_i}(L_i) \leqslant i$. We also have $\mathrm{w}_{G_i}(L_i) \leqslant i$ for $i = 0$, as $L_0 = \emptyset$. Thus, $\mathrm{w}_{G_i}(L_{i+1}) \leqslant i+1$, as the vertices of P are pairwise comparable in G_i.

We apply Lemma 4 to G_i, setting $X = L_{i+1}$. We obtain an st-graph G' and a set of layers $\mathcal{E}'$. We set $G_{i+1} = G'$ and $\mathcal{E}_{i+1} = \mathcal{E}'$. It is straightforward to verify (1), (3), and (4), as these properties are exactly the ones in the statement of Lemma 4. This ends our construction.

Let t denote the largest index i for which G_i, L_i, and $\mathcal{E}_i$ are defined, that is, there is no st-path in G_t that contains at least ℓ vertices of the initial graph G that are not contained in L_t. Note that L_i contains at least $i\ell$ vertices of G, and thus $t \leqslant \frac{n}{\ell} = n^{1/3}\log^{1/3} n$. Denote $S = V(G) \setminus L_t$. By the definition of t, we know that $\mathrm{h}_{G_t}(S) \leqslant \ell$.

Applying Lemma 3 to G_t with $X = S$, we obtain a topological ordering $<$ of G_t such that there is no $(4\ell + 1)$-twist with respect to $<$ consisting of edges with at least one endpoint in S. Applying Observation 2 to the graph consisting of these edges we obtain a family $\mathcal{D}$ of stacks (with respect to $<$) such that $|\mathcal{D}| \leqslant 14 \cdot 4\ell \log(4\ell) = \mathcal{O}(\ell \log \ell)$. Note that in order to apply Lemma 3, we needed the height of S to be bounded in G_t. This is why we needed the iterative process, as Lemma 4 gives us a supergraph in which new comparabilities might occur.

Let $\mathcal{T}' = \bigcup_{i=0}^{t} \mathcal{E}_i \cup \mathcal{D}$. We claim that any $T \in \mathcal{T}'$ is a stack with respect to $<$. Let us prove this. Fix $T \in \mathcal{T}'$. If $T \in \mathcal{D}$, then we conclude by the definition of $\mathcal{D}$. Now notice that (1) implies that every topological ordering of G_{i+1} is also a topological ordering of G_i (after we restrict it to the vertices of G_i). This means that $<$, as a topological ordering of G_t, is a topological ordering of any of the graphs $G_0, \ldots, G_t$. If $T \in \mathcal{E}_i$, then this observation, together with (4), allows us to conclude.

Now we claim that every edge of G belongs to some $T \in \mathcal{T}'$. Let us fix $e \in E(G)$. We consider three cases.

- If $e \notin E(G_t)$, then there exists $i \in \{0, \ldots, t-1\}$ such that $e \in E(G_i)$ and $e \notin E(G_{i+1})$. This means $e \in E_\Delta(i+1)$ and by (3) we have $e \in T$ for some $T \in \mathcal{E}_{i+1}$.
- If $e \in E(G_t)$ and both endpoints of e are in L_t, then by (3) we have $e \in T$ for some $T \in \mathcal{E}_t$.
- If $e \in E(G_t)$ and at least one endpoint of e is not in L_t, then at least one is in S. We have $e \in T$ for some $T \in \mathcal{D}$ by the definition of $\mathcal{D}$.

These are indeed all the cases and our claim is proved.

Now we restrict the layers of $\mathcal{T}'$ to the edges of G (that is, we remove from them all the edges that are not in $E(G)$). We also remove duplicate edges – if one edge belongs to more than one layer, we remove it from all but one layer, choosing this layer arbitrarily. Let us denote the resulting family by $\mathcal{T}$. $(<, \mathcal{T})$ is a stack embedding of G, that is, $\mathcal{T}$ is a partition of the edges of G such that the edges of any $T \in \mathcal{T}$ are pairwise non-crossing with respect to $<$. This follows immediately from the two properties of $\mathcal{T}'$ we just proved.

What remains is to bound the size of $\mathcal{T}$. Recall that $\ell = \frac{n^{2/3}}{\log^{1/3} n}$, $t \leqslant \frac{n}{\ell} = n^{1/3} \log^{1/3} n$, $|\mathcal{D}| = \mathcal{O}(\ell \log \ell)$, and $|\mathcal{E}_i| \leqslant 14i$. We have

$$|\mathcal{T}| \leqslant |\mathcal{T}'| \leqslant \sum_{i=0}^{t} |\mathcal{E}_i| + |\mathcal{D}| = \mathcal{O}\left(t^2 + \ell \log \ell\right) =$$

$$= \mathcal{O}\left(n^{2/3} \log^{2/3} n + \frac{n^{2/3}}{\log^{1/3} n} (\log n - \log \log n)\right) = \mathcal{O}\left(n^{2/3} \log^{2/3} n\right).$$

$(<, \mathcal{T})$ is a stack embedding of G witnessing $\operatorname{sn}(G) = \mathcal{O}\left(n^{2/3} \log^{2/3} n\right)$. This completes the proof. □

6. CONCLUSIONS

While the result of Theorem 5 is a cause for celebration, a significant gap remains to be bridged. It is conjectured [19] that the stack number of upward planar graphs is bounded by a constant. This conjecture is supported by the fact that constructing lower bounds for this problems proves to be notoriously hard. The current best known construction [15] yields a lower bound of 5. The main bottleneck in the proof of Theorem 5 is the repeated application of Lemma 4. Future research could aim to improve upon this approach, perhaps by establishing a more specialised variant of Lemma 4. In the context of our proof of Theorem 5, it would be particularly useful to develop a method for managing edges such that one endpoint lies within a given set of vertices and the other resides on a fixed path. Partitioning such edges into a constant number of stacks would yield an upper bound of $\widetilde{\mathcal{O}}(\sqrt{n})$. We believe more novel ideas are required to push this bound even further. One should most likely aim to leverage some global argument that considers the entire graph at once, rather than restricting one's attention to a linear number of simple substructures, like we did in the proof of Lemma 4.

References

[1] Mohammad Alzohairi and Ivan Rival. Series-parallel planar ordered sets have pagenumber two. In *International Symposium on Graph Drawing*, pages 11–24. Springer, 1996.

[2] Mohammad Alzohairi, Ivan Rival, and Alexandr Kostochka. The pagenumber of spherical lattices is unbounded. *Arab J. Math. Sci.*, 7(1):79–82, 2001.

[3] Giuseppe Di Battista, Roberto Tamassia, and Ioannis G Tollis. Area requirement and symmetry display of planar upward drawings. *Discrete & Computational Geometry*, 7(4):381–401, 1992.

[4] Frank Bernhart and Paul C Kainen. The book thickness of a graph. *Journal of Combinatorial Theory, Series B*, 27(3):320–331, 1979.

[5] Jonathan F Buss and Peter W Shor. On the pagenumber of planar graphs. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing*, pages 98–100, 1984.

[6] Fan RK Chung, Frank Thomson Leighton, and Arnold L Rosenberg. Embedding graphs in books: a layout problem with applications to vlsi design. *SIAM Journal on Algebraic Discrete Methods*, 8(1):33–58, 1987.

[7] James Davies. Improved bounds for colouring circle graphs. *Proceedings of the American Mathematical Society*, 150(12):5121–5135, 2022.

[8] Robert P. Dilworth. A decomposition theorem for partially ordered sets. *Annals of Mathematics*, 51(1):161–166, 1950.

[9] Paul Erdös and George Szekeres. A combinatorial problem in geometry. *Compositio mathematica*, 2:463–470, 1935.

[10] Fabrizio Frati, Radoslav Fulek, and Andres Ruiz-Vargas. On the page number of upward planar directed acyclic graphs. *Journal of Graph Algorithms and Applications*, 17(3):221–244, 2013.

[11] Zvi Galil, Ravi Kannan, and Endre Szemerédi. On nontrivial separators for k-page graphs and simulations by nondeterministic one-tape turing machines. In *Proceedings of the eighteenth annual ACM symposium on Theory of computing*, pages 39–49, 1986.

[12] Francesco Giordano, Giuseppe Liotta, Tamara Mchedlidze, and Antonios Symvonis. Computing upward topological book embeddings of upward planar digraphs. In *International Symposium on Algorithms and Computation*, pages 172–183. Springer, 2007.

[13] Lenwood S Heath and Sriram V Pemmaraju. Stack and queue layouts of posets. *SIAM Journal on Discrete Mathematics*, 10(4):599–625, 1997.

[14] Lenwood S. Heath, Sriram V. Pemmaraju, and Ann N. Trenk. Stack and queue layouts of directed acyclic graphs: Part i. *SIAM Journal on Computing*, 28(4):1510–1539, 1999.

[15] Paul Jungeblut, Laura Merker, and Torsten Ueckerdt. A sublinear bound on the page number of upward planar graphs. *SIAM Journal on Discrete Mathematics*, 37(4):2312–2331, 2023.

[16] Paul Jungeblut, Laura Merker, and Torsten Ueckerdt. Directed acyclic outerplanar graphs have constant stack number. *TheoretiCS*, 4, 2025.

[17] David Kelly. Fundamentals of planar ordered sets. *Discrete Mathematics*, 63(2-3):197–216, 1987.

[18] Martin Nöllenburg and Sergey Pupyrev. On families of planar dags with constant stack number. In *International Symposium on Graph Drawing and Network Visualization*, pages 135–151. Springer, 2023.

[19] Richard Nowakowski and Andrew Parker. Ordered sets, pagenumbers and planarity. *Order*, 6(3):209–218, 1989.

[20] Taylor L Ollmann. On the book thicknesses of various graphs. In *Proc. 4th Southeastern Conference on Combinatorics, Graph Theory and Computing*, volume 8. Utilitas Math., 1973.

[21] Sriram Venkata Pemmaraju. *Exploring the powers of stacks and queues via graph layouts.* Virginia Polytechnic Institute and State University, 1992.

[22] Mihalis Yannakakis. Embedding planar graphs in four pages. *Journal of Computer and System Sciences*, 38(1):36–67, 1989.

[23] Mihalis Yannakakis. Planar graphs that need four pages. *Journal of Combinatorial Theory, Series B*, 145:241–263, 2020.