The Page Number of Monotone Directed Acyclic Outerplanar Graphs Is Four or Five

Jawaherul Md. Alam
Amazon Inc., Tempe, AZ, USA

Michael A. Bekos
Department of Mathematics, University of Ioannina, Greece

Martin Gronemann
Independent Researcher, Dortmund, Germany

Michael Kaufmann
Institut für Informatik, Universität Tübingen, Germany

Abstract

A *k-page book embedding* of a directed acyclic graph consists of a topological order of its vertices and a k-coloring of its edges, such that no two edges of the same color cross, that is, their endpoints do not alternate in the order. The minimum value of k for which such an embedding exists is referred to as the *page number* of the graph. In contrast to general directed acyclic planar graphs, which may have unbounded page number [SIAM J. Comput. 28(5), 1999], it was recently shown that directed acyclic outerplanar graphs have bounded page number. In particular, Jungeblut, Merker and Ueckerdt provided an upper bound of 24,776 on their page number [FOCS 2023: 1937-1952].

In this work, we focus on so-called *monotone* directed acyclic outerplanar graphs. Starting from a single edge, these graphs are constructed by iteratively connecting a new vertex to the endpoints of an existing edge on the outer face using either two incoming or two outgoing edges incident to it. These graphs have *twist-number* 4 [GD 2023: 135-151] (i.e., they admit a topological order in which no more than four edges pairwise cross), a property, which was leveraged by Jungeblut, Merker and Ueckerdt to show that their page number is at most 128. We lower this upper bound to 5 and we also provide a lower bound of 4. A notable consequence of our result is a significant improvement of the upper bound on the page number of general directed outerplanar graphs from 24,776 to 1,160.

2012 ACM Subject Classification Mathematics of computing → Graph algorithms; Mathematics of computing → Graph theory

Keywords and phrases Book embeddings, page number, directed outerplanar graphs

Digital Object Identifier 10.4230/LIPIcs.GD.2025.9

Funding The work of M.B. has been supported by HFRI Grant No: 26320.

Acknowledgements We would like to thank the anonymous reviewers of this submission for valuable comments; in particular, for bringing to our attention the details of Remark 5.

1 Introduction

Embedding graphs in books forms a central topic in topological Graph Theory and Graph Drawing with early results dating back to the 70's [4, 19]. Primarily motivated by applications in VLSI design [7, 20, 21], book embeddings of graphs have been the subject of intensive research [1, 2, 6, 11, 12, 22, 23]. Formally, a *k-page book embedding* of a graph consists of a linear order of its vertices and a k-coloring of its edges, such that no two edges of the same color cross, that is, their endpoints do not alternate in the order; refer, e.g., to [10] for a thorough introduction. When the graph is directed, the underline linear order must coincide with a topological order of the graph [18]; see, e.g., Figure 1. Given a graph (directed or not),

33rd International Symposium on Graph Drawing and Network Visualization (GD 2025).
Editors: Vida Dujmović and Fabrizio Montecchiani; Article No. 9; pp. 9:1–9:17
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

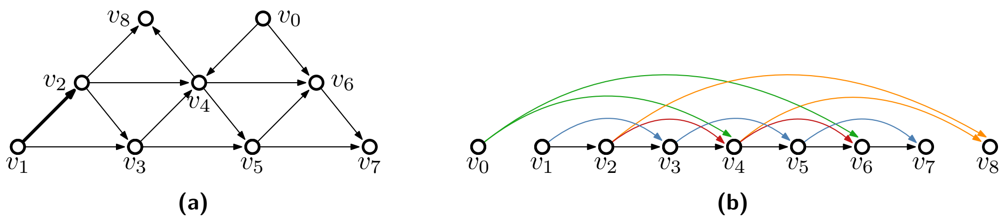

Figure 1 (a) A directed acyclic monotone outerplanar graph with base edge (v_1, v_2), and (b) a 4-page book embedding of it.

the minimum value of k for which a k-page book embedding exists is commonly referred to as the *page number* (a.k.a. *book thickness* and *stack number*) of the graph. Accordingly, the page number of a graph family is defined as the maximum page number among its members.

In this context, a particularly intriguing research branch that has received significant attention over the years involves planar graphs and their various subclasses. While the page number of planar graphs has been recently resolved in the undirected setting [2, 23, 22], several open problems persist in the directed setting despite numerous efforts; see, e.g., [15]. Undoubtedly, the most notable one, which dates back to 1989 [18], is the one of specifying the page number of upward planar graphs; a sublinear upper bound was only recently proved in the literature [14]. In general, directed acyclic planar graphs may have unbounded page number as shown by Nowakowski and Parker [18]. This negative result was recently further strengthened by Jungeblut, Merker and Ueckerdt [15], who showed that even graphs with treewidth 2 may have unbounded page number. These negative results suggest that, in the directed setting, only very restricted subclasses of planar graphs are expected to have bounded page number. Indicatively, we mention the classes of directed trees [13], two-terminal series-parallel graphs [9], N-free graphs [16], directed acyclic outerpaths [17] and planar graphs whose faces have a special structure [5].

It is worth making a separate reference to the class of directed acyclic outerplanar graphs. For this elementary class, Heath, Pemmaraju and Trenk [13] back in 1999 conjectured a constant upper bound on their page number. Surprisingly enough, their conjectured remained open for more than twenty years. An important step towards settling it was done by Nöllenburg and Pupyrev [17] in 2023, who showed that the *monotone* directed acyclic outerplanar graphs have bounded page number; these graphs are constructed starting from a single so-called *base edge* by iteratively connecting a new vertex to the endpoints of an existing edge on the outer face using either two incoming or two outgoing edges incident to it (and thus, they are by definition maximal). More precisely, an explicit upper bound of 128 can be derived by combining the work by Nöllenburg and Pupyrev [17] with a known result by Davies [8]. Based on this result, Jungeblut, Merker and Ueckerdt [15] managed to settle in the positive the conjecture by Heath, Pemmaraju and Trenk [13]. In particular, they first observed the following relationship between the page number of monotone and general directed acyclic outerplanar graphs:

▶ **Theorem 1** (Jungeblut, Merker and Ueckerdt [15]). *If s is an upper bound on the page number of directed acyclic monotone outerplanar graphs, then $8 \cdot (12 \cdot (2s + 2) + 1)$ is an upper bound on the page number of general directed acyclic outerplanar graphs.*

Theorem 1 together with the aforementioned bound of 128 by Nöllenburg and Pupyrev [17], enabled Jungeblut, Merker and Ueckerdt [15] to prove that the page number of directed acyclic outeplanar graphs is at most 24,776. Of course, they asked for improvements to this bound.

Our Contribution. In this work, we present the first such improvement. In view of the multiplicative factors of the relationship given in Theorem 1, we naturally turn our attention to the class of directed acyclic monotone outerplanar graphs, for which we lower the upper bound on their page number from 128 [8, 17] to 5; as a side result of independent interest, we also present a corresponding lower bound of 4. A notable consequence of our upper bound is that, when it is coupled with Theorem 1, it yields a significant improvement of the upper bound on the page number of general directed acyclic outerplanar graphs from 24,776 to 1,160. A summary of our findings is given in the following theorem and its corollary.

▶ **Theorem 2.** *The page number of directed acyclic monotone outerplanar graphs is* 4 *or* 5*.*

▶ **Corollary 3.** *The page number of directed acyclic outerplanar graphs is at most* 1,160*.*

2 Preliminaries

In this paper, we consider *directed graphs*, i.e., graphs in which every edge e between two vertices u and v has an orientation, either from u to v, or from v to u; in the former case, we denote e by (u, v), while in the latter case by (v, u). When neglecting the orientation of e, we may refer to it with $\{u, v\}$. An edge in a directed graph is *transitive* if the graph contains a directed path of length at least 2 from its source to its target. A directed graph is *acyclic* if it contains no directed cycle. It is well-known that every directed acyclic graph admits a topological order.

A *linear order* $\prec$ of a directed acyclic graph G is a topological order of its vertices. For any two vertices u and v of G, we write $u \prec v$ if and only if (u, v) belongs to G and u precedes v in the topological order. We say that two independent edges (u, v) and (w, z) with $u \prec v$, $w \prec z$ and $u \prec w$ *cross* in $\prec$ if and only if $u \prec w \prec v \prec z$ holds. A *page* is a set of pairwise non-crossing edges with respect to $\prec$. A *k-page book embedding* of G consists of a linear order $\prec$ of its vertices and a partition of its edges into k pages with respect to $\prec$. The *page number* of G is the minimum value of k such that G admits a k-page book embedding.

Our focus is on *directed acyclic maximal outerplanar graphs*, which admit planar drawings with all vertices being incident to the outer face, and no edge can be added such that the resulting graph is still directed acyclic outerplanar. Such a graph is *monotone* if it can be defined through the following construction sequence: **(i)** A single edge (u, v) is a monotone directed acyclic outerplanar graph; we refer to this first edge in the construction sequence as *base edge*. **(ii)** If H is a monotone directed acyclic outerplanar graph and (x, y) is an edge of H incident to its outer face, then the graph obtained from H by adding a new vertex z and either the edges (x, z) and (y, z) or the edges (z, x) and (z, y) is again a monotone directed acyclic outerplanar graph; in both cases, we say that z is *stacked* on the edge (x, y). By the next lemma, we assume that the base edge of G is incident to the outer face.

▶ **Lemma 4.** *Every directed acyclic monotone outerplanar graph G with at least three vertices contains exactly two base edges that are incident to its outer face.*

Proof. If G has exactly three vertices, then its two non-transitive edges are base edges; see Figures 2a and 2b. So, assume that G has more than three vertices. Let z be the last vertex in the construction sequence of G and let (x, y) be the edge on which z is stacked. It follows that z is incident to exactly two edges, namely, either the edges (z, x) and (z, y), or the edges (x, z) and (y, z). We consider the case in which the edges (x, z) and (y, z) belong to G; the case in which the edges (z, x) and (z, y) belong to G is symmetric. Let H be the graph obtained by removing z from G. Since z was the last vertex in the construction sequence of G, it follows that H is a directed acyclic monotone outerplanar graph. Hence, by induction, we may assume that H contains exactly two base edges that are incident to its outer face.

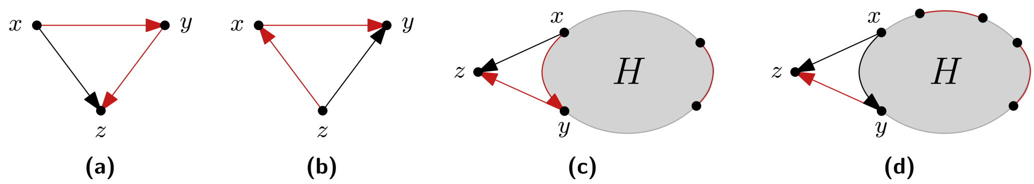

Figure 2 Illustrations for the proof of Lemma 4.

We distinguish two cases depending on whether (x, y) is a base edge of H or not. Note that, by definition of H, edge (x, y) is on the outer face of H. Assume first that (x, y) is indeed a base edge of H. Since (x, y) is a base edge of H, it follows that the edge (y, z) is a base edge of G; see Figure 2c. Since z is the last vertex in the construction sequence of G, since the edge (x, y) is an inner edge of G and since H has exactly two base edges on its outer face, it follows that G has exactly two base edges on its outer face. To complete the proof, assume that (x, y) is not a base edge of H. In this case, we prove that neither (x, z) nor (y, z) is a base of G, proving that G has exactly two base edges on its outer face, namely, the two of H. Since the edge (x, z) cannot be a base edge of G, assume for a contradiction that (y, z) is a base edge of G; see Figure 2d. It follows that there is a construction sequence that starting from (y, z) yields G. Neglecting z, this sequence implies a construction sequence for H, contradicting the fact that H has exactly two base edges on its outer face. ◀

▶ Remark 5. Note that there is a technical difference in the definition of directed acyclic monotone outerplanar graphs used in [15] and [17]. The latter requires the base edge to be on the outer face, while the former relaxes this requirement. In view of Lemma 4, the two definitions become equivalent. A notable implication of this equivalence is that substituting the twist number of 4 into the expression given by Davies [8] yields an upper bound of 64 for the page number of directed acyclic monotone outerplanar graphs (rather than 128, as used in [15]).

3 Embedding directed acyclic monotone outerplanar graphs in 5 pages

In the following, we describe how to embed a directed acyclic monotone outerplanar graph $G = (V, E)$ into five pages. Our approach consists of three main steps.

1. We define a *canonical* construction sequence for G that yields a rooted spanning tree T of the underlying undirected graph of G (see Section 3.1).
2. Based on these, we specify the linear order $\prec$ of the vertices of G, such that the edges of T do not cross in $\prec$ (see Section 3.2).
3. The remaining edges of G, namely, those that do not belong to T, can be 4-colored, such that the edges of the same color do not cross in $\prec$ (see Section 3.3).

It should be emphasized that T is rooted, but the orientation of the tree edges T in G does not reflect this hierarchy. Since T is rooted and spanning, we refer to a subtree of T rooted at a vertex v by $T(v)$ and to the set of non-tree edges by $N = E \setminus E(T)$.

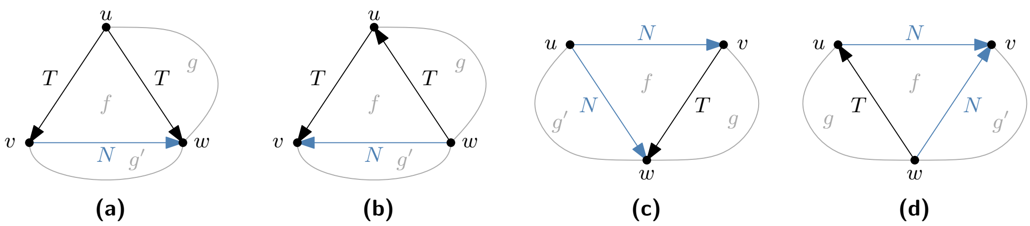

Figure 3 Illustrations for the definition of the canonical construction sequence.

3.1 A canonical construction sequence

Even though G has exactly two base edges incident to its outer face, there may exist several construction sequences that result in G, even if the starting base edge is fixed. We next describe a so-called *canonical* construction sequence that is uniquely defined if we fix the starting base edge to one of the two base edges incident to the outer face of G. To do so, we consider the weak dual G^* of G, which has a face-vertex u_f for each bounded face f of G and an edge between two vertices u_f and u_g if and only if the corresponding faces f and g of G are adjacent. G^* is a tree, which we assume to be rooted at the bounded face of G having the fixed base edge of G on its boundary. Since the fixed base edge is incident to the outer face of G, this face is uniquely defined. Further, since G is maximal, G^* is binary.

Starting from the root of G^*, we perform a specific DFS traversal of the vertices of G^*. When visiting a new vertex u_f of G^*, we assume that the edge (u,v) of G corresponding to the edge connecting u_f with its parent u_p in G^* has been assigned to T or to N. At the root of G^*, we assume that this edge is the fixed base edge which is assigned to T, such that its source is the parent of its target in T. The task is to assign the other two edges of f to T and N, and guide the traversal of G^* to each of the (at most) two subtrees of G^* rooted at u_f.

Let w be the third vertex of f. Assume first that the edge (u,v) belongs to T; see Figures 3a and 3b. Since (u,v) belongs to T, it follows that either u is the parent of v or v is the parent of u in T; recall that the edge orientations do not reflect the hierarchy in T. Consider first the case, in which f contains the edges (u,w) and (v,w). If u is the parent of v in T, then we assign the edge (u,w) to T and the edge (v,w) to N. We further assume that w is a child of u in T. On the other hand, if v is the parent of u in T, then we assign the edge (v,w) to T and the edge (u,w) to N. We further assume that w is a child of v in T. Consider now the case, in which f contains the edges (w,u) and (w,v). If u is the parent of v in T, then we assign the edge (w,u) to T and the edge (w,v) to N. We further assume that w is a child of u in T. On the other hand, if v is the parent of u in T, then we assign the edge (w,v) to T and the edge (w,u) to N. We further assume that w is a child of v in T. This completes the case where edge (u,v) belongs to T. Observe that in all four cases, we assigned to T the edges towards the vertex that is the parent among u and v.

Assume now that the edge (u,v) belongs to N; see Figures 3c and 3d. Consider first the case, in which f contains the edges (u,w) and (v,w). In this case, we assign (v,w) to T and (u,w) to N. We further assume that w is a child of v in T. Otherwise, f contains the edges (w,u) and (w,v) and we assign (w,u) to T and (w,v) to N. We further assume that w is a child of u in T. Observe that in both cases we assigned to N the transitive edge of face f.

So far, we have assigned the edges of f that are different from (u,v) to T and N. We next describe how the traversal of G^* proceeds. Let u_g and $u_{g'}$ be the (at most) two children of u_f in G^*, such that f and g share the edge assigned to T, while f and g' share the edge assigned to N in G. We recursively traverse the subtree of G^* rooted at u_g starting from u_g and then recursively traverse the subtree of G^* rooted at $u_{g'}$ starting from $u_{g'}$.

The procedure described above uniquely defines a canonical construction sequence $\pi = \{v_1, v_2, \dots, v_n\}$ of G, where (v_1, v_2) is the fixed starting base edge of G. For $i \geq 3$, let G_i be the subgraph of G induced by $v_1, \dots, v_i$. It follows that vertex v_{i+1} is connected with exactly two (neighboring) vertices in G_i, while the edges realizing these connections have been assigned to T and N. If the edge between the neighbors of v_{i+1} in G_i belongs to N, then the non-tree edge incident to v_{i+1} is transitive. The fact that T is a spanning tree of G follows by construction: each time that we define the next vertex v_{i+1} in the canonical construction sequence exactly one of the two edges connecting v_{i+1} with its two neighbors (incoming or outgoing) in G_i is assigned to T; that is, v_{i+1} is added as a leaf to T. We summarize these observations below and then present two properties of the canonical construction sequence.

T.1 The edges of T induce a spanning tree of G

T.2 The root of T is the source of the base edge of G.

T.3 Vertex v_{i+1} is a leaf in the restriction of T to G_{i+1}.

▶ **Property 1.** *Assume that the edge between the neighbors v_c and v_p of v_{i+1} in G_i belongs to T, such that v_p is the parent of v_c in T. Then $i = c$ holds, that is, v_{i+1} and v_c are consecutive in the canonical construction sequence.*

Proof. We argue along the traversal of the dual G^*. Let f be the face of G bounded by v_{i+1}, v_p and v_c. Consider the parent vertex u_g of u_f in G^*. Then, faces f and g share the edge between v_p and v_c. Since this edge belongs to T, it follows that once the visit of u_g is completed, the traversal of G^* will immediately continue to u_f. In other words, after face g the third vertex of face f is appended to the canonical construction sequence. This implies that v_c and v_{i+1} are consecutive in the canonical construction sequence. Hence, $i = c$. ◀

▶ **Property 2.** *Let (v_k, v_l) be a non-tree edge of G, that is, $(v_k, v_l) \in N$. Then, the endpoints of (v_k, v_l) belong to disjoint subtrees, that is, $T(v_k) \cap T(v_l) = \emptyset$ holds.*

Proof. We argue by induction on the length of the canonical construction sequence. The property holds in G_2 (i.e., for the base case $i = 2$), since $N = \emptyset$. Assume that for all non-tree edges of G_i with $2 < i < n$ the property holds. We prove that the property holds in G_{i+1}. Consider vertex v_{i+1}. We distinguish two cases based on whether the edge between the neighbors of v_{i+1} in G_i belongs to T or to N. In the former case, let v_p and v_c be the neighbors of v_{i+1} in G_{i+1}, such that v_p is the parent of v_c in T. Then, the edge between v_{i+1} and v_c is the non-tree edge incident to v_{i+1}. This edge connects to two siblings of v_p in T. Therefore, the endpoints of this edge belong to disjoint subtrees, namely, $T(v_{i+1}) \cap T(v_c) = \emptyset$. Consider now the case where the edge between the neighbors of v_{i+1} in G_i belongs to N. Let v_p and v_c be the neighbors of v_{i+1} in G_{i+1} and assume that v_{i+1} becomes a child of v_p in G_{i+1}. The edge between v_{i+1} and v_c is the non-tree edge incident to v_{i+1} in G_{i+1}. By induction hypothesis the endpoints of the edge between v_c and v_p are in disjoint subtrees. Since $v_{i+1} \in T(v_p)$, the same holds for the endpoints of the edge between v_{i+1} and v_c, i.e., $T(v_{i+1}) \cap T(v_c) = \emptyset$. ◀

3.2 Computing the linear order

We next describe the linear order $\prec$ of G. A key ingredient in our approach is the so-called *consecutive subtree property*, which requires for every vertex v in G the vertices in the subtree $T(v)$ to appear consecutively in $\prec$. We denote the corresponding interval of $\prec$ induced by $T(v)$ by $I_\prec(v)$. Note that by definition v is part of $I_\prec(v)$. To guarantee this property, for $2 \leq i < n$, we assume that we have recursively computed a linear order $\prec_i$ of G_i, such that:

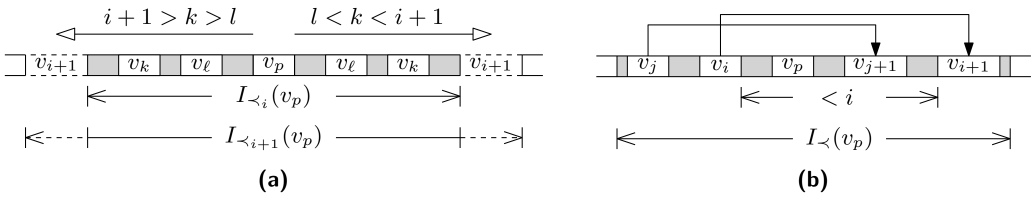

Figure 4 (a) Inserting v_{i+1} in $\prec_i$ for G_i to derive $\prec_{i+1}$ for G_{i+1}. (b) The case in which two sibling edge (v_i, v_{i+1}) and (v_j, v_{j+1}) with common parent v_p cross.

L.1 $\prec_i$ is a topological order of G_i

L.2 $\prec_i$ obeys the *consecutive subtree property* in G_i

L.3 Let v_p be the parent of two distinct children v_k and v_ℓ in T. Then, the following holds:

$$v_k \prec_i v_\ell \prec_i v_p \quad \text{or} \quad v_p \prec_i v_\ell \prec_i v_k \quad \Rightarrow \quad k > \ell.$$

In the base case $i = 2$, G_2 consists of the base edge (v_1, v_2), which is part of T; thus L.1–L.3 are trivially satisfied. For $i > 2$, we show how to insert v_{i+1} into $\prec_i$ to derive a linear order $\prec_{i+1}$ of G_{i+1} that satisfies L.1 - L.3; see Figure 4a. Assume that v_{i+1} is a child of v_p in the restriction of T in G_{i+1}. We consider two cases based on the orientation of the edge between v_{i+1} and v_p; recall that the edge orientations do not reflect the hierarchy in T. Assume first that this edge is the edge (v_p, v_{i+1}). We insert v_{i+1} directly after the last vertex of $I_{\prec_i}(v_p)$. Otherwise, we insert v_{i+1} directly before the first vertex of $I_{\prec_i}(v_p)$. This ensures that L.2 holds for G_{i+1}, since $I_{\prec_{i+1}}(v_p) = I_{\prec_i}(v_p) \cup \{v_{i+1}\}$ is consecutive. For L.3, we observe that for every child v_c of v_p in $I_{\prec_i}(v_p)$, it holds that $i+1 > c$, thereby satisfying the invariant. It remains to show that $\prec_{i+1}$ is indeed a topological order of G_{i+1}. Let v_j be the neighbor of v_{i+1} in G_i that is different from v_p. We consider two cases based on whether the edge between v_j and v_p belongs to T or to N. In the former case, assume w.l.o.g. that the edges from v_{i+1} to v_j and v_p are incoming to v_{i+1}; the case, where these edges are outgoing, is symmetric. Since we inserted v_{i+1} after $I_{\prec_i}(v_p)$ and $v_j \in I_{\prec_i}(v_p)$, we have $v_j \prec_{i+1} v_{i+1}$ and $v_p \prec_{i+1} v_{i+1}$ as desired. Assume that the edge between v_j and v_p belongs to N. Since the edge between v_{i+1} and v_p belongs to T, the edge between v_j and v_{i+1} belongs to N and thus is transitive in the face formed by v_{i+1}, v_j and v_p. Again, assume w.l.o.g. that the edges from v_{i+1} to v_j and v_p are incoming to v_{i+1}; the case, where these edges are outgoing, is symmetric. In this case, it follows that $v_p \prec_{i+1} v_{i+1}$ holds. By transitivity of (v_j, v_{i+1}), we get that (v_j, v_p) exists. Hence, $v_j \prec_{i+1} v_p \prec_{i+1} v_{i+1}$ holds and L.1–L.3 are satisfied by $\prec_{i+1}$.

We next present three key properties of the linear order $\prec$ of G. For any two vertices u and v of G, we write $\mathrm{LCA}(u, v)$ to denote the lowest common ancestor of u and v in T.

▶ **Property 3.** *For every two vertices u and v of G, $u \in I_\prec(v)$ if and only if $u \in T(v)$.*

Proof. It follows from the fact that $\prec$ obeys the consecutive subtree property (L.2). ◀

▶ **Property 4.** *Let (u, v) and (w, z) be two edges that cross in $\prec$. Then, both of the following two conditions hold:* **(i)** *u or v is in $T(LCA(w, z))$ and* **(ii)** *w or z is in $T(LCA(u, v))$.*

Proof. W.l.o.g. assume that $u \prec w \prec v \prec z$ holds. Since (w, z) belongs to G and since $w \prec v \prec z$, by L.2 we obtain $v \in I_\prec(\mathrm{LCA}(w, z))$. Symmetrically, $w \in I_\prec(\mathrm{LCA}(u, v))$. Then, the property follows from Property 3. ◀

Notice that Property 4 holds for all edges, including edges of T. However, for a pair of two independent tree edges only one of the two conditions can be satisfied, which directly implies the following property guaranteeing that all edges of T fit into a single page as desired.

▶ **Property 5.** *No two edges of T cross in $\prec$.*

3.3 Assigning edges to pages

The overall approach for the edge assignment is as follows. We dedicate a single page for edges belonging to the tree T (by Property 5). The algorithm for coloring the non-tree edges is notably simple (though the proof of correctness will be tedious). More precisely, the non-tree edge incident to the next vertex v_{i+1} ($i = 3, \ldots, n-1$) in the canonical construction sequence is colored either **(i)** with the color that is not used by its two neighbors in G_i, if v_{i+1} is stacked on a tree edge, or **(ii)** with the color of the non-tree edge it is stacked on, otherwise. In the following subsections, we will prove that this coloring yields a valid 5-page book embedding.

3.3.1 Partitioning the non-tree edges to groups

For the non-tree edges in N, we will first partition them into groups. These groups will then be 4-colored such that groups of the same color can be assigned to the same page, i.e., no two edges of the same color cross in $\prec$. In order to form the desired partition of N, we need a few more notions. We start with the notion of *sibling edge*. An edge (v_i, v_j) that belongs to N with v_i and v_j having the same parent v_p in T is called a *sibling edge* with *parent* v_p. Sibling edges satisfy the following property in the canonical construction sequence.

▶ **Property 6.** *If (v_i, v_j) is a sibling edge with parent v_p, then $i = j + 1$ or $j = i + 1$ holds.*

Proof. Follows directly from Property 1 of the canonical construction sequence. ◀

Using Property 6, we next prove that no two sibling edges with the same parent cross in $\prec$.

▶ **Property 7.** *If $\{v_i, v_{i+1}\}$ and $\{v_j, v_{j+1}\}$ are two sibling edges with common parent v_p, then they do not cross in $\prec$.*

Proof. Assume to the contrary that they cross; see Figure 4b. We distinguish cases based on the relative order of the endpoints of these edges with respect to the common parent v_p. Let us first consider the case where the endpoints of at least one of the two edges either both precede v_p or follow v_p. W.l.o.g. assume v_p follows both v_i and v_{i+1} in $\prec$. By L.3, we have $v_{i+1} \prec v_i \prec v_p$. Since both v_j and v_{j+1} are children of v_p and $j \neq i$, neither of them can be between v_i and v_{i+1} in $\prec$. Therefore, $\{v_i, v_{i+1}\}$ and $\{v_j, v_{j+1}\}$ cannot cross. We conclude that in order for $\{v_i, v_{i+1}\}$ and $\{v_j, v_{j+1}\}$ to cross, the endpoints of each of these two edges have to be on different sides of v_p in $\prec$, that is, v_p is between the endpoints of each of these edges in $\prec$. Since $\{v_i, v_{i+1}\}$ and $\{v_j, v_{j+1}\}$ cross, they are independent and therefore $i + 1 \neq j$. Assume w.l.o.g. that $i + 1 < j$. In the following, we focus on the case where $v_i \prec v_p \prec v_{i+1}$; the case in which $v_{i+1} \prec v_p \prec v_i$ is symmetric. From L.3 we get that for every child v_k of v_p with $v_i \prec v_k \prec v_{i+1}$, we have $k < i$. For $\{v_i, v_{i+1}\}$ and $\{v_j, v_{j+1}\}$ to cross either v_j or v_{j+1} has to be between v_i and v_{i+1} in $\prec$ contradicting that $i + 1 < j$. ◀

So far we have shown that tree edges can fit together in a single page and that the same holds for sibling edges having the same parent. Two questions that arise at this point are how to handle sibling edges with different parents and non-tree edges that are not sibling

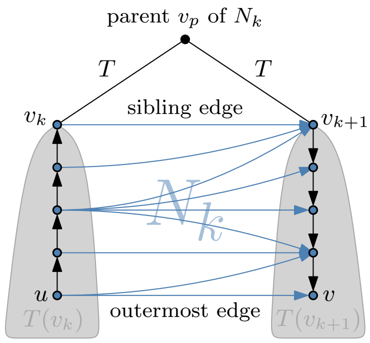

Figure 5 Illustration for Invariants N.1–N.3.

edges. For the latter, the idea is to assign these edges to *groups* such that each group has a designated sibling edge as *representative*. As we will see, a crossing between two non-tree edges of different groups can be reduced to a crossing between the groups' representatives.

Let us formally introduced the concept of groups. As already mentioned, each *group* of non-tree edges has exactly one representative, which is a sibling edge. In other words, groups and sibling edges are in a one-to-one correspondence. Further, a sibling edge belongs to the group it represents. By Property 6, we denote with $N_i \subset N$ the group that is represented by the sibling edge $\{v_i, v_{i+1}\}$. It follows that $\{v_i, v_{i+1}\} \in N_i$. Note that the index i stems from the rank of v_i and v_{i+1} in the canonical construction sequence π. Since not every pair of consecutive vertices in π yields a sibling edge, it follows that there does not exist a group N_i for every integer i in $[1, n-1]$. We further require that the groups form a partition of N.

We now describe how to obtain such a partition of N. We do so by using the canonical construction sequence and by imposing further invariant properties that will help us later to show that no more than four colors are needed in order to color the groups that we will form. With this in mind, we also need to ensure that non-tree edges of the same group can always be assigned to the same page. We capture these properties with the following invariants.

Assume that all non-tree edges of G_i with $2 < i < n$ have been partitioned to groups $N_1, \ldots, N_{i-1}$, such that the following hold for each non-empty group N_k with $k \in [1, i-1]$:

N.1 N_k has exactly one representative sibling edge $\{v_k, v_{k+1}\}$ of G_i.

N.2 For every edge (u, v) of G_i that belongs to group N_k one of the following holds:

- **a.** If (v_k, v_{k+1}) is the representative of N_k, then there is a directed path in $T(v_k)$ from u to v_k and a directed path in $T(v_{k+1})$ from v_{k+1} to v.
- **b.** If (v_{k+1}, v_k) is the representative of N_k, then there is a directed path in $T(v_{k+1})$ from u to v_{k+1} and a directed path in $T(v_k)$ from v_k to v.

Each vertex of these directed paths is incident to at least one edge of N_k.

N.3 Each non-empty group N_k is also associated with an *outermost edge* (u, v), which is the solely edge on the outer face of G_i among the edges of N_k and *covers* all edges of N_k in $\prec$, that is, for each $(u', v') \in N_k$, we have $u \preceq u' \preceq v' \preceq v$.

We prove that the aforementioned invariants hold also for the non-tree edges of G_{i+1} using the following approach. Consider vertex v_{i+1}. We distinguish cases based on whether the edge between the neighbors of v_{i+1} in G_i belongs to T or to N.

We first consider the case in which the edge between the neighbors of v_{i+1} in G_i belongs to T. Let v_p and v_c be these neighbors, such that v_p is the parent of v_c in T. In this case, vertex v_{i+1} is a child of v_p in G_{i+1}, that is, the edge between v_{i+1} and v_p belongs to T, while the edge between v_{i+1} and v_c belongs to N. In particular, the edge between v_{i+1} and v_c is a sibling edge. We assign it to a new group N_i and we make it both the representative and the outermost of this group. In other words, in this case we are creating a new group that solely consists of a single sibling edge. N.1–N.3 are satisfied because the edge between v_{i+1} and v_c cannot be associated with any of the groups N_k with $k \in [1, i-1]$, belongs to the outer face of G_{i+1} and trivially covers all edges of N_i, as there are no other edges in this group.

Consider now the case where the edge between the neighbors of v_{i+1} in G_i belongs to N. Let v_l and v_r be these neighbors. To simplify the presentation, we assume that the edge between v_l and v_r is oriented from v_l to v_r; the other case is symmetric. Since (v_l, v_r) belongs to G_i, by N.1 the edge (v_l, v_r) has been assigned to a group, say N_j with $j \in [1, i-1]$, whose representative edge is (v_j, v_{j+1}) by N.2. Assume that the edges connecting v_{i+1} with v_l and v_r in G_{i+1} are outgoing from v_{i+1}; the case in which the edges connecting v_{i+1} with v_l and v_r in G_{i+1} are incoming to v_{i+1} is symmetric. Under this assumption, the edge (v_{i+1}, v_r) is the non-tree edge incident to v_{i+1} in G_{i+1}. We proceed by assigning this edge to N_j. Regarding N.1 there is nothing to be proven, as N_j is an existing group. Since (v_j, v_{j+1}) is the representative of N_j, we know that by N.2 applied on G_i that there exist a directed path in $T(v_j)$ from v_l to v_j and a directed path in $T(v_{j+1})$ from v_{j+1} to v_r. Since v_{i+1} has an outgoing edge to v_l, it follows that there is a directed path from v_{i+1} to v_l, as well, guaranteed that N.2 is satisfied from G_{i+1}. To prove that N.3 is satisfied in G_{i+1}, we first observe that (v_l, v_r) is the outermost edge of N_j, since it is the one incident to the outer face of G_i among the edges of N_j by N.2. As a result, it covers all edges of G_i belonging N_j in $\prec$. In G_{i+1}, the edge (v_{i+1}, v_r) becomes that outermost edge of N_j, as it is the solely edge of N_j incident to the outer face of G_{i+1}. Furthermore, (v_{i+1}, v_r) covers all edges of G_{i+1} belonging N_j in $\prec$, since v_{i+1} appear right before v_l in $\prec$. This guarantees that N.3 is satisfied in G_{i+1}.

3.3.2 Properties of the formed groups

Having proved that N.1–N.3 hold for the non-tree edges of G_{i+1}, and thus by induction of G, we proceed to investigate properties of the groups that are formed with this approach.

▶ **Property 8.** *No two edges assigned to the same group cross in $\prec$.*

Proof. Assume that no two (non-tree) edges of the same group cross in $\prec$ for the subgraph G_i of G. By Invariant N.3, it follows that the non-tree edge incident to v_{i+1} in G_{i+1} covers all edges of the same group in G_i, as it becomes the outermost of this group. This implies that no two edges of the same group cross in $\prec$ for the subgraph G_{i+1} of G. This completes the proof, since in G_2 the property trivially holds. ◀

▶ **Property 9.** *Let N_i and N_j be two distinct groups with group parents v_k and v_l, respectively, such that neither v_k is an endpoint of an edge in N_j nor v_l is an endpoint of an edge in N_i. If $v_k \neq v_l$ or if the representative sibling edges of N_i and N_j are independent, then no two edges in N_i and N_j cross in $\prec$.*

Proof. Let (u, v) and (x, y) be two edges of N_i and N_j, respectively. Let also $\{v_i, v_{i+1}\}$ and $\{v_j, v_{j+1}\}$ be the representative sibling edges of N_i and N_j, respectively. We assume that the former edge is oriented from v_i to v_{i+1}, while the latter from v_j to v_{j+1}; the remaining cases are handled symmetrically. It follows by N.2 that $u \in T(v_i)$, $v \in T(v_{i+1})$, $x \in T(v_j)$ and $y \in T(v_{j+1})$. Since v_k and v_l are the group parents of N_i and N_j, these relationships imply that $u, v \in T(v_k)$ and $x, y \in T(v_l)$. Assume to the contrary that (u, v) and (x, y) cross in $\prec$. Consider first the case where the two subtrees $T(v_k)$ and $T(v_l)$ are vertex-disjoint. L.2 implies that both u and v either precede or follow both x and y in $\prec$. This implies that (u, v) and (x, y) do not cross in $\prec$; a contradiction. Hence, $T(v_k)$ and $T(v_l)$ are not vertex-disjoint.

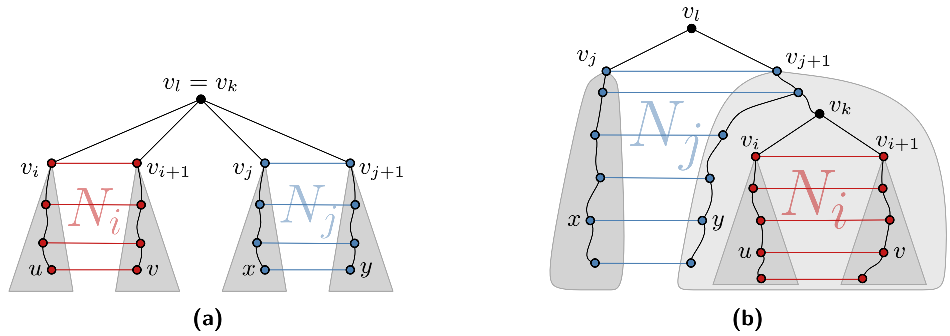

Figure 6 Illustrations for the proof of Property 9: (a) $T(v_k)$ and $T(v_l)$ share the same root; (b) $T(v_k)$ is a proper subtree of $T_(v_l)$.

Assume first that $T(v_k)$ and $T(v_l)$ share the same root, that is, $v_k = v_l$; see Figure 6a. It follows by the assumptions of the lemma that the representative sibling edges (v_i, v_{i+1}) and (v_j, v_{j+1}) of N_i and N_j are independent, implying that the subtrees $T(v_i)$, $T(v_{i+1})$, $T(v_j)$ and $T(v_{j+1})$ are vertex-disjoint. By N.2, we know that the vertices of each of these four subtrees are consecutive in $\prec$. Since (u, v) and (x, y) cross in $\prec$, it follows that either $u \prec x \prec v \prec y$ or $x \prec u \prec y \prec v$. W.l.o.g. consider the former case; the latter follows by the same argument. Since $u \in T(v_i)$, $v \in T(v_{i+1})$, $x \in T(v_j)$ and $y \in T(v_{j+1})$ and since the vertices of each of these four subtrees are consecutive in $\prec$, it follows that $v_i \prec v_j \prec v_{i+1} \prec v_{j+1}$, that is, the representative edges of N_i and N_j cross in $\prec$; a contradiction to Property 7.

Hence, we deduce that $v_k \neq v_l$. Since $T(v_k)$ and $T(v_l)$ are not vertex-disjoint and since $v_k \neq v_l$, it follows that either $T(v_k)$ is a proper subtree of $T(v_l)$ or vice versa. Assume the former, that is, $T(v_k) \subsetneq T(v_l)$; see Figure 6b. Since the vertices of $T(v_k)$ appear consecutively in $\prec$, it follows that in order for (u, v) and (x, y) to cross in $\prec$, one endpoint of (x, y), say w.l.o.g. x, has to be contained in $T(v_k)$. Then, by N.2, there exists a path between x and v_l in T, such that with the exception of v_l all vertices on this path are incident to an edge of N_j. Since x belongs also to $T(v_k)$, this path necessarily contains vertex v_k, which is a contradiction to the fact that v_k is not an endpoint of an edge in N_j. ◀

3.3.3 The groups can be 4-colored

We next color the groups with four colors, such that edges in groups with the same color do not cross in $\prec$ (recall Property 9), thus completing the proof of the upper bound of Theorem 2.

▶ **Property 10.** *There exists a 4-coloring of the groups that contain the non-tree edges of G, such that no two edges assigned to two groups of the same color cross in $\prec$.*

Proof. Our proof is by induction on the length of canonical construction sequence. Assume that we have computed a 4-coloring of the non-tree edges of G_i satisfying the next invariants:

C.1 Edges of the same group in G_i are of the same color.

C.2 Each vertex in G_i is incident to at most two edges belonging to groups of different colors, that is, the non-tree edges incident to it are either single-colored or bicolored.

C.3 For every tree edge $\{v_p, v_c\}$ on the outer face of G_i, with v_c being a child of v_p, all non-tree edges incident to v_c belong to the same group, that is, they are single-colored.

C.4 If N_k and N_l are two distinct groups with group parents v_p and v_q, such that either v_p is an endpoint of an edge in N_l or v_q is an endpoint of an edge in N_k, then N_k and N_l are of different colors.

Figure 7 Illustration for two main cases of the proof of Property 10, in which the edge between neighbors of v_{i+1} in G_i is: (a) a non-tree edge (v_l, v_r) belonging to N_j, (b) the tree edge $\{v_p, v_i\}$, whereas $\{v_i, v_{i+1}\}$ is the representative sibling edge of the new group N_i.

C.5 If N_k and N_l are two distinct groups with common group parent v_p, such that their representative sibling edges $\{v_k, v_{k+1}\}$ and $\{v_l, v_{l+1}\}$ share an endpoint, then N_k and N_l are of different colors.

Clearly, for G_2 all invariants are satisfied since $N = \emptyset$. So, we assume that $i > 2$ and we will prove that we can appropriately color the edges incident to v_{i+1} in G_{i+1}, such that C.1–C.5 are satisfied by for the non-tree edges in G_{i+1}. We distinguish cases based on whether the edge between the neighbors of v_{i+1} in G_i belongs to T or to N, that is, v_{i+1} is stacked on a tree or a non-tree edge of G_i, respectively.

Stacking on a non-tree edge. Assume first that the edge, say (v_l, v_r), between the neighbors v_l and v_r of v_{i+1} in G_i belongs to some group N_j; see Figure 7a. By C.1, we can assume that the color assigned to the edges of N_j is w.l.o.g. c_1. To simplify the presentation, we assume that the two edges incident to v_{i+1} in G_{i+1} are outgoing from v_{i+1}, that is, (v_{i+1}, v_l) and (v_{i+1}, v_l) belong to G; the case of these edges being incoming is symmetric by exchanging the roles of v_l and v_r. Under this assumption, the non-tree edge incident to v_{i+1} in G_{i+1} is the edge (v_{i+1}, v_r); the edge (v_{i+1}, v_r) is the tree edge incident to v_{i+1} in G_{i+1}. Since (v_{i+1}, v_r) belongs to N_j, we color it with color c_1, so as to guarantee C.1. Since the set of colors that is used for the edges incident to v_r and v_l does not change and since there is only one non-tree edge incident to v_{i+1} in G_{i+1} (namely, the edge (v_{i+1}, v_r)), it follows that C.2 is guaranteed in G_{i+1}. Regarding C.3, we observe that the tree edge (v_{i+1}, v_r) incident to v_{i+1} in G_{i+1} is an edge on the outer face of G_{i+1}. Since for the tree edge (v_{i+1}, v_r) v_{i+1} is the child of v_l in T and since there is only one non-tree edge incident to v_{i+1} in G_{i+1} (namely, the edge (v_{i+1}, v_r)), it follows that C.3 is maintained in G_{i+1}. Since v_{i+1} is a leaf in the restriction of T to G_{i+1}, C.4 is satisfied by the inductive hypothesis. By the same argument and the fact that no new sibling edges are introduced in G_{i+1}, C.5 holds in G_{i+1}. This completes the case where the edge between the neighbors of v_{i+1} in G_i belongs N.

Stacking on a tree edge. We proceed to the more tedious case, where the edge between the neighbors of v_{i+1} in G_i belongs to T; see Figure 7b. Let v_p and v_c be the endpoints of this edge, such that v_p is the parent of v_c in T. By Property 1, $c = i$ holds. Also, by T.3, it follows that $v_c = v_i$ is a leaf in the restriction of T to G_i. In this case, the edge $\{v_i, v_{i+1}\}$ is the representative of the new group N_i in G_{i+1} solely consisting of $\{v_i, v_{i+1}\}$, while $\{v_p, v_{i+1}\}$ belongs to T. The former implies that regardless of the color that we will assign to N_i, C.1 holds in G_{i+1}. We further observe that since v_i and v_{i+1} are siblings and v_i is a leaf in the restriction of T to G_i, it remains a leaf also in the restriction of T to G_{i+1}. In particular, this implies that v_i is not the group parent of some group in G_{i+1}.

To guarantee C.4 and C.5, we have to choose the color for the representative sibling edge $\{v_i, v_{i+1}\}$ in N_i appropriately. In particular, we have to ensure that the color chosen is different from the colors incident to v_p and v_i in G_i. By C.2, the non-tree edges incident to v_p in G_i have at most two colors, say c_1 and c_2. Since $\{v_p, v_i\}$ is incident to the outer face of G_i and since v_i is a child of v_p in T, by Item C.3 the non-tree edges incident to v_i in G_i belong to the same group; let c_3 be the color assigned to this group. The color that we assign to N_i is one not belonging in $\{c_1, c_2, c_3\}$, say c_4.

Invariant C.2 is trivially satisfied at v_{i+1} because of its degree in G_{i+1}. Since by C.3 the non-tree edges incident to v_i in G_i belong to a single group, C.2 is satisfied for v_i in G_{i+1}. The same holds for v_p by the inductive hypothesis, since the edge $\{v_p, v_{i+1}\}$ belongs to T. Hence, C.2 is maintained G_{i+1}. Since the edge $\{v_p, v_{i+1}\}$ belongs to the outer face of G_{i+1} and since v_{i+1} has only one non-tree edge incident to it in G_{i+1} (i.e., the edge $\{v_i, v_{i+1}\}$), C.3 is maintained in G_{i+1}. Note that the edge (v_p, v_i) is not on the outer face of G_{i+1}.

For C.4, consider a group N_k with $k < i$ and group parent v_q and assume that either v_p is an endpoint of an edge in N_k or v_q is an endpoint of an edge in N_i. For the latter case, note that since N_i consists only of its representative sibling edge $\{v_i, v_{i+1}\}$, the only possible candidate for v_q is v_i. However, v_i is a leaf in G_{i+1} and therefore not a parent of a group. For the former case, that is, there is a non-tree edge incident to v_p (the parent of N_i) having the same color as N_i, we argue with our definition of c_4 that this is not the case. For C.5 it is sufficient to consider a possible sibling edge sharing an endpoint with $\{v_i, v_{i+1}\}$. The only possible candidate in G_{i+1} would be a sibling edge $\{v_{i-1}, v_i\}$. However, we ensure that c_4 is different from the color of any other non-tree edge incident to v_i which is sufficient to maintain C.5 in G_{i+1}. This completes the proof that C.1–C.5 are maintained in G_{i+1}.

Putting everything together. To prove that no two non-tree edges assigned to two groups of the same color cross in $\prec$, we combine all properties that we have obtained so far. Let e and e' be two such edges and assume that e and e' belong to N_i and N_j and that v_p and v_q are the group parents of N_i and N_j, respectively. By C.1 and Property 8, it follows that if $N_i = N_j$, then e and e' do not cross in $\prec$, as desired. Hence, we may assume that N_i and N_j are two distinct groups. If v_p is an endpoint of an edge in N_i or v_q is an endpoint of an edge in N_j, then N_i and N_j are of different colors by C.4; a contradiction to the fact that e and e' belong to groups of the same color. It follows that neither v_p is an endpoint of an edge in N_i nor v_q is an endpoint of an edge in N_j. In this case, if $v_p \neq v_q$, then by Property 9, e and e' do not cross in $\prec$. Therefore, we may assume that N_i and N_j share the same parent, that is, $v_p = v_q$. If the representative sibling edges of N_i and N_j share an endpoint, then by C.5, it follows that N_i and N_j are of different colors; a contradiction. Hence, the representative sibling edges of N_i and N_j are independent. Then, again Property 9 applies which states that e and e' do not cross in $\prec$. This concludes the proof of the property, which in turn coupled with Property 5 concludes the proof of the upper bound of Theorem 2. ◀

3.4 Time complexity

Given an directed acyclic graph, we prove in the following lemma that we can determine whether it is monotone outerplanar and compute a 5-page book embedding of it, if so, in linear time.

▶ **Lemma 6.** *Given a directed acyclic graph, there is a linear time algorithm to determine whether it is monotone outerplanar and to compute a* 5*-page book embedding of it, in the positive case.*

Proof. Let G be a directed acyclic n-vertex graph. We first check in $O(n)$ time whether the underlying undirected graph of G is maximal outerplanar. If so, we assume a maximal outerplanar embedding and proceed to check whether G is monotone. In view of Lemma 4, this can be trivially done in $O(n^2)$ time, since given an edge e on the boundary of G one can check whether G is monotone with e being a base edge in $O(n)$ time by performing a BFS traversal of the weak dual G^* of G starting from the bounded face having e on its boundary.

For an $O(n)$ implementation, let f be a face that corresponds to a leaf u_f of the weak dual G^* of G. If f is the solely bounded face in G, then, since G is directed acyclic, it follows that G is monotone and we can also report the two base edges on its boundary (as observed in Lemma 4). Otherwise, let f' be the bounded face neighboring f in G, which is uniquely defined, since u_f is a leaf in G^*. Denote by (x, y) the edge shared by f and f' in G, by z the third vertex of f (which is of degree 2, since u_f is a leaf in G^*), and by H the subgraph of G obtained by the removal of z from G. There exist two cases for G to be monotone; **(i)** either one of the two edges incident to z is a base edge of G, or **(ii)** none of these two edges is a base edge of G. In Case (i), the edges incident to z are the edges (x, z) and (z, y); further, the edge (x, y) must be a base edge of H. In Case (ii), the edges incident to z are either the edges (x, z) and (y, z), or the edges (z, x) and (z, y); further, the edge (x, y) cannot be a base edge of H. Now, we recursively check whether H is monotone. In the negative case, we report that G is not monotone, as well. In the positive case, we may assume that the recursion also reports the two base edges of H. With this information, we can check in constant time, whether G is monotone by just observing the orientation of the edges incident to z depending on whether the edge (x, y) is a base edge of H or not. Since each such check takes constant time and since we perform at most $n - 4$ such checks (that is, the number of bounded faces of G), it follows that in $O(n)$ time we can check whether G is monotone.

If G meets the requirement of being monotone outerplanar and with a base edge on the outer face given, we can easily construct the canonical construction sequence and the spanning tree T of G in $O(n)$ time by a DFS traversal of the weak dual G^* of G as described in Section 3.1. For the linear order $\prec$, one may use the canonical construction sequence and maintain for every vertex v_p two ordered lists of its children in T. The first list contains the children that have to precede v_p in $\prec$, while the second contains the ones that follow v_p in $\prec$ as described in Section 3.2. Once these lists have been obtained, a simple in-order traversal of T yields $\prec$. Therefore, $\prec$ can also be computed in $O(n)$ time. As a final step, the coloring of the non-tree edges (that is, their page assignment) can be trivially done in $O(n)$ time by the simple approach described in the introduction of Section 3.3. This concludes the description of the linear-time implementation. ◀

4 Lower bounds

In this section, we first demonstrate a relatively small directed acyclic monotone outerplanar graph and prove combinatorially that it does not admit a book embedding with two pages.

▶ **Theorem 7.** *There exist directed acyclic monotone outerplanar graphs that do not admit book embeddings in two pages.*

Proof. The graph given in Figure 8a consists of nine vertices $x, y, v_1, \ldots, v_7$ and it is indeed a directed acyclic monotone outerplanar graph, whose base edge is the edge (v_1, v_3). Assume for a contradiction that there exists a book embedding $\mathcal{L}$ of it with two pages p_0 and p_1. Since the graph contains the directed path $v_1 \to \ldots \to v_7$, these vertices appear in this order in $\mathcal{L}$. Furthermore, w.l.o.g., we may assume that the edges (v_a, v_{a+2}) are assigned to page p_0, if a is even, and to page p_1, if a is odd (see the blue and red edges in Figure 8b).

Due to the edges (v_4, x) and (y, v_4), vertices x and y have to appear after and before vertex v_4 in $\mathcal{L}$, respectively. In particular, vertex x cannot appear between v_4 and v_6 in $\mathcal{L}$, as otherwise the edge (v_2, x) crosses (v_1, v_3) and (v_4, v_6), which are assigned to p_0 and p_1,

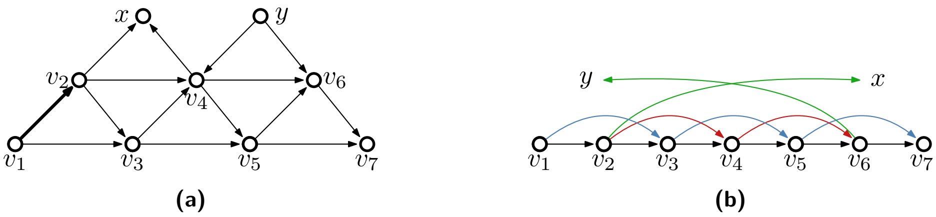

Figure 8 A directed acyclic monotone outerplanar graph not admitting a 2-page book embedding.

respectively. Symmetrically, vertex y cannot appear between v_2 and v_4. This implies that (v_2, x) and (y, v_6) cross in $\mathcal{L}$, which, in turn, implies that they have to be assigned to different pages. If (v_2, x) is assigned to p_0, then it crosses (v_1, v_3), which is also in p_0; a contradiction. If (y, v_6) is assigned to p_0, then it crosses (v_5, v_7), which is also in p_0; a contradiction. It follows that there is no 2-page book embedding for the graph, completing the proof. ◀

▶ Remark 8. We remark that there exist directed acyclic monotone outerplanar graphs that do not admit book embeddings with three pages; for an example refer to Figure 9. This statement is the outcome of the application of a well-known approach [3], which using SAT solving allows to test whether a given graph can be embedded in a book with a certain number of pages. Note that we are not aware of a directed acyclic monotone outerplanar graph that cannot be embedded in 4 pages. This motivated the title of our paper.

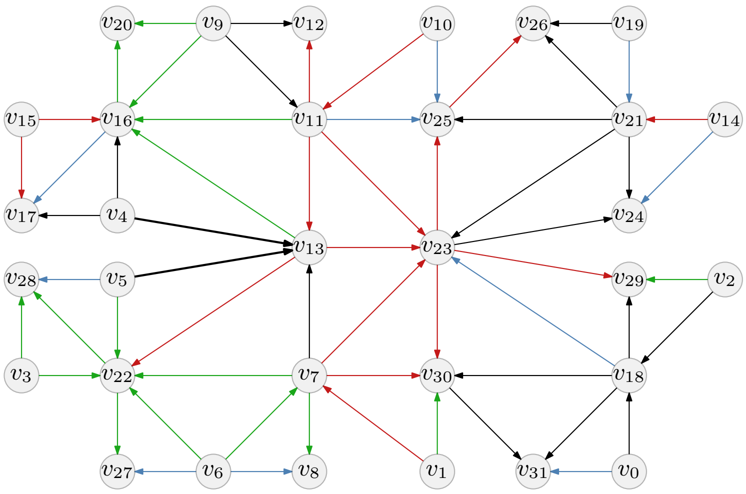

Figure 9 A directed acyclic monotone outerplanar graph that does not admit a 3-page book embedding. Note that the graph is 4-page book embeddable; the linear order of one such 4-page embedding is given by the vertex labels, while the page assignment by the coloring of the edges.

5 Open problems

The most natural open question raised from our work is to close the gap of Theorem 2. Towards another improvement on the page number of general directed acyclic outerplanar graphs, one needs either to introduce improvements for each of the remaining steps in the approach by Jungeblut, Merker and Ueckerdt [15] that supports Theorem 1 or to come up with a new, novel approach. We deem both problems intriguing.

References

1 Michael A. Bekos, Martin Gronemann, and Chrysanthi N. Raftopoulou. Two-page book embeddings of 4-planar graphs. *Algorithmica*, 75(1):158–185, 2016. doi:10.1007/s00453-015-0016-8.

2 Michael A. Bekos, Michael Kaufmann, Fabian Klute, Sergey Pupyrev, Chrysanthi N. Raftopoulou, and Torsten Ueckerdt. Four pages are indeed necessary for planar graphs. *J. Comput. Geom.*, 11(1):332–353, 2020. doi:10.20382/JOCG.V11I1A12.

3 Michael A. Bekos, Michael Kaufmann, and Christian Zielke. The book embedding problem from a SAT-solving perspective. In Emilio Di Giacomo and Anna Lubiw, editors, *Graph Drawing and Network Visualization*, volume 9411 of *LNCS*, pages 125–138. Springer, 2015. doi:10.1007/978-3-319-27261-0_11.

4 Frank Bernhart and Paul C. Kainen. The book thickness of a graph. *J. Comb. Theory, Ser. B*, 27(3):320–331, 1979. doi:10.1016/0095-8956(79)90021-2.

5 Carla Binucci, Giordano Da Lozzo, Emilio Di Giacomo, Walter Didimo, Tamara Mchedlidze, and Maurizio Patrignani. Upward book embeddability of st-graphs: Complexity and algorithms. *Algorithmica*, 85(12):3521–3571, 2023. doi:10.1007/S00453-023-01142-Y.

6 Jonathan F. Buss and Peter W. Shor. On the pagenumber of planar graphs. In Richard A. DeMillo, editor, *ACM Symposium on Theory of Computing*, pages 98–100. ACM, 1984. doi:10.1145/800057.808670.

7 Fan R. K. Chung, Frank Thomson Leighton, and Arnold L. Rosenberg. Embedding graphs in books: A layout problem with applications to VLSI design. *SIAM Journal on Algebraic and Discrete Methods*, 8(1):33–58, 1987. arXiv:http://dl.acm.org/citation.cfm?id=21936.25439.

8 James Davies. Improved bounds for colouring circle graphs. *Proceedings of the American Mathematical Society*, 2022. doi:10.1090/proc/16044.

9 Emilio Di Giacomo, Walter Didimo, Giuseppe Liotta, and Stephen K. Wismath. Book embeddability of series-parallel digraphs. *Algorithmica*, 45(4):531–547, 2006. doi:10.1007/S00453-005-1185-7.

10 Vida Dujmović and David R. Wood. On linear layouts of graphs. *Discrete Mathematics & Theoretical Computer Science*, 6(2):339–358, 2004. doi:10.46298/DMTCS.317.

11 Hikoe Enomoto, Tomoki Nakamigawa, and Katsuhiro Ota. On the pagenumber of complete bipartite graphs. *J. Comb. Theory, Ser. B*, 71(1):111–120, 1997. doi:10.1006/jctb.1997.1773.

12 Lenwood S. Heath. Embedding planar graphs in seven pages. In *FOCS*, pages 74–83. IEEE Computer Society, 1984. doi:10.1109/SFCS.1984.715903.

13 Lenwood S. Heath, Sriram V. Pemmaraju, and Ann N. Trenk. Stack and queue layouts of directed acyclic graphs: Part I. *SIAM J. Comput.*, 28(4):1510–1539, 1999. doi:10.1137/S0097539795280287.

14 Paul Jungeblut, Laura Merker, and Torsten Ueckerdt. A sublinear bound on the page number of upward planar graphs. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *ACM-SIAM SODA*, pages 963–978. SIAM, 2022. doi:10.1137/1.9781611977073.42.

15 Paul Jungeblut, Laura Merker, and Torsten Ueckerdt. Directed acyclic outerplanar graphs have constant stack number. In *FOCS*, pages 1937–1952. IEEE, 2023. doi:10.1109/FOCS57990.2023.00118.

16 Tamara Mchedlidze and Antonios Symvonis. Crossing-free acyclic hamiltonian path completion for planar *st*-digraphs. In Yingfei Dong, Ding-Zhu Du, and Oscar H. Ibarra, editors, *ISAAC*, volume 5878 of *LNCS*, pages 882–891. Springer, 2009. doi:10.1007/978-3-642-10631-6_89.

17 Martin Nöllenburg and Sergey Pupyrev. On families of planar dags with constant stack number. In Michael A. Bekos and Markus Chimani, editors, *Graph Drawing*, volume 14465 of *LNCS*, pages 135–151. Springer, 2023. doi:10.1007/978-3-031-49272-3_10.

18 Richard Nowakowski and Andrew Parker. Ordered sets, pagenumbers and planarity. *Order*, 6:209–218, 1989. doi:10.1007/BF00563521.

19 T. Ollmann. On the book thicknesses of various graphs. In F. Hoffman, R.B. Levow, and R.S.D. Thomas, editors, *Southeastern Conference on Combinatorics, Graph Theory and Computing*, volume VIII of *Congressus Numerantium*, page 459, 1973.

20 Arnold L. Rosenberg. The diogenes approach to testable fault-tolerant arrays of processors. *IEEE Trans. Computers*, 32(10):902–910, 1983. doi:10.1109/TC.1983.1676134.

21 Arnold L. Rosenberg. Book embeddings and wafer-scale integration. In *Southeastern Conference on Combinatorics, Graph Theory and Computing*, volume 54 of *Congressus Numerantium*, pages 217–224, 1986.

22 Mihalis Yannakakis. Four pages are necessary and sufficient for planar graphs (extended abstract). In Juris Hartmanis, editor, *ACM Symposium on Theory of Computing*, pages 104–108. ACM, 1986. doi:10.1145/12130.12141.

23 Mihalis Yannakakis. Planar graphs that need four pages. *J. Comb. Theory, Ser. B*, 145:241–263, 2020. doi:10.1016/j.jctb.2020.05.008.