



Upward Book Embeddability of st -Graphs: Complexity and Algorithms

Carla Binucci¹ · Giordano Da Lozzo² · Emilio Di Giacomo¹ ·
Walter Didimo¹  · Tamara Mchedlidze³ · Maurizio Patrignani²

Received: 26 October 2021 / Accepted: 9 June 2023 / Published online: 14 July 2023
© The Author(s) 2023

Abstract

A k -page upward book embedding (k UBE) of a directed acyclic graph G is a book embeddings of G on k pages with the additional requirement that the vertices appear in a topological ordering along the spine of the book. The k UBE TESTING problem, which asks whether a graph admits a k UBE, was introduced in 1999 by Heath, Pemmaraju, and Trenk (SIAM J Comput 28(4), 1999). In a companion paper, Heath and Pemmaraju (SIAM J Comput 28(5), 1999) proved that the problem is linear-time solvable for $k = 1$ and NP-complete for $k = 6$. Closing this gap has been a central question in algorithmic graph theory since then. In this paper, we make a major contribution towards a definitive answer to the above question by showing that k UBE TESTING is NP-complete for $k \geq 3$, even for st -graphs, i.e., acyclic directed graphs with a single source and a single sink. Indeed, our result, together with a recent work of Bekos et al. (Theor Comput Sci 946, 2023) that proves the NP-completeness of 2UBE for planar st -graphs, closes the question about the complexity of the k UBE problem for any k . Motivated by this hardness result, we then focus on the 2UBE TESTING for planar st -graphs. On the algorithmic side, we present an $O(f(\beta) \cdot n + n^3)$ -time algorithm for 2UBE TESTING, where β is the branchwidth of the input graph and f is a singly-exponential function on β . Since the treewidth and the branchwidth of a graph are within a constant factor from each other, this result immediately yields an FPT algorithm for st -graphs of bounded treewidth. Furthermore, we describe an $O(n)$ -time

Carla Binucci, Giordano Da Lozzo, Emilio Di Giacomo, Walter Didimo, Tamara Mchedlidze and Maurizio Patrignani have contributed equally to this work

A preliminary shorter version of this paper appeared in the Proceedings of the 35th International Symposium on Computational Geometry (SoCG) [30]. The current journal version significantly extends [30], reporting all technical details, full proofs for all the results, and discussing several additional references.

✉ Walter Didimo
walter.didimo@unipg.it

Extended author information available on the last page of the article

algorithm to test whether a plane st -graph whose faces have a special structure admits a 2UBE that additionally preserves the plane embedding of the input st -graph. On the combinatorial side, we present two notable families of plane st -graphs that always admit an embedding-preserving 2UBE.

Keywords Upward book embeddings · st -graphs · SPQR-trees · Branchwidth · Treewidth · Sphere-cut decomposition

1 Introduction

A k -page book embedding (π, σ) of an undirected graph $G = (V, E)$ consists of a vertex ordering $\pi : V \leftrightarrow \{1, 2, \dots, |V|\}$ and of an assignment $\sigma : E \rightarrow \{1, \dots, k\}$ of the edges of G to one of k sets, called *pages*, so that for any two edges (a, b) and (c, d) in the same page, with $\pi(a) < \pi(b)$ and $\pi(c) < \pi(d)$, we have neither $\pi(a) < \pi(c) < \pi(b) < \pi(d)$ nor $\pi(c) < \pi(a) < \pi(d) < \pi(b)$. From a geometric perspective, a k -page book embedding can be represented as a drawing $\Gamma(\pi, \sigma)$ of G where the k pages correspond to k half-planes sharing a vertical line, called the *spine*. Each vertex v is a point on the spine with y -coordinate $\pi(v)$; each edge e is a circular arc on page $\sigma(e)$, and no edges in the same page cross. As an example, Fig. 1a represents a 3-page book embedding of the complete graph K_6 .

For directed graphs (digraphs), Heath, Pemmaraju, and Trenk introduced a variant of k -page book embeddings in which all the edges are oriented in the upward direction [69]. This implies that the input digraph is acyclic and all its vertices appear along the spine in a topological ordering. This type of book embedding for digraphs is called a k -page upward book embedding of G (k UBE for short). An example of a 4UBE is shown in Fig. 1c (throughout the paper, in the figures showing a k UBE, the edges are oriented from bottom to top; for simplicity the orientation of the edges is not explicitly shown). Note that, if (π, σ) is a 2UBE and the two pages are co-planar, then $\Gamma(\pi, \sigma)$ is an *upward planar drawing* of G , i.e., a crossing-free drawing with all edges monotonically increasing in the upward direction. Upward planar drawings represent a classical subject in topological graph theory [7–9, 26, 27, 32–34, 36, 39, 40, 42, 44, 50, 63, 85].

The *page number* of a (di)graph G (also called *book thickness* or *stack number*) is the minimum number k such that G admits a k -page (upward) book embedding. Computing the page number of undirected and directed graphs is a widely studied problem, with applications in a variety of domains, including VLSI design, fault-tolerant processing, parallel process scheduling, sorting networks, parallel matrix computations [37, 66, 84], computational origami [2], and graph drawing [29, 46, 64, 92]. See [53, 54] for additional references and applications.

Book embeddings of undirected graphs Seminal results on book embeddings of undirected graphs are described in the paper of Bernhart and Kainen [25]. They prove that the graphs with page number one are exactly the outerplanar graphs, while graphs with page number two are the sub-Hamiltonian graphs. This second result implies that it is NP-complete to decide whether a graph admits a 2-page book embedding

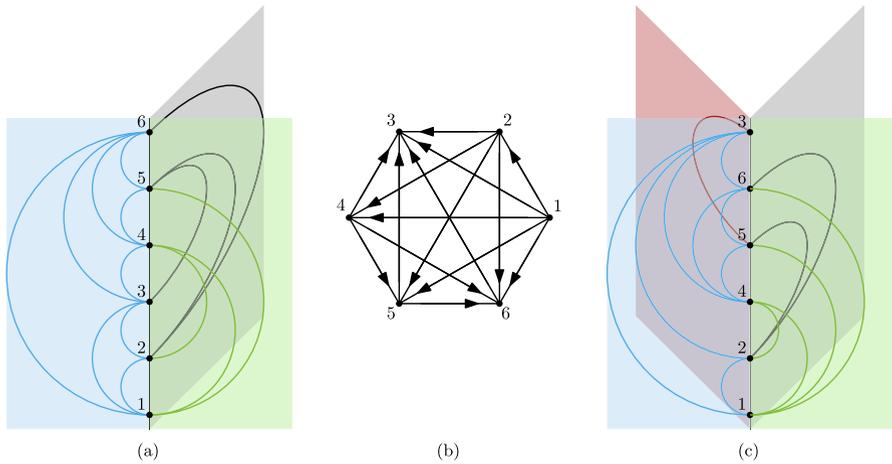


Fig. 1 **a** A 3-page book embedding of K_6 . **b** An orientation of K_6 . **c** A 4-page upward book embedding of K_6 oriented as in **(b)**

[91]. Yannakakis [93] proved that every planar graph has a 4-page book embedding, while the question whether four pages are sometimes necessary for planar graphs has only been recently settled in the affirmative [24, 94]. The aforementioned works have inspired several papers about the page number of specific families of undirected graphs (e.g., [18, 21, 23, 37, 58]) and about the relationship between the page number and other graph parameters (e.g., [55, 62, 76, 77]). Different authors studied constrained versions of k -page book embeddings where either the vertex ordering π is (totally or partially) fixed [10, 11, 38, 78, 89, 90], or the page assignment σ for the edges is given [12–14, 70], or further conditions are required for the edges that occur in the same page [3, 25]. Relaxed versions of book embeddings where edge crossings are allowed (called k -page drawings) or where edges can cross the spine (called *topological book embeddings*) have also been considered (e.g., [1, 17, 31, 35, 45, 56, 57]). Finally, 2-page (topological) book embeddings find applications to point-set embedding and universal point set (e.g., [15, 16, 48, 49, 59, 75]).

Book embeddings of directed graphs. As for undirected graphs, there are many papers devoted to the study of upper and lower bounds on the page number of directed graphs. Heath, Pemmaraju, and Trenk [69] show that directed trees and unicyclic digraphs have page number one and two, respectively. Alzohairi and Rival [5], and later Di Giacomo et al. [46] with an improved linear-time construction, show that series-parallel digraphs have page number two. Mchedlidze and Symvonis [79] generalize this result and prove that N -free upward planar digraphs, which include series-parallel digraphs, also have page number two (a digraph is upward planar if it admits an upward planar drawing). Bhore et al. [28] provide constant upper bounds on the page number of several subfamilies of upward outerplanar graphs. Frati et al. [61] give several conditions under which upward planar triangulations have bounded page number. Recently, Jungeblut, Merker, and Ueckerdt [73] proved the first sublinear upper bound on the page number of upward planar graphs. Overall, the question asked by Nowakowski

and Parker [82] more than 30 years ago of whether the page number of planar posets, and more generally of upward planar digraphs, is bounded remains unanswered. The page number of acyclic digraphs in terms of posets is also investigated by Alhashem, Jourdan, and Zaguia [4].

About lower bounds, Nowakowski and Parker [82] give an example of an upward planar digraph that requires 3 pages in any upward book embedding (see Fig. 10a). Subsequently, Hung [72] shows an upward planar digraph with page number 4. This bounds has been improved to 5 pages by Jungeblut, Merker, and Ueckerdt [73]. Heath and Pemmaraju [67] describe acyclic n -vertex planar digraphs that are not upward planar and whose upward book embeddings require $\lfloor n/2 \rfloor$ pages. Syslo [88] provides a lower bound on the page number of a poset in terms of its bump number.

Besides the study of upper and lower bounds on the page number of digraphs, several papers concentrate on the design of testing algorithms for the existence of k UBEs. Prior to this paper, the problem was known to be NP-complete only for $k = 6$ [68]. Only very recently, the case $k = 2$ has been settled by Bekos et al. [19], who show that the 2UBE problem is NP-complete for planar st -graphs and planar posets¹. Furthermore, for $k = 2$, Mchedlidze and Symvonis give linear-time testing algorithms for outerplanar and planar triangulated st -graphs [81]. They also present an $O(w^2 n^w)$ -time testing algorithm for 2UBEs of n -vertex planar st -graphs of width w , where the *width* is the minimum number of directed paths that cover all the vertices [79]. Heath and Pemmaraju [68] describe a linear-time algorithm to recognize digraphs that admit 1UBEs. As for the undirected case, constrained or relaxed variants of k UBEs for digraphs are studied [2, 47, 64], as well as applications to the point-set embedding problem [46, 64].

Contributions. Our paper is motivated by the gap present in the literature about the computational complexity of the k UBE TESTING problem. Namely, as discussed above, polynomial-time algorithms are known only for $k = 1$, or for $k = 2$ and subclasses of planar digraphs, while NP-completeness is known only for exactly 6 pages. Closing this gap has been a central question in algorithmic graph theory for over 20 years. In this paper, we make a major step towards a definitive answer to the above question, and address the research direction proposed by Heath and Pemmaraju [68] about identifying graph classes for which the existence of k UBEs can be solved efficiently. Our results are as follows:

- We prove that testing whether a digraph G admits a k UBE is NP-complete for every $k \geq 3$, even if G is an st -graph (Section 3). An analogous result was previously known only for the constrained version in which the page assignment is given [2]. Our result and the recent work of Bekos et al. [19] close the question about the complexity of the k UBE problem for any k .
- We describe a new meaningful subclass of upward planar digraphs that admit a 2UBE (Section 4). This class is structurally different from the N -free upward planar digraphs, the largest class of upward 2-page book embeddable digraphs previously known.

¹ We remark that the result of Bekos et al. [19, 20] appeared after the conference version of our research was published [30] and after the present paper was submitted.

- We give a linear-time algorithm to test the existence of an embedding-preserving 2UBE for a family of plane st -graphs whose faces have a special structure, namely each face is either a generalized triangle (i.e., it contains a transitive edge) or a rhombus (i.e., its left and right paths have length two); see Section 5.
- For planar n -vertex st -graphs of bounded branchwidth fi , we describe an FPT algorithm for the 2UBE TESTING problem that runs in $O(f(fi) \cdot n + n^3)$ time, where f is a singly-exponential function (Section 6). The algorithm works for both variable and fixed embedding, and it implies a sub-exponential-time algorithm for general planar st -graphs. Since the treewidth and the branchwidth of a graph are within a constant factor from each other, this result immediately yields an FPT algorithm for st -graphs of bounded treewidth.
- The approach used in the design of the FPT algorithm also leads to an $O(n)$ -time testing algorithm for st -graphs of treewidth at most two (also known as two-terminal series-parallel graphs) in the fixed embedding setting (Theorem 7). We recall that in the variable embedding setting an st -graphs of treewidth at most two always admits a 2UBE, which can be found in $O(n)$ time [46].

2 Preliminaries

We only consider (di)graphs without loops and multiple edges, and we denote by $V(G)$ and $E(G)$ the sets of vertices and edges of a (di)graph G . A graph G is *1-connected* if there is a path between any two vertices. G is *k -connected*, for $k \geq 2$, if the removal of $k - 1$ vertices leaves the graph 1-connected. A 2-connected (3-connected) graph is also called *biconnected* (*triconnected*).

A *planar drawing* of G is a geometric representation in the plane such that: (i) each vertex $v \in V(G)$ is drawn as a distinct point p_v ; (ii) each edge $e = (u, v) \in E(G)$ is drawn as a simple curve connecting p_u and p_v ; (iii) no two edges intersect in Γ except at their common end-vertices (if they are adjacent). A graph is *planar* if it admits a planar drawing. A planar drawing Γ of G divides the plane into topologically connected regions, called *faces*. The *outer face* of Γ is the region of unbounded size; the other faces are *internal*. A *planar embedding* of G is an equivalence class of planar drawings that define the same set of (internal and outer) faces, and it can be described by the clockwise sequence of vertices and edges on the boundary of each face plus the choice of the outer face. Graph G together with a given planar embedding is an *embedded planar graph*, or simply a *plane graph*: If Γ is a planar drawing of G whose set of faces is described by the planar embedding of G , we say that Γ *preserves* this embedding, or also that Γ is an *embedding-preserving drawing* of G .

A digraph G is a *planar st -graph* if: (i) it is acyclic; (ii) it has a single source s and a single sink t ; (iii) it admits a planar embedding \mathcal{E} with s and t on the outer face. A planar st -graph G together with \mathcal{E} is a *planar embedded st -graph* or a *plane st -graph*.

Let G be a plane st -graph and let $e = (u, v)$ be an edge of G . The *left face* (resp. *right face*) of e is the face to the left (resp. right) of e while moving from u to v . The boundary of every face f of G consists of two directed paths p_l and p_r from a common source s_f to a common sink t_f . The paths p_l and p_r are the *left path* and the

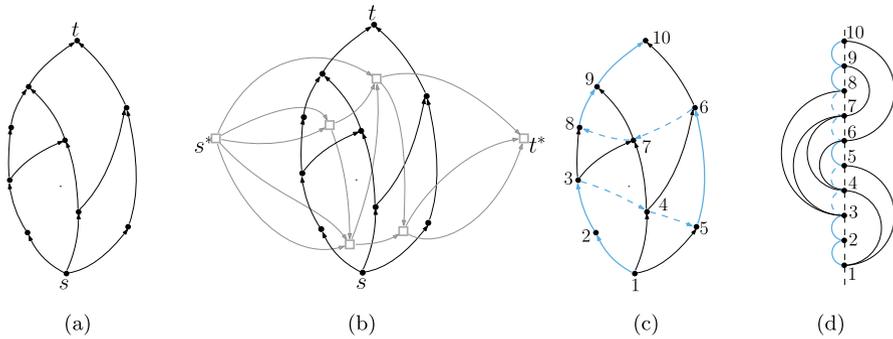


Fig. 2 **a** A plane st -graph G . **b** The dual of G is shown in gray. **c** An embedding-preserving HP-completion of G ; the light blue edges form a directed Hamiltonian st -path, and the dashed edges are augmenting edges. **d** An embedding-preserving 2UBE of G corresponding to (c)

right path of f , respectively. The vertices s_f and t_f are the source and the sink of f , respectively. If f is the outer face, p_l (resp. p_r) consists of the edges for which f is the left face (resp. right face); in this case p_l and p_r are also called the left boundary and the right boundary of G , respectively. If f is an internal face, p_l (resp. p_r) consists of the edges for which f is the right face (resp. left face).

The dual graph G^* of a plane st -graph G is a plane st -graph (possibly with multiple edges) such that: (i) G^* has a vertex associated with each internal face of G and two vertices s^* and t^* associated with the outer face of G , that are the source and the sink of G^* , respectively; (ii) for each internal edge e of G , G^* has a dual edge from the left to the right face of e ; (iii) for each edge e in the left boundary of G , there is an edge from s^* to the right face of e ; and (iv) for each edge e in the right boundary of G , there is an edge from the left face of e to t^* .

Consider a planar st -graph G and let \overline{G} be a planar st -graph obtained by augmenting G with directed edges in such a way that it contains a directed Hamiltonian st -path $P_{\overline{G}}$. The graph \overline{G} is an HP-completion of G . Consider now a plane st -graph G and let \mathcal{E} be a planar embedding of G . Let $\overline{\mathcal{E}}$ be an embedded HP-completion of G whose embedding $\overline{\mathcal{E}}$ is such that its restriction to G is \mathcal{E} . We say that $\overline{\mathcal{E}}$ is an embedding-preserving HP-completion of G . Bernhart and Kainen [25] prove that an undirected planar graph admits a 2-page book embedding if and only if it is sub-Hamiltonian, i.e., it can be made Hamiltonian by adding edges while preserving its planarity. Theorem 1 is an immediate consequence of the result in [25] for planar digraphs (see also Fig. 2); when we say that a 2UBE $\langle \pi, \sigma \rangle$ is embedding-preserving we mean that the drawing $\Gamma(\pi, \sigma)$ preserves the planar embedding of G .

Theorem 1 A planar (plane) st -graph G admits a (embedding-preserving) 2UBE $\langle \pi, \sigma \rangle$ if and only if G admits an (embedding-preserving) HP-completion \overline{G} . Also, the order π coincides with the order of the vertices along $P_{\overline{G}}$.

3 NP-Completeness for k UBE ($k \geq 3$)

In this section, we prove that k UBE TESTING, i.e., the problem of deciding whether a digraph G admits an upward k -page book embedding, is NP-complete for each fixed $k \geq 3$. The proof uses a reduction from the BETWEENNESS problem [83], which is defined as follows:

BETWEENNESS

Instance: A finite set S of elements and a set $R \subseteq S \times S \times S$ of triplets.

Question: Does there exist an ordering $\tau : S \rightarrow \mathbb{N}$ of the elements of S such that for any element $(a, b, c) \in R$ either $\tau(a) < \tau(b) < \tau(c)$ or $\tau(c) < \tau(b) < \tau(a)$?

We incrementally define a set of families of digraphs and prove some properties of these digraphs. Then, we use the digraphs of these families to reduce a generic instance of BETWEENNESS to an instance of 3UBE TESTING, thus proving the hardness result for $k = 3$. We then explain how the proof can be easily adapted to work for $k > 3$.

For a digraph G , we denote by $u \rightsquigarrow v$ a directed path from a vertex u to a vertex v in G . Let $\gamma = \langle \pi, \sigma \rangle$ be a 3UBE of G . Two edges (u, v) and (w, z) of G conflict if either $\pi(u) < \pi(w) < \pi(v) < \pi(z)$ or $\pi(w) < \pi(u) < \pi(z) < \pi(v)$. Two conflicting edges cannot be assigned to the same page. The next property will be used in the following; it is immediate from the definition of book embedding and from the pigeonhole principle.

Property 1 *In a 3UBE there cannot exist 4 edges that mutually conflict.*

Shell digraphs. The first family that we define are the *shell digraphs*, recursively defined as follows. Digraph G_0 , depicted in Fig. 3a, consists of a directed path P with eight vertices denoted as $s_0, q_0, p_{-1}, t_{-1}, s'_0, q'_0, t'_0$, and p_0 in the order they appear along P . Besides the edges of P , the following directed edges exist in G_0 : (s_0, s'_0) , (q_0, q'_0) , (t_{-1}, p_0) . Finally, there is a vertex t_0 connected to P by means of the two directed edges (p_{-1}, t_0) and (t'_0, t_0) . Graph G_h is obtained from G_{h-1} with additional vertices and edges as shown in Fig. 3b. A new directed path of two vertices s_h and q_h is connected to G_{h-1} with the edge (q_h, s_{h-1}) ; a second path of four vertices s'_h, q'_h, t'_h , and p_h is connected to G_h with the edge (t_{h-1}, s'_h) . The following edges exist between these new vertices: (s_h, s'_h) , (q_h, q'_h) , (t_{h-1}, p_h) . Finally, there is a vertex t_h connected to the other vertices by means of the two directed edges (p_{h-1}, t_h) and (t'_h, t_h) . For any $h \geq 0$, the edges (s_h, s'_h) and (q_h, q'_h) are called the *forcing edges* of G_h ; the edges (p_{h-1}, t_h) and (t_{h-1}, p_h) are the *channel edges* of G_h ; the edge (t'_h, t_h) is the *closing edge* of G_h . The vertices and edges of $G_h \setminus G_{h-1}$ are the *exclusive vertices and edges* of G_h . The following lemma establishes basic properties of the shell digraphs.

Lemma 1 *Every shell digraph G_h for $h \geq 0$ admits a 3UBE. In any 3UBE $\gamma = \langle \pi, \sigma \rangle$ of G_h the following conditions hold for every $i = 0, 1, \dots, h$:*

- S1 *all vertices of G_i are between s_i and t_i in π ;*
- S2 *the channel edges of G_i are in the same page;*

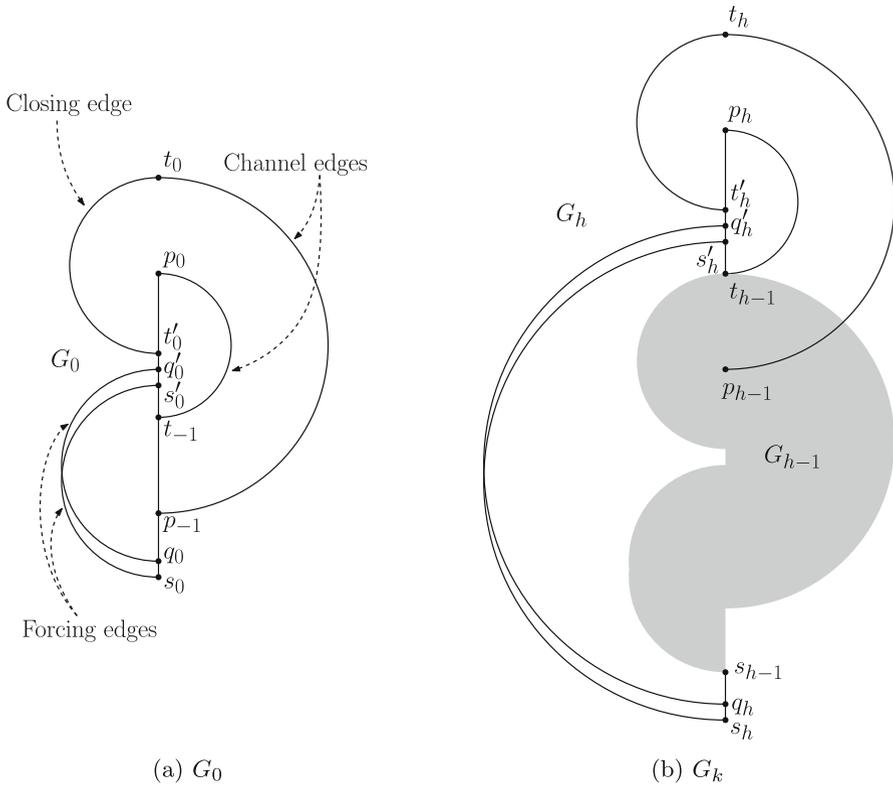


Fig. 3 Definition of shell digraphs

53 if $i > 0$, the channel edges of G_i and those of G_{i-1} are in different pages.

Proof The proof is by induction on h .

Base case $h = 0$. We describe how to define a 3UBE $\gamma = \langle \pi, \sigma \rangle$ of G_0 . The eight vertices of the directed path $s_0 \rightsquigarrow p_0$ must appear in π in the same order they appear along the path. Consider now t_0 . Because of the closing edge (t'_0, t_0) , we have $\pi(t'_0) < \pi(t_0)$. If we put t_0 between t'_0 and p_0 , the channel edges (p_{-1}, t_0) and (t_{-1}, p_0) and the forcing edges (s_0, s'_0) and (q_0, q'_0) would mutually conflict. But then a 3UBE would not exist by Property 1. Thus, the only possibility is that t_0 is the last vertex in π . This uniquely defines the order π and implies condition S1. As for the page assignment σ , the two forcing edges must be in different pages because they conflict. Since each of the two channel edges conflicts with both forcing edges, the channel edges cannot be assigned to the pages used for the forcing edges. Thus, they must be in the same page, which is possible because the two channel edges do not conflict (this proves condition S2). Finally, the closing edge conflicts with the channel edge (t_{-1}, p_0) and thus it cannot be in the same page as the channel edges; since however it does not conflict with any other edge it can be assigned to one of the pages used for the forcing edges.

This concludes the proof that a 3UBE of G_0 exists and that it must satisfy conditions S1 and S2. Condition S3 does not apply in this case.

Inductive case $h > 0$. By induction, G_{h-1} admits a 3UBE $\gamma' = \langle \pi', \sigma' \rangle$ that satisfies S1–S3. We extend γ' to a 3UBE $\gamma = \langle \pi, \sigma \rangle$ of G_h as follows. Since γ' satisfies S1, s_{h-1} is the first vertex in π' and t_{h-1} is the last one. The vertices of path $s_h \rightsquigarrow s_{h-1}$ must appear in π in the same order they appear along the path. Analogously, the vertices of $t_{h-1} \rightsquigarrow p_h$ must appear in π in the order they have along the path. Because of the closing edge (t'_h, t_h) , we have $\pi(t'_h) < \pi(t_h)$. Therefore, s_h must be the first vertex along π . Consider now t_h . If we put t_h between t'_h and p_h , the channel edges (p_{h-1}, t_h) and (t_{h-1}, p_h) and the forcing edges (s_h, s'_h) and (q_h, q'_h) would mutually conflict. But then a 3UBE would not exist by Property 1. Thus, t_h must be the last vertex in π . This uniquely defines the order π and implies condition S1 for G_h . As for the page assignment σ , observe that the only exclusive edge of G_h that conflicts with some edge of G_{h-1} is the edge (p_{h-1}, t_h) , which only conflicts with the channel edge (p_{h-2}, t_{h-1}) of G_{h-1} . This implies that (p_{h-1}, t_h) must be in a page different from the one of (p_{h-2}, t_{h-1}) . The two forcing edges of G_h must be in a page different from the channel edge (p_{h-1}, t_h) and since they conflict, they must be in different pages. The channel edge (t_{h-1}, p_h) conflicts with the forcing edges but not with the other channel edge (p_{h-1}, t_h) . Thus, the channel edges must be in the same page (which proves condition S2). The fact that the page of (p_{h-1}, t_h) must be different from that of (p_{h-2}, t_{h-1}) , implies condition S3. Finally, the closing edge conflicts with the channel edge (t_{h-1}, p_h) and thus it cannot be in the same page as the channel edges; since however it does not conflict with any other edge, it can be assigned to one of the pages used for the forcing edges. This concludes the proof that a 3UBE of G_h exists and that it satisfies conditions S1, S2, and S3. \square

Note that Condition S1 uniquely defines the vertex ordering of G_h in every 3UBE. Namely, the path $s_h \rightsquigarrow p_0$ precedes each path $t_{i-1} \rightsquigarrow p_i$ (for $i = 1, \dots, h$), and each path $t_{i-1} \rightsquigarrow p_i$ precedes the path $t_i \rightsquigarrow p_{i+1}$ (for $i = 1, \dots, h - 1$) (see Fig. 4a for an example with $h = 2$).

Filled shell digraphs. Let G_h be a shell digraph. A *filled shell digraph* $H_{h,s}$ (for $h \geq 0$ and $s \geq 1$) is obtained from G_h by adding $h + 2$ groups $\alpha_{-1}, \alpha_0, \dots, \alpha_h$ of s vertices each; see Fig. 4b for an illustration. The vertices of group α_i are denoted as $v_{i,1}, v_{i,2}, \dots, v_{i,s}$. These vertices will be used to map the elements of the set S of an instance of BETWEENNESS to an instance of 3UBE TESTING. For each vertex $v_{-1,j}$ of the set α_{-1} there is a directed edge $(p_{-1}, v_{-1,j})$ and a directed edge $(v_{-1,j}, t_{-1})$. For each vertex $v_{i,j}$ of the set α_i with $i \geq 0$ and i even, there is a directed edge $(p_i, v_{i,j})$. Finally, for each vertex $v_{i,j}$ of the set α_i with $i \geq 0$, there is a directed edge $(v_{i-1,j}, v_{i,j})$.

Lemma 2 *Every filled shell digraph $H_{h,s}$ for $s > 0$ and even $h \geq 0$ admits a 3UBE. In any 3UBE $\gamma = \langle \pi, \sigma \rangle$ of $H_{h,s}$ the following conditions hold for every $i = -1, 0, 1, \dots, h$:*

F1 the vertices of the group α_i are between p_i and t_i in π ;

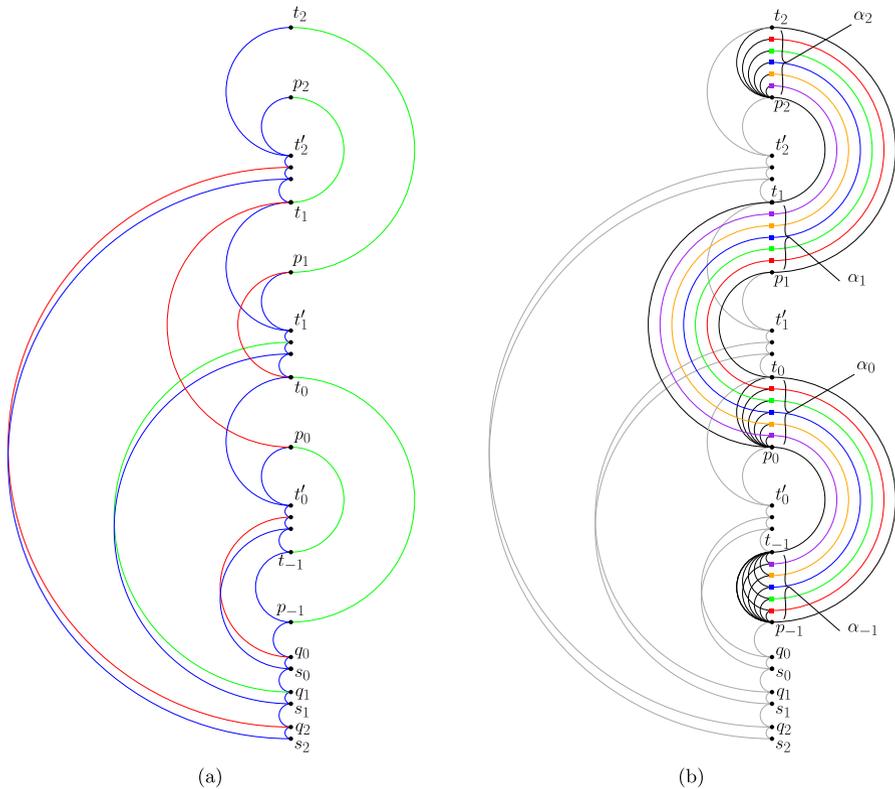


Fig. 4 **a** A 3UBE of the shell digraph G_2 ; the colors of the edges represent the pages. **b** Definition of $H_{n,s}$ for $h = 2$ and $s = 5$

- F2 if $i \geq 0$ each edge $(v_{i-1,j}, v_{i,j})$ is in the page of the channel edges of G_i (for $j = 1, \dots, s$);
- F3 if $i \geq 0$ the vertices of α_i are in reverse order with respect to those of α_{i-1} in π .

Proof The proof is by induction on h .

Base case $h = 0$. We describe how to define a 3UBE $\gamma = \langle \pi, \sigma \rangle$ of $H_{0,s}$. By Lemma 1, every 3UBE of the subgraph G_0 of $H_{0,s}$ satisfies conditions S1–S3.

Let $v_{-1,j}$ (with $1 \leq j \leq s$) be a vertex of α_{-1} . The edges $(p_{-1}, v_{-1,j})$ and $(v_{-1,j}, t_{-1})$ imply $\pi(p_{-1}) < \pi(v_{-1,j})$ and $\pi(v_{-1,j}) < \pi(t_{-1})$, which proves condition F1 for α_{-1} . Consider now the group α_0 ; the edge $(p_0, v_{0,j})$ implies $\pi(p_0) < \pi(v_{0,j})$. On the other hand, if we put $v_{0,j}$ after t_0 , the edge $(v_{-1,j}, v_{0,j})$, the channel edge (p_{-1}, t_0) , and the two forcing edges (s_0, s'_0) and (q_0, q'_0) would mutually conflict. But then a 3UBE would not exist by Property 1. Thus, each vertex of group α_0 must be between p_0 and t_0 in π , which implies condition F1 for the group α_0 . As for the page assignment σ , each $(v_{-1,j}, v_{0,j})$ conflicts with each forcing edge of G_0 , and thus it must be in the page of the channel edges of G_0 . This implies condition F2. The edges $(v_{-1,j}, v_{0,j})$ can be assigned to the same page only if the vertices of α_{-1} appear in reverse order with respect to those of α_0 in π . Thus, condition F3 holds and

a 3UBE of $H_{0,s}$ can be defined by choosing an arbitrary order for the vertices of α_{-1} and the reverse order for the vertices of α_0 .

Inductive case $h > 0$. Consider the subgraph $H'_{h,s}$ of $H_{h,s}$ consisting of $H_{h-2,s}$ plus the exclusive vertices and edges of G_{h-1} and of G_h . By induction and by Lemma 1, every 3UBE $\gamma' = \langle \pi', \sigma' \rangle$ of $H'_{h,s}$ satisfies conditions F1–F3 and conditions S1–S3. We extend γ' to a 3UBE $\gamma = \langle \pi, \sigma \rangle$ of $H_{h,s}$ as follows. By condition F1 of γ' , each vertex $v_{h-2,j}$ is before t_{h-2} in π' ; on the other hand, because of the edges $(p_h, v_{h,j})$, each vertex $v_{h,j}$ must follow p_h in π . This implies that each $v_{h-1,j}$ is between p_{h-1} and t_{h-1} in π . Indeed, if $v_{h-1,j}$ was before p_{h-1} in π , the edge $(v_{h-1,j}, v_{h,j})$, the channel edge (p_{h-1}, t_h) and the two forcing edges of G_h would mutually conflict and therefore a 3UBE would not exist by Property 1. On the other hand, if $v_{h-1,j}$ was after t_{h-1} , the edge $(v_{h-2,j}, v_{h-1,j})$, the channel edges (p_{h-2}, t_{h-1}) and the two forcing edges of G_{h-1} would mutually conflict and again a 3UBE would not exist by Property 1. Thus, each vertex of group α_{h-1} must be between p_{h-1} and t_{h-1} , which proves condition F1 for the group α_{h-1} . Consider now a vertex $v_{h,j}$. If it was after t_h in π , then the edge $(v_{h-1,j}, v_{h,j})$, the channel edge (p_{h-1}, t_h) and the two forcing edges of G_h would mutually conflict – again a 3UBE would not exist by Property 1. Hence, each vertex of α_h is between p_h and t_h , which proves condition F1 also for α_h .

As for the page assignment σ , each $(v_{h-2,j}, v_{h-1,j})$ conflicts with each forcing edge of G_{h-1} and hence it must be in the page of the channel edges of G_{h-1} . The same argument applies to the edges $(v_{h-1,j}, v_{h,j})$ with respect to the forcing edges of G_h . Thus the edges $(v_{h-1,j}, v_{h,j})$ must be in the page of the channel edges of G_h , which proves condition F2.

The edges $(v_{h-2,j}, v_{h-1,j})$ can be assigned to the same page only if the vertices of α_{h-2} appear in reverse order with respect to those of α_{h-1} in π . Similarly, the edges $(v_{h-1,j}, v_{h,j})$ can be assigned to the same page only if the vertices of α_{h-1} appear in reverse order with respect to those of α_h in π . Thus, condition F3 holds and a 3UBE of $H_{h,s}$ can be defined by ordering the vertices of α_{h-1} in reverse order with respect to those of α_{h-2} and the vertices of α_h with the same order as those of α_{h-2} . \square

Observe that, by Condition F3, all groups α_i with even index have the same ordering in π and all groups with odd index have the opposite order. As mentioned above the vertices in the groups α_i will correspond to the elements of the set S of an instance of BETWEENNESS in the reduced instance of 3UBE TESTING. If the reduced instance admits a 3UBE, the order of the odd-indexed groups in π will give the desired order for the instance of BETWEENNESS.

3-filled shell digraphs and hardness proof. Starting from a filled shell digraph $H_{h,s}$, a Λ -filled shell digraph $\widehat{H}_{h,s}$ is obtained by replacing some edges with a gadget that has two possible configurations in any 3UBE of $\widehat{H}_{h,s}$. More precisely, we replace each edge (t'_i, p_i) of $H_{h,s}$ for i odd with the gadget shown in Fig. 5a. The gadget replacing (t'_i, p_i) will be denoted as Λ_i . Notice that this replacement preserves Conditions F1–F3 of Lemma 2.

Lemma 3 *Every Λ -filled shell digraph $\widehat{H}_{h,s}$ for $s > 0$ and even $h \geq 0$ admits a 3UBE. In any 3UBE $\gamma = \langle \pi, \sigma \rangle$ of $\widehat{H}_{h,s}$ the following conditions hold for every odd $i = 1, 3, \dots, h - 1$:*

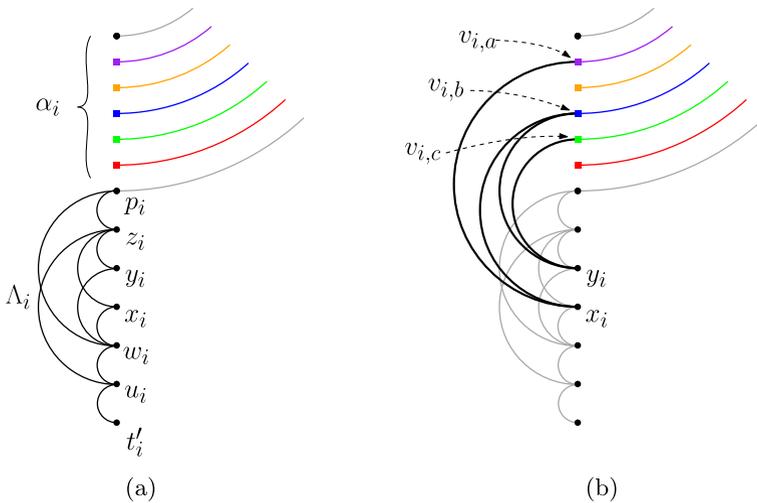


Fig. 5 **a** A gadget Λ_i (black edges). **b** The triplet edges of G_i (bold edges)

- G1 the vertices of the gadget Λ_i are between t'_i and p_i in π ;
- G2 the vertices x_i and y_i are between w_i and z_i in π and there exists a 3UBE $\gamma' = \langle \pi', \sigma' \rangle$ of $\widehat{H}_{h,s}$ where the order of x_i and y_i is exchanged in π' .

Proof By Lemma 2, every 3UBE $\gamma' = \langle \pi', \sigma' \rangle$ of $H_{h,s}$ is such that t'_i and p_i are consecutive in π' for each $i = 0, 1, \dots, h$. Notice that the gadget Λ_i admits a 3UBE (actually a 2UBE) γ_i . If we replace the edge (t'_i, p_i) with γ_i , we do not create any conflict between the edges of Λ_i and the other edges of $H_{h,s}$. This proves that $\widehat{H}_{h,s}$ has a 3UBE.

About condition G1, observe that since any vertex of the gadget Λ_i belongs to a directed path from t'_i to p_i , then the vertices of Λ_i must be between t'_i and p_i . About condition G2, observe that x_i and y_i both appear in a directed path from w_i to z_i and therefore they must be between w_i and z_i . Also, suppose that $\pi(x_i) < \pi(y_i)$ (the other case is symmetric). If we exchange the order of x_i and y_i in π' we introduce a conflict between (w_i, x_i) and (y_i, z_i) , which do not conflict with any other edges. If they are in the same page in γ it is sufficient to change the page of one of them in γ' . \square

Theorem 2 3UBE TESTING is NP-complete even for *st*-graphs.

Proof 3UBE TESTING is clearly in NP. To prove the hardness we describe a reduction from BETWEENNESS. From an instance $I = \langle S, R \rangle$ of BETWEENNESS we construct an instance G_I of 3UBE TESTING that is an *st*-graph; we start from the Λ -filled shell digraph $\widehat{H}_{h,s}$ with $h = 2\lfloor |R| \rfloor$ and $s = |S|$. Let v_1, v_2, \dots, v_s be the elements of S . They are represented in $\widehat{H}_{h,s}$ by the vertices $v_{i,1}, v_{i,2}, \dots, v_{i,s}$ of the groups α_i , for $i = -1, 0, 1, \dots, h$. In the reduction each group α_i with odd index is used to encode one triplet and, in a 3UBE of G_I , the order of the vertices in these groups (which is the same by condition F3) corresponds to the desired order of the elements of S for

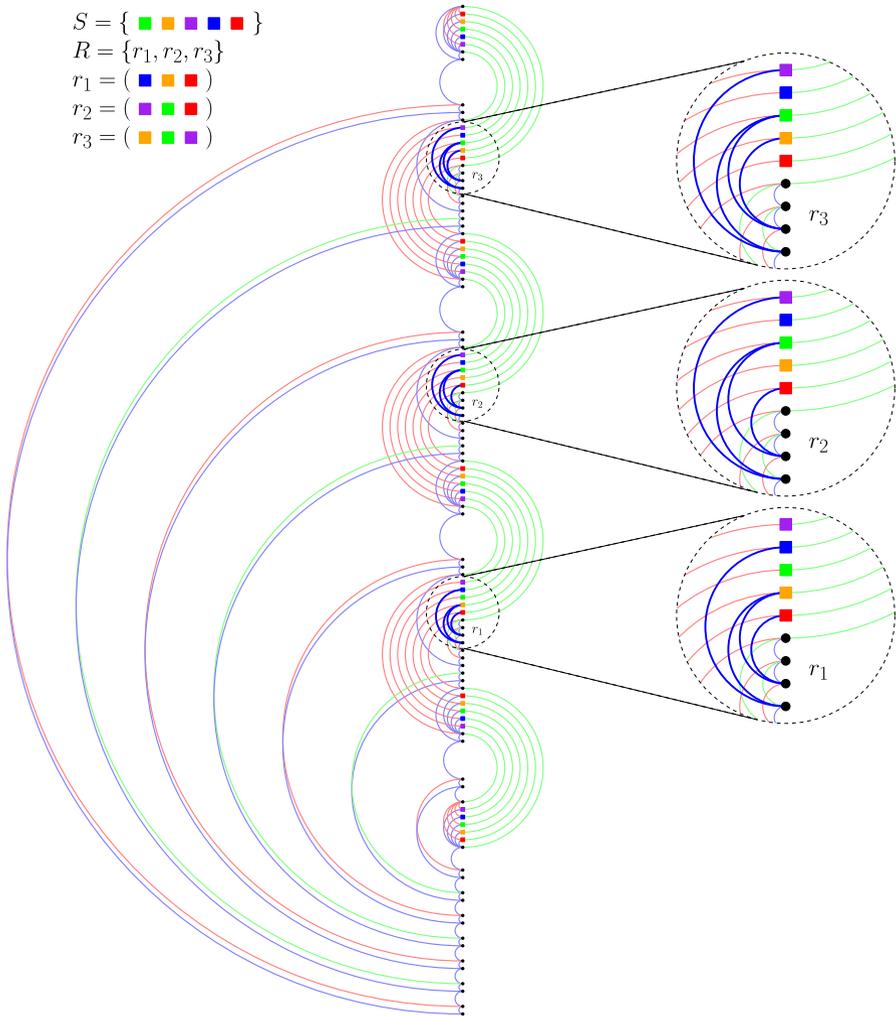


Fig. 6 A 3UBE of the st -graph G_I reduced from a positive instance $I = \langle S, R \rangle$ of BETWEENNESS; the edge colors represent the corresponding pages

the instance I . Number the triplets of R from 1 to $|R|$ and let (v_a, v_b, v_c) be the j -th triplet. We use the group α_i and the gadget Λ_i with $i = 2j - 1$ to encode the triplet (v_a, v_b, v_c) . More precisely, we add to $\hat{H}_{h,s}$ the edges $(x_i, v_{i,a}), (x_i, v_{i,b}), (y_i, v_{i,b}),$ and $(y_i, v_{i,c})$ (see Fig. 5b). These edges are called *triplet edges* and are denoted as T_i . We will show that in any 3UBE of G_I the triplet edges are forced to be in the same page and this is possible if and only if the constraints defined by the triplets in R are respected. The digraph obtained by the addition of the triplet edges is not an st -graph because the vertices of the last group α_h are all sinks. The desired instance G_I of 3UBE TESTING is the st -graph obtained by adding the edges $(v_{h,j}, t_h)$ (for

$j = 1, 2, \dots, s$). Fig. 6 shows a 3UBE of the st -graph G_I reduced from a positive instance I of BETWEENNESS.

We now show that I is a positive instance of BETWEENNESS if and only if the st -graph G_I constructed as described above admits a 3UBE. Suppose first that I is a positive instance of BETWEENNESS, i.e., there exists an ordering τ of S that satisfies all triplets in R . The subgraph $\widehat{H}_{h,s}$ of G_I admits a 3UBE that satisfies conditions S1–S3, F1–F3, and G1–G2 by Lemmas 1, 2 and 3. Observe that the order of the vertices of the groups α_i can be arbitrarily chosen (provided that all groups with even index have the same order and the groups with odd index have the reverse order). Thus we can choose the order of the groups with odd index to be equal to τ . Let $\gamma = \langle \pi, \sigma \rangle$ be the resulting 3UBE of $\widehat{H}_{h,s}$. We now show that if we add the triplet edges to γ , these edges do not conflict. Let (v_a, v_b, v_c) be the triplet encoded by the triplet edges T_i and suppose that $\tau(v_a) < \tau(v_b) < \tau(v_c)$ (the other case is symmetric). Since the vertices of the groups with odd index are ordered in π as in τ , we have $\pi(v_{i,a}) < \pi(v_{i,b}) < \pi(v_{i,c})$. If $\pi(x_i) < \pi(y_i)$ then the edges T_i do not conflict. If otherwise $\pi(y_i) < \pi(x_i)$, by condition G2 we can exchange the order of x_i and y_i , thus guaranteeing again that the triplet edges T_i do not conflict. On the other hand, the triplet edges T_i conflict with the edges E_i^Δ of Λ_i , with the channel edges E_i^{ch} of G_i , and with all the edges E_i^α connecting group α_{i-1} to group α_i . All the edges in $E_i^\Delta \cup E_i^{ch} \cup E_i^\alpha$ can be assigned to only two pages. Indeed, the edges E_i^Δ require two pages, while one page is enough for the edges of $E_i^{ch} \cup E_i^\alpha$. Also, since the edges of E_i^Δ do not conflict with those in $E_i^{ch} \cup E_i^\alpha$, two pages suffice for all of them. Hence, the triplet edges T_i can all be assigned to the third page. Since this is true for all the triplet edges, G_I admits a 3UBE.

Suppose now that G_I admits a 3UBE $\gamma = \langle \pi, \sigma \rangle$. By Lemmas 1, 2 and 3, γ satisfies conditions S1–S3, F1–F3, and G1–G2. By condition F2 the order of the vertices of the groups α_i with odd index is the same for all groups. We claim that all triplets in R are satisfied if this order is used as the order τ for the elements of S . Let (v_a, v_b, v_c) be the triplet encoded by the triplet edges T_i . By condition G2, the vertices x_i and y_i of the gadget Λ_i are between w_i and z_i in π . Thus the triplet edges T_i conflict with the edges (u_i, z_i) and (w_i, p_i) . These two edges must be in two different pages because they conflict. It follows that the triplet edges T_i must all be in the same page, i.e., the third one. Since the three edges of T_i are in the same page we have either $\pi(x_i) < \pi(y_i)$ and $\pi(v_{i,a}) < \pi(v_{i,b}) < \pi(v_{i,c})$ or $\pi(y_i) < \pi(x_i)$ and $\pi(v_{i,c}) < \pi(v_{i,b}) < \pi(v_{i,a})$ (any other order would cause a crossing between the edges of T_i). In both cases vertex $v_{i,b}$ is between $v_{i,a}$ and $v_{i,c}$, i.e., v_b is between v_a and v_c in τ . Since this is true for all triplets, I is a positive instance of BETWEENNESS. \square

For $k > 3$, the reduction from an instance I of BETWEENNESS to an instance G_I of k UBE TESTING is similar. In the shell digraph every pair of forcing edges is replaced by a bundle of $k - 1$ edges that mutually conflict (see Fig. 7a). The edges in each bundle require $k - 1$ pages and force all edges that conflict with them to use the k -th page. Analogously, the two edges (u_i, z_i) and (w_i, p_i) of the gadget Λ_i are replaced by a bundle of $k - 1$ edges that mutually conflict (see Fig. 7b); this forces the triplet edges to be in the k -th page. Finally, observe that our reduction does not work for

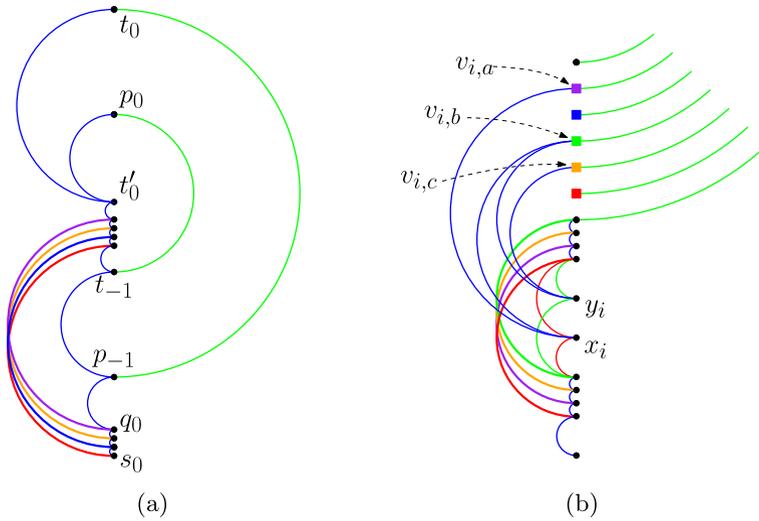


Fig. 7 Reduction for k UBE TESTING (example with $k = 5$). **a** Replacement of the forcing edges. **b** Replacement of the gadget Δ_j . In both figures colors represent the pages

$k = 2$, as it would yield an instance of 2UBE that never admits a book embedding on two pages, i.e., it is always a negative instance.

Corollary 1 k UBE TESTING is NP-complete for every $k \geq 3$, even for st -graphs.

4 Existential Results for 2UBE

Let f be an internal face of a plane st -graph, and let p_l and p_r be the left and the right path of f ; f is a *generalized triangle* if either p_l or p_r is a single edge (i.e., a transitive edge), and it is a *rhombus* if each of p_l and p_r consists of exactly two edges (see Figs. 8a and 8b).

Let G be a plane st -graph. A *forbidden configuration* of G consists of a transitive edge $e = (u, v)$ shared by two internal faces f and g such that $s_f = s_g = u$ and $t_f = t_g = v$ (i.e., two generalized triangles sharing the transitive edge); see Fig. 8c. The absence of forbidden configurations is necessary for the existence of an embedding-preserving 2UBE. If G is triangulated, the absence of forbidden configurations is also a sufficient condition [81].

Theorem 3 Any plane st -graph such that the left and the right path of every internal face contain at least two and three edges, respectively, admits an embedding-preserving 2UBE.

Proof Let G be a plane st -graph such that the left and the right path of every internal face contain at least two and three edges, respectively. We prove how to construct an embedding preserving HP-completion. The idea is to construct \overline{G} by adding a face

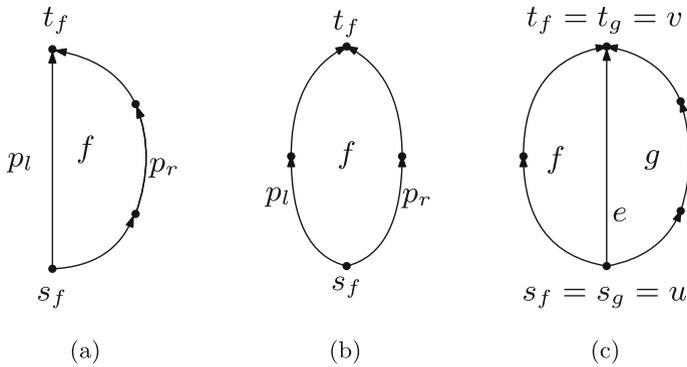


Fig. 8 **a** A generalized triangle G . **b** A rhombus **c** A forbidden configuration

of G per time from left to right. Namely, the faces of G are added according to a topological ordering of the dual graph of G . When a face f is added, its right path is attached to the right boundary of the current digraph. We maintain the invariant that at least one edge e in the left path of f belongs to the Hamiltonian path of the current digraph. The Hamiltonian path is extended by replacing e with a path that traverses the vertices of the right path of f . To this aim, dummy edges are suitably inserted inside f . When all faces are added, the resulting graph is an HP-completion \overline{G} of G .

More precisely, let N be the number of internal faces of G , and let G^* be the dual graph of G . Let $s^* = f_0, f_1, \dots, f_N, f_{N+1} = t^*$ be a topological ordering of G^* . Denote by G_0 the left boundary of G and by G_i , for $i = 1, 2, \dots, N$, the subgraph of G consisting of the faces f_1, f_2, \dots, f_i . G_i can be obtained by adding the right path p_r^i of face f_i to G_{i-1} . We construct a sequence $\overline{G}_0, \overline{G}_1, \dots, \overline{G}_N$ of st -graphs such that \overline{G}_i is an HP-completion of G_i . Clearly, \overline{G}_N will be an HP-completion of G . While constructing the sequence, we maintain the following invariant: given any two consecutive edges along the right boundary of \overline{G}_i , at least one of them belongs to the Hamiltonian path $P_{\overline{G}_i}$ of \overline{G}_i . \overline{G}_0 coincides with G_0 and all its edges are in $P_{\overline{G}_0}$, so the invariant holds. Suppose then that \overline{G}_{i-1} , with $i > 1$, satisfies the invariant. To construct \overline{G}_i we must add the right path p_r^i of f_i plus possibly some dummy edges inside f_i . Let $s_f = v_0, v_1, v_2, \dots, v_{k-1}, t_f = v_k$ be the right path p_r^i of f_i and let $s_f = u_0, u_1, u_2, \dots, u_{h-1}, t_f = u_h$ be the left path p_l^i of f_i . By hypothesis p_r^i has at least three edges, and therefore $k \geq 3$; moreover, since G has no transitive edge, $h \geq 2$. Notice that p_l^i is a subpath of the right boundary of \overline{G}_{i-1} and that the right boundary of \overline{G}_i is obtained from the right boundary of \overline{G}_{i-1} by replacing p_l^i with p_r^i . Let (u_{-1}, s_f) be the edge along the right boundary of \overline{G}_{i-1} entering s_f and let (t_f, u_{h+1}) be the edge along the right boundary of \overline{G}_{i-1} exiting t_f ; each of these edges may not exist if $s_f = s$ or $t_f = t$. We have different cases depending on whether (s_f, u_1) and (u_{h-1}, t_f) belong to $P_{\overline{G}_{i-1}}$ or not.

Case 1: neither (s_f, u_1) nor (u_{h-1}, t_f) belong to $P_{\overline{G}_{i-1}}$. See Fig. 9 for an illustration. By the invariant there is an edge (u_j, u_{j+1}) with $0 \leq j \leq h - 1$ between s_f and t_f that belongs to $P_{\overline{G}_{i-1}}$. We add the two dummy edges (u_j, v_1) and (v_{k-1}, u_{j+1}) ,

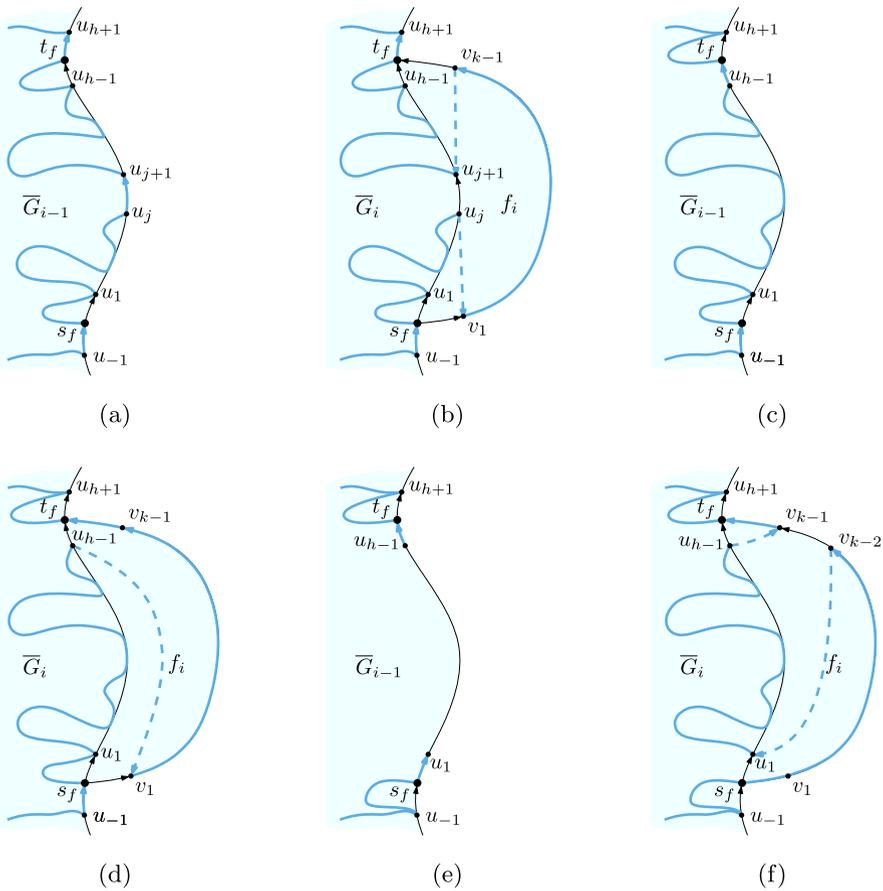


Fig. 9 Illustration for Theorem 3: a, b Case 1; c, d Case 2; e, f Case 4

thus “extending” $P_{\overline{G}_{i-1}}$ to a Hamiltonian path $P_{\overline{G}_i}$ of \overline{G}_i ; namely, the edge (u_j, u_{j+1}) is bypassed by the path $u_j, v_1, v_2, \dots, v_{k-1}, u_{j+1}$. The only edges of p_r^i that do not belong to $P_{\overline{G}_i}$ are (s_f, v_1) and (v_{k-1}, t_f) . If (u_{-1}, s_f) exists, then by the invariant it belongs to $P_{\overline{G}_{i-1}}$ (because (s_f, u_1) does not) and thus the invariant is preserved for (s_f, v_1) and its preceding and following edges. Analogously, if (t_f, u_{h+1}) exists, then by the invariant it belongs to $P_{\overline{G}_{i-1}}$ and thus the invariant is preserved for (v_{k-1}, t_f) and its preceding and following edges.

Case 2: (s_f, u_1) does not belong to $P_{\overline{G}_{i-1}}$, while (u_{h-1}, t_f) does. See Fig. 9 for an illustration. We add the dummy edge (u_{h-1}, v_1) . This “extends” $P_{\overline{G}_{i-1}}$ to $P_{\overline{G}_i}$ of \overline{G}_i bypassing the edge (u_{h-1}, t_f) with the path $u_{h-1}, v_1, v_2, \dots, v_{k-1}, t_f$. The only edge of p_r^i that does not belong to $P_{\overline{G}_i}$ is (s_f, v_1) . If (u_{-1}, s_f) exists, then by the invariant it belongs to $P_{\overline{G}_{i-1}}$ (because (s_f, u_1) does not) and thus the invariant is preserved for (s_f, v_1) and its preceding and following edges.

Case 3: (s_f, u_1) belongs to $P_{\overline{G}_{i-1}}$, while (u_{h-1}, t_f) does not. This case is symmetric to the previous one.

Case 4: both (s_f, u_1) and (u_{h-1}, t_f) belong to $P_{\overline{G}_{i-1}}$. See Fig. 9 for an illustration. We add the two dummy edges (v_{k-2}, u_1) and (u_{h-1}, v_{k-1}) . In this case we “extend” $P_{\overline{G}_{i-1}}$ to $P_{\overline{G}_i}$ bypassing (s_f, u_1) with the path s_f, v_1, \dots, v_{k-2} and bypassing the edge (u_{h-1}, t_f) with the path u_{h-1}, v_{k-1}, t_f . The only edge of p_r^i that does not belong to $P_{\overline{G}_i}$ is (v_{k-2}, v_{k-1}) ; moreover, all the edges $(s_f, v_1), (v_1, v_2), \dots, (v_{k-3}, v_{k-2})$ and the edge (v_{k-1}, t_f) belong to $P_{\overline{G}_i}$. Thus the invariant is preserved. \square

The next theorem is proved with a construction similar to that of Theorem 3.

Theorem 4 *Let G be a plane st -graph such that every internal face of G is a rhombus. Then G admits an embedding-preserving 2UBE.*

Proof During the construction of \overline{G}_i from \overline{G}_{i-1} , we maintain the same invariant as in Theorem 3. If all faces are rhombi, when we construct \overline{G}_i from \overline{G}_{i-1} , we have that p_l^i is a path s_f, u_1, t_f and p_r^i is a path s_f, v_1, t_f . At least one of (s_f, u_1) and (u_1, t_f) belongs to $P_{\overline{G}_{i-1}}$. If (s_f, u_1) belongs to $P_{\overline{G}_{i-1}}$ we add the dummy edge (v_1, u_1) . In this case we bypass the edge (s_f, u_1) with the path s_f, v_1, u_1 . If (u_1, t_f) belongs to $P_{\overline{G}_{i-1}}$ we add the dummy edge (u_1, v_1) . In this case we bypass the edge (u_1, t_f) with the path u_1, v_1, t_f . In both cases it is easy to see that the invariant is maintained. \square

5 Testing 2UBE for Plane Graphs with Special Faces

By Theorem 3, if all internal faces of a plane st -graph G are such that their left and right path contain at least two and at least three edges, respectively, G admits an embedding-preserving 2UBE. If these conditions do not hold, an embedding-preserving 2UBE may not exist (see Fig. 10a). We now describe an efficient testing algorithm for a plane st -graph $G = (V, E)$ whose internal faces are generalized triangles or rhombi (see Fig. 10b). We construct a *mixed* graph $G_M = (V, E \cup E_U)$, where E_U is a set of undirected edges and $(u, v) \in E_U$ if u and v are the two vertices of a rhombus face f distinct from s_f and t_f (blue edges in Fig. 10c). For a rhombus face f , the graph obtained from G by adding the directed edge (u, v) inside f is still a plane st -graph (see, e.g. [22, 41]). Since there is only one edge of E_U inside each rhombus face of G , the following holds.

Property 2 *Every orientation of the edges in E_U transforms G_M into an acyclic digraph.*

Theorem 5 *Let G be a plane st -graph such that every internal face of G is either a generalized triangle or a rhombus. There is an $O(n)$ -time algorithm that decides whether G admits an embedding-preserving 2UBE, and which computes it in the positive case.*

Proof The edges of E_U are the only edges that can be used to construct an embedding-preserving HP-completion of G . This, together with Theorem 1, implies that G admits

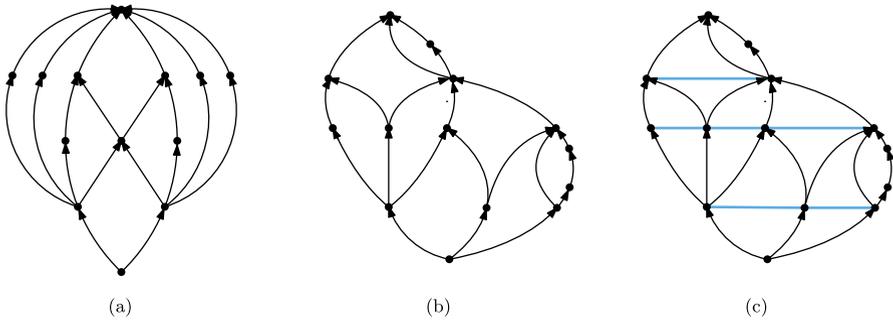


Fig. 10 **a** A plane st -graph that does not admit a 2UBE [82]. **b** A plane st -graph G whose faces are generalized triangles or rhombi. **c** The mixed graph $G_M = (V, E, E_U)$

a 2UBE if and only if the undirected edges of G_M can be oriented so that the resulting digraph $\overrightarrow{G_M}$ has a directed Hamiltonian path from s to t . By Property 2, any orientation of the undirected edges of G_M gives rise to an acyclic digraph. On the other hand an acyclic digraph is Hamiltonian if and only if it is unilateral (see, e.g. [6, Theorem 4]); we recall that a digraph is *unilateral* if each pair of vertices is connected by a directed path (in at least one of the two directions) [80]. Testing whether the undirected edges of G_M can be oriented so that the resulting digraph $\overrightarrow{G_M}$ is unilateral, and computing such an orientation if it exists, can be done in time $O(|V| + |E| + |E_U|) = O(n)$ [80, Theorem 4]. A Hamiltonian path of $\overrightarrow{G_M}$ is given by a topological ordering of its vertices. □

6 Testing Algorithms for 2UBE Parameterized by the Branchwidth

In this section, we show that the 2UBE TESTING problem is fixed-parameter tractable with respect to the branchwidth of the input st -graph both in the fixed and in the variable embedding setting. Since the treewidth $tw(G)$ and the branchwidth $bw(G)$ of a graph G are within a constant factor from each other (i.e., $bw(G) - 1 \leq tw(G) \leq \lfloor \frac{3}{2}bw(G) \rfloor - 1$ [86]), our FPT algorithm also extends to graphs of bounded treewidth.

We use the SPQR-tree data structure [43] to efficiently handle the planar embeddings of the input digraphs, and sphere-cut decompositions [87] to develop a dynamic-programming approach on the skeletons of the rigid components.

Before presenting the testing algorithm (Subsection 6.3), we briefly recall the definition of the SPQR-tree \mathcal{T} of a biconnected graph (Subsection 6.1) and the definitions of *branchwidth* and *sphere-cut decomposition* (Subsection 6.2).

6.1 SPQR-trees of Planar st -Graphs

Let G be a biconnected graph. An SPQR-tree \mathcal{T} of G is a tree-like data structure that represents the decomposition of G into its triconnected components, which can be computed in linear time [43, 65, 71]. See Fig. 11 for an illustration. Each node μ

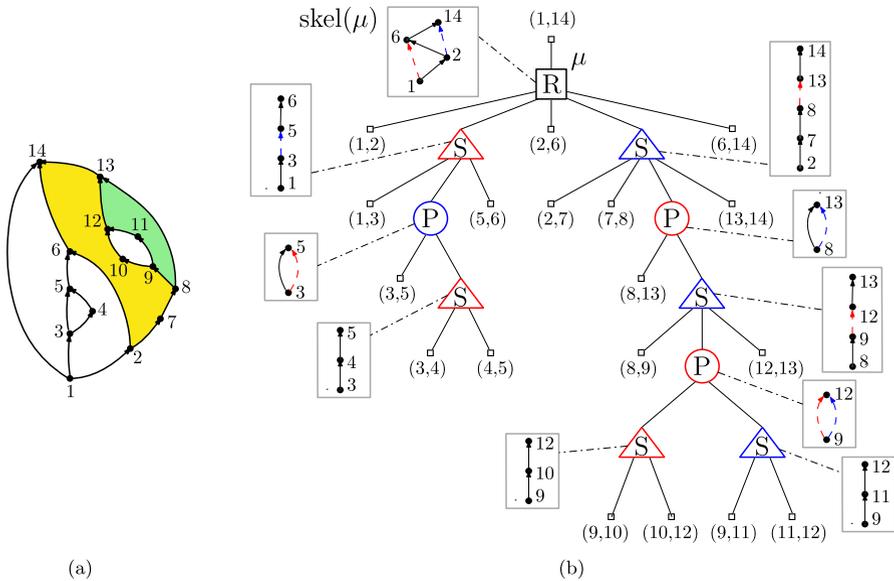


Fig. 11 **a** A biconnected planar *st*-graph G , where s and t are labeled 1 and 14, respectively. **b** The SPQR-tree of G with respect to the edge (s, t) ; for each S-, P-, or R-node of the tree, the figure depicts the skeleton of the node. In **a**, the left and the right outer face of the S-node with poles 8 and 13 are yellow and green, respectively

of \mathcal{T} of degree greater than one corresponds to a triconnected component of G (as defined in [71]) and has two special vertices, called *poles*. Namely, let μ be a node of \mathcal{T} of degree greater than one, and let u and v be its poles; μ is associated with a multigraph $\text{skel}(\mu)$, called the *skeleton* of μ . The multigraph $\text{skel}^+(\mu)$ obtained by adding an edge (u, v) to $\text{skel}(\mu)$ describes the triconnected component corresponding to μ . Node μ is either an *R*-node, an *S*-node, or a *P*-node, based on the following cases:

- R-node: If $\text{skel}^+(\mu)$ is triconnected;
- S-node: If $\text{skel}^+(\mu)$ is a cycle of length at least three;
- P-node: If $\text{skel}^+(\mu)$ is a bundle of at least three parallel edges;

Also, for each edge of G , there is a corresponding node in \mathcal{T} having degree one, which is called a *Q*-node. If μ is a Q-node that corresponds to an edge e , the *poles* of μ coincide with the end-vertices of e , and $\text{skel}(\mu)$ coincides with e .

A *virtual edge* e_v in $\text{skel}(\mu)$ corresponds to a tree node v adjacent to μ in \mathcal{T} . In the following, we will regard \mathcal{T} as rooted at an arbitrary Q-node ρ . The edge e of G corresponding to ρ is the *reference edge* of G , and \mathcal{T} is the SPQR-tree of G with respect to e . For every node μ of \mathcal{T} , the edges corresponding to the leaves (i.e., the Q-nodes) of the subtree \mathcal{T}_μ of \mathcal{T} rooted at μ induce a subgraph $\text{pert}(\mu)$ of G , called the *pertinent graph* of μ . The graph $\text{pert}(\mu)$ is described by \mathcal{T}_μ in the decomposition.

If G is planar, the SPQR-tree \mathcal{T} of G with respect to e implicitly describes all planar embeddings of G in which e is incident to the outer face. All such embeddings are obtained by combining the different planar embeddings of the skeletons of P- and

R-nodes: For a P-node μ , the different embeddings of $\text{skel}(\mu)$ are the different permutations of its edges. If μ is an R-node, $\text{skel}(\mu)$ has two possible planar embeddings, obtained by flipping $\text{skel}(\mu)$ at its poles. Let μ be a node of \mathcal{T} , let \mathcal{E} be an embedding of G , let \mathcal{E}_μ be the embedding $\text{pert}(\mu)$ in \mathcal{E} , and let f_μ^O be the outer face of \mathcal{E}_μ . The path along f_μ^O between s_μ and t_μ that leaves f_μ^O to its left (resp. to its right) when traversing the boundary of f_μ^O from s_μ to t_μ is the *left outer path* (resp. the *right outer path*) of \mathcal{E}_μ . The *left* (resp., *right*) *outer face* of \mathcal{E}_μ is the face of \mathcal{E} , not belonging to \mathcal{E}_μ , that is incident to the left (resp., to the right) outer path of \mathcal{E}_μ . In Fig. 11a, the left and the right outer face of the S-node with poles 8 and 13 are yellow and green, respectively.

To ease the description, we can assume that each S-node has exactly two children [51]. In particular, we will exploit the following property of \mathcal{T} when G is an *st*-graph containing the edge $e = (s, t)$ and \mathcal{T} is rooted at the Q-node of e .

Property 3 [43] *Let μ be a node of \mathcal{T} with poles u and v . We have that $\text{pert}(\mu)$ is either a uv -graph or a vu -graph.*

6.2 Branch decomposition and sphere-cut decomposition

A *branch decomposition* $\langle T, \xi \rangle$ of a graph G consists of an unrooted binary tree T (i.e., each internal node of T has degree three) and of a bijection $\xi : \mathcal{L}(T) \leftrightarrow E(G)$ between the leaf set $\mathcal{L}(T)$ of T and the edge set $E(G)$ of G . For each arc a of T , let T_1^a and T_2^a be the two connected components of $T - a$, and, for $i = 1, 2$, let G_i^a be the subgraph of G that consists of the edges corresponding to the leaves of T_i^a , i.e., the edge set $\{\xi(\mu) : \mu \in \mathcal{L}(T) \cap V(T_i^a)\}$. The *middle set* $\text{mid}(a) \subseteq V(G)$ is the intersection of the vertex sets of G_1^a and G_2^a , i.e., $\text{mid}(a) := V(G_1^a) \cap V(G_2^a)$. The *width* $\text{fi}(\langle T, \xi \rangle)$ of $\langle T, \xi \rangle$ is the maximum size of the middle sets over all arcs of T , i.e., $\text{fi}(\langle T, \xi \rangle) := \max\{|\text{mid}(a)| : a \in E(T)\}$. An *optimal branch decomposition* of G is a branch decomposition with minimum width; this width is called the *branchwidth* of G and is denoted by $\text{fi} = \text{bw}(G)$.

Let Σ be a sphere. A Σ -*plane* graph G is a planar graph G embedded (i.e., topologically drawn) on Σ . A *noose* of a Σ -plane graph G is a closed simple curve on Σ that (i) intersects G only at vertices and (ii) traverses each face at most once.

The *length* of a noose O is the number of vertices it intersects. Every noose O bounds two closed discs Δ_O^1 and Δ_O^2 in Σ , i.e., $\Delta_O^1 \cap \Delta_O^2 = O$ and $\Delta_O^1 \cup \Delta_O^2 = \Sigma$.

Let G be a Σ -plane graph and let $\langle T, \xi \rangle$ be a branch decomposition of G . Suppose that, for each arc a of T , there exists a noose O_a that traverses all and only the vertices of $\text{mid}(a)$ and whose closed discs $\Delta_{O_a}^1$ and $\Delta_{O_a}^2$ enclose the drawing of G_1^a and of G_2^a , respectively. Denote by π_a the circular order of the vertices in $\text{mid}(a)$ as they appear along O_a and by $\Pi = \{\pi_a : a \in E(T)\}$ the set of all the circular orders π_a . The triple $\langle T, \xi, \Pi \rangle$ is a *sphere-cut decomposition* of G . A clockwise traversal of O_a in the drawing of G defines the cyclic ordering π_a of $\text{mid}(a)$. We always assume that the vertices of each middle set $\text{mid}(a) = V(G_1^a) \cap V(G_2^a)$ are enumerated according to π_a . Observe that, Condition II of the definition of noose implies that, for any arc $a \in E(T)$, the graphs G_1^a and G_2^a are both connected. Also, the nooses are *pairwise non-crossing*, i.e., for any two nooses O_a and O_b , it holds that O_b lies either entirely

inside $\Delta_{O_a}^1$ or entirely inside $\Delta_{O_a}^2$. We will exploit the following result by Dorn et al [52, Theorem 1].

Theorem 6 [52] *Let G be a connected n -vertex Σ -plane graph having branchwidth fi and no vertex of degree one. There exists a sphere-cut decomposition of G having width fi which can be constructed in $O(n^3)$ time.*

6.3 Algorithm description

Without loss of generality, we assume that the input st -graph G contains the edge (s, t) , which guarantees that G is biconnected. In fact, in any 2UBE of G vertices s and t have to be the first and the last vertex of the spine, respectively. Thus, either (s, t) is an edge of G or it can be added to any of the two pages of a 2UBE of G to obtain a 2UBE $\langle \pi, \sigma \rangle$ of $G \cup (s, t)$. Clearly, the edge (s, t) will be incident to the outer face of $\Gamma(\pi, \sigma)$.

Overview. Our approach builds upon a classification of the embeddings of each triconnected component of the biconnected graph G . Intuitively, such a classification is based on the visibility of the spine that the embedding “leaves” on its outer face. Let \mathcal{T} be the SPQR-tree of G with respect to the reference edge (s, t) . For any node μ of \mathcal{T} , we show that the planar embeddings of $\text{pert}(\mu)$ that yield a 2UBE of $\text{pert}(\mu)$ can be partitioned into a finite number of equivalence classes, called *embedding types*. By visiting \mathcal{T} bottom-up, we describe how to compute all the *realizable* embedding types of the pertinent graph of each node, that is, those embedding types that are allowed by some 2UBE of it. To this aim we will exploit the realizable embedding types of the pertinent graphs of its children. If there exists at least one embedding type that is realizable by the pertinent graph of the unique child of the root of \mathcal{T} , which coincides with the whole graph G minus the edge (s, t) , then G admits a 2UBE. The most challenging part of this approach is handling the triconnected components that correspond to the P-nodes, where the problem is reduced to a maximum flow problem on a capacitated flow network with edge demands, and to the R-nodes, where a sphere-cut decomposition of bounded width is used to efficiently compute the feasible embedding types.

Embedding Types Given a 2UBE $\langle \pi, \sigma \rangle$, the two pages will be called the *left page* (the one to the left of the spine) and the *right page* (the one to the right of the spine), respectively. We write $\sigma(e) = L$ (resp. $\sigma(e) = R$) if the edge e is assigned to the left page (resp. to the right page). A point p of the spine is *visible from the left (right) page* if it is possible to shoot a horizontal ray originating from p and directed leftward (rightward) without intersecting any edge in $\Gamma(\pi, \sigma)$.

Let μ be a node of the SPQR-tree \mathcal{T} of G rooted at (s, t) . Recall that, by Property 3, since \mathcal{T} has been rooted at (s, t) , the pertinent graph $\text{pert}(\mu)$ and the skeleton $\text{skel}(\mu)$ of μ are $s't'$ -graphs, where s' and t' are the poles of μ . We denote by s_μ (t_μ) the pole of μ that is the source (resp. the sink) of $\text{pert}(\mu)$ and of $\text{skel}(\mu)$. Let $\langle \pi_\mu, \sigma_\mu \rangle$ be a 2UBE of $\text{pert}(\mu)$. In the following we use the symbols L , R , N , and B as abbreviations for *left*, *right*, *none*, and *both*, respectively. We say that $\langle \pi_\mu, \sigma_\mu \rangle$ has *embedding type* (for short,

has Type or is of Type) $\langle s_vis, spine_vis, t_vis \rangle$, with $s_vis, t_vis \in \{L, R, N\}$ and $spine_vis \in \{L, R, B, N\}$, where:

1. s_vis is L (resp., R), if in $\Gamma(\pi_\mu, \sigma_\mu)$ there is a portion of the spine incident to s and between s and t that is visible from the left page (resp., from the right page).
If s_vis is neither L nor R , then s_vis is N .
2. t_vis is L (resp., R), if in $\Gamma(\pi_\mu, \sigma_\mu)$ there is a portion of the spine incident to t and between s and t that is visible from the left page (resp., from the right page).
If t_vis is neither L nor R , then t_vis is N .
3. $spine_vis$ is L (resp., R), if in $\Gamma(\pi_\mu, \sigma_\mu)$ there is a portion of the spine between s and t that is visible from the left page (resp., from the right page); $spine_vis$ is B , if in $\Gamma(\pi_\mu, \sigma_\mu)$ there is a portion of the spine between s and t that is visible from the left page, and a portion of the spine between s and t that is visible from the right page. If $spine_vis$ is neither L , nor R , nor B , then $spine_vis$ is N .

Further, we say that a node μ and $\text{pert}(\mu)$ admit embedding type (for short, admit Type) $\langle s_vis, spine_vis, t_vis \rangle$, if $\text{pert}(\mu)$ admits a 2UBE of Type $\langle s_vis, spine_vis, t_vis \rangle$; in this case, we also say that the Type $\langle s_vis, spine_vis, t_vis \rangle$ is realizable by μ and by $\text{pert}(\mu)$. We have the following.

Lemma 4 *Let μ be a node of \mathcal{T} and let $\langle \pi_\mu, \sigma_\mu \rangle$ be a 2UBE of $\text{pert}(\mu)$.*

Then, $\langle \pi_\mu, \sigma_\mu \rangle$ has exactly one embedding type, where the possible embedding types are the 18 depicted in Fig. 12.

Proof The first part of the statement follows from the definition of embedding type. To see that the number of possible embedding types allowed by a 2UBE of $\text{pert}(\mu)$ is at most 18, it suffices to consider the following facts. First, the number of different embedding types admitted by $\text{pert}(\mu)$ is at most 36. Further, some combinations are “impossible”, in the sense that not all combinations of values for s_vis , $spine_vis$, and t_vis appear in a 2UBE of $\text{pert}(\mu)$. In particular, we have that the following values for $spine_vis$ are forbidden: (a) $spine_vis$ cannot be L , if either $s_vis = R$ or $t_vis = R$; (b) $spine_vis$ cannot be R , if either $s_vis = L$ or $t_vis = L$; (c) $spine_vis$ cannot be N , if either $s_vis \neq N$ or $t_vis \neq N$. Condition a and Condition b rule out 5 combinations each, while Condition c rules out 8 more combinations (see also Table 1). This leaves us with the 18 embedding types shown in Fig. 12. □

Recall that, for each node μ of \mathcal{T} , $\text{pert}(\mu)$ may have exponentially many planar embeddings, given by the permutations of the children of the P-nodes and by the flips of the R-nodes. Moreover, there may exist exponentially many 2UBEs associated with each such planar embedding. The next lemma will be the key tool to help us control this combinatorial explosion.

Let $\langle \pi, \sigma \rangle$ be a 2UBE of G , let μ be a node of \mathcal{T} , and let $\langle \pi_\mu, \sigma_\mu \rangle$ be the restriction of $\langle \pi, \sigma \rangle$ to $\text{pert}(\mu)$. Also, let $\langle \pi'_\mu, \sigma'_\mu \rangle \neq \langle \pi_\mu, \sigma_\mu \rangle$ be a 2UBE of $\text{pert}(\mu)$. We prove the following.

Lemma 5 *If $\langle \pi'_\mu, \sigma'_\mu \rangle$ and $\langle \pi_\mu, \sigma_\mu \rangle$ have the same embedding type, then G admits a 2UBE whose restriction to $\text{pert}(\mu)$ is $\langle \pi'_\mu, \sigma'_\mu \rangle$.*

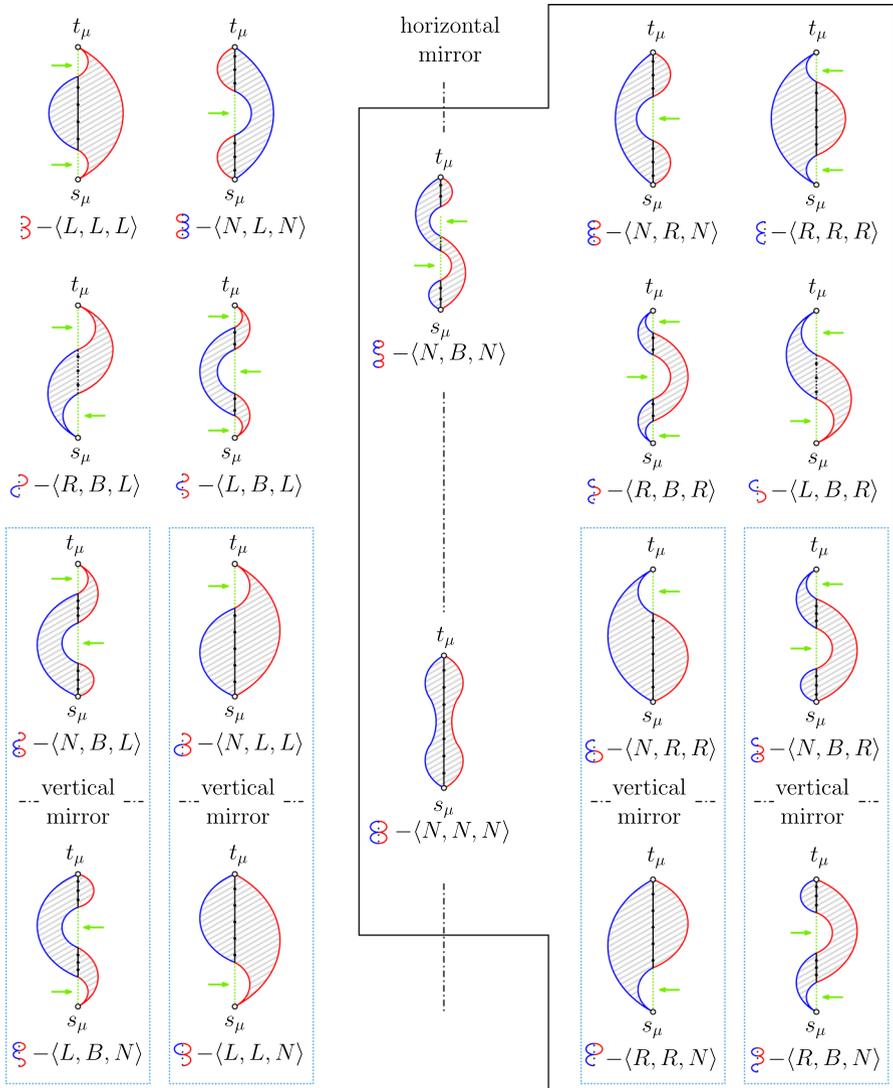


Fig. 12 Illustration of the possible embedding types of a 2UBE of $\text{pert}(\mu)$ for a node μ with poles s_μ and t_μ ; the portions of the spine that are visible from the left page or from the right page are green. Embedding types that are in the same dotted box are one the vertically-mirrored copy of the other. Embedding types on the left are the horizontally-mirrored copies of the ones on the right. Embedding types $\mathfrak{B}-\langle N, B, N \rangle$ and $\mathfrak{G}-\langle N, N, N \rangle$ are the horizontally-mirrored and vertically-mirrored copies of themselves

Table 1 Embedding types that are not possible

	Ruled out by Condition a					Condition b					Condition c							
<i>t_vis</i>	R	R	R	L	N	L	L	L	R	N	L	L	R	R	L	R	N	N
<i>spine_vis</i>	L	L	L	L	L	R	R	R	R	R	N	N	N	N	N	N	N	N
<i>s_vis</i>	L	R	N	R	R	R	L	N	L	L	L	R	L	R	N	N	L	R

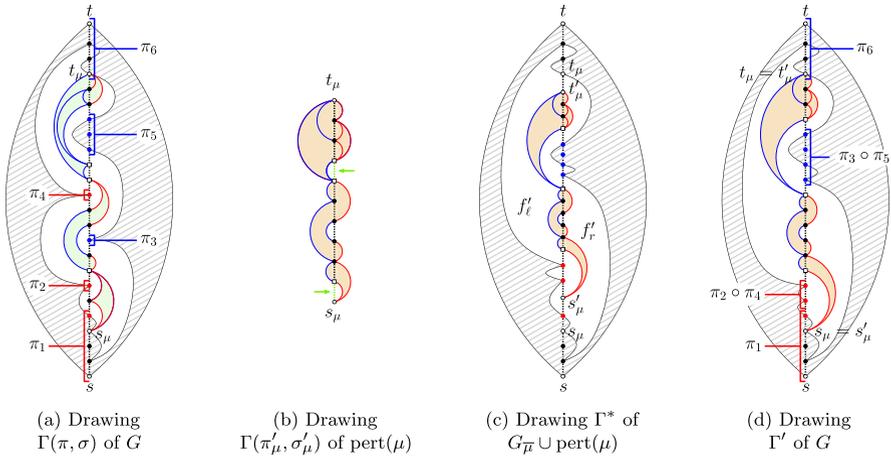


Fig. 13 Illustrations for Lemma 5. The 2UBE of $\text{pert}(\mu)$ are of Type $\mathbb{S}-(L, B, N)$

Proof Notice that if $\text{pert}(\mu)$ is an edge, the statement does not apply, because $\text{pert}(\mu)$ does not admit two distinct 2UBEs of the same type. Thus we assume that $\text{pert}(\mu)$ has at least three vertices. We show how to construct a 2UBE $\langle \pi', \sigma' \rangle$ of G whose restriction to $\text{pert}(\mu)$ is $\langle \pi'_\mu, \sigma'_\mu \rangle$.

For the ease of description we actually show how to construct an upward planar drawing Γ' of G in which each vertex v lies along the spine in the same bottom-to-top order determined by π' and each edge is drawn on the page assigned by σ' . Clearly, such a drawing implies the existence of $\langle \pi', \sigma' \rangle$.

Consider a drawing $\Gamma(\pi, \sigma)$ of G ; refer to Fig. 13a. Let $G_{\bar{\mu}}$ be the subgraph of G composed of the edges of G not in $\text{pert}(\mu)$ and of their endpoints. Consider the drawing $\Gamma_{\bar{\mu}}$ of $G_{\bar{\mu}}$ obtained by restricting $\Gamma(\pi, \sigma)$ to $G_{\bar{\mu}}$. Denote by $\pi_{\bar{\mu}}$ the bottom-to-top order of the vertices of $G_{\bar{\mu}}$ in $\Gamma_{\bar{\mu}}$. Let $\pi_1, \pi_2, \dots, \pi_k$ be the maximal subsequences of $\pi_{\bar{\mu}}$ between s and t and composed of consecutive vertices in $\pi_{\bar{\mu}}$ that are also consecutive in π (refer to Fig. 13a). Observe that sequences $\pi_i, i = 1, \dots, k$, may be formed by a single vertex or by multiple vertices. Also, the first sequence π_1 includes s_μ and the last sequence π_k includes t_μ . Further, since $\text{pert}(\mu)$ has at least three vertices, there exists at least a vertex of $\text{pert}(\mu)$ between s_μ and t_μ in $\Gamma(\pi, \sigma)$; therefore we have that $k \geq 2$.

Drawing $\Gamma_{\bar{\mu}}$ contains a face f_μ that is incident to the ending vertex of the sequence π_1 , to the starting vertex of the sequence π_k , and, if $k \geq 3$, to all the starting and

ending vertices of the sequences π_i , with $i = 2, \dots, k - 1$. Face f_μ always contains the vertices s_μ and t_μ . Some starting and ending vertices of the sequences π_i are encountered when traversing f_μ clockwise from s_μ to t_μ (*left vertices of f_μ*) and some of them are encountered when traversing f_μ counter-clockwise from s_μ to t_μ (*right vertices of f_μ*). The left (resp. right) vertices of f_μ are red (resp. blue) in Fig. 13.

We show how to insert into face f_μ the drawing $\Gamma(\pi'_\mu, \sigma'_\mu)$ of $\text{pert}(\mu)$, producing the desired drawing Γ' of the 2UBE $\langle \pi', \sigma' \rangle$ of G . First, consider the drawing $\Gamma(\pi'_\mu, \sigma'_\mu)$ of $\text{pert}(\mu)$ (see Fig. 13b) and insert it, possibly after squeezing it, into f_μ in such a way that its spine lays entirely on the line of the spine of $\Gamma_{\bar{\mu}}$ and in such a way that the vertices of μ do not fall in between the vertices of any maximal sequence π_i , $i = 1, \dots, k$. This is always possible, because $k \geq 2$ and there exists a portion of the spine of $\Gamma_{\bar{\mu}}$ that is in the interior of f_μ . Observe that this implies that s_μ and t_μ have now a double representation, since the drawing of the source s'_μ and sink t'_μ of $\Gamma(\pi'_\mu, \sigma'_\mu)$ do not coincide with the drawing of s_μ and t_μ in $\Gamma_{\bar{\mu}}$. Denote by Γ^* the resulting drawing (refer to Fig. 13c).

Suppose $\langle \pi_\mu, \sigma_\mu \rangle$ (and, hence, also $\langle \pi'_\mu, \sigma'_\mu \rangle$) is of Type $\langle x, y, z \rangle$. Observe that if $y \in \{N, R\}$ then there are no left vertices of f_μ . Otherwise, if $y \in \{L, B\}$, then it is possible to identify two vertices $v_{\ell,b}$ and $v_{\ell,t}$ of $\Gamma(\pi'_\mu, \sigma'_\mu)$ such that there exists a portion of the spine between $v_{\ell,b}$ and $v_{\ell,t}$ that is visible from left. Consider all the vertices of the sequences π_i whose starting and ending vertices are left vertices of f_μ , and move them between $v_{\ell,b}$ and $v_{\ell,t}$ along the spine of Γ^* , while preserving their relative order. Analogously, if $y \in \{N, L\}$, there are no right vertices of f_μ . Otherwise, if $y \in \{R, B\}$, it is possible to identify two vertices $v_{r,b}$ and $v_{r,t}$ of $\Gamma(\pi'_\mu, \sigma'_\mu)$ such that there exists a portion of the spine between $v_{r,b}$ and $v_{r,t}$ that is visible from right. Consider all the vertices of the sequences π_i whose starting and ending vertices are right vertices of f_μ , and move them between $v_{r,b}$ and $v_{r,t}$ along the spine of Γ^* , while preserving their relative order.

Γ^* is an upward planar drawing of G except for the fact that there are two duplicated vertices s'_μ and t'_μ . Thus, by identifying s_μ with s'_μ and t_μ with t'_μ , we obtain the desired upward drawing Γ' of G in which each vertex v lies along the spine and each edge is drawn as a y -monotone curve on one of the two half-planes defined by the spine. Next we prove that Γ' is planar. Clearly, identifying s_μ with s'_μ (t_μ with t'_μ) does not introduce crossings when s_μ and s'_μ (t_μ and t'_μ) are consecutive along the spine in Γ^* (see t'_μ and t_μ in Fig. 13c). In the case in which s_μ and s'_μ (t_μ and t'_μ) are not consecutive along the spine in Γ^* , necessarily $x \in \{L, R\}$ ($z \in \{L, R\}$). Suppose $x = L$ ($z = L$), the case when $x = R$ ($z = R$) being analogous. There cannot exist in $\Gamma_{\bar{\mu}}$ an edge $e = (a, b)$ such that $\pi(a) < \pi(s_\mu) < \pi(b)$ and such that $\sigma(e) = R$. Therefore, we can continuously move s'_μ , together with its incident edges, toward s_μ , remaining inside the region bounded by f_μ without intersecting any edge of $\Gamma_{\bar{\mu}}$ (see Fig. 13d). This concludes the proof. \square

We now describe an algorithm to decide if G admits a 2UBE and its running time. The same procedure can be easily refined to actually compute a 2UBE of G , with no

additional cost, by decorating each node $\mu \in \mathcal{T}$ with the embedding choices performed at μ , for each of its $O(1)$ possible embedding types.

Testing Algorithm. The algorithm is based on visiting the SPQR-tree \mathcal{T} of G and on computing, for each non-root node μ of \mathcal{T} , the set of embedding types realizable by μ , assuming to have already computed the set of embedding types realizable by each of the children of μ . Lemma 5 is the reason why we only need to maintain a single 2UBE of $\text{pert}(\mu)$, for each embedding type realizable by μ , i.e., a constant number of 2UBEs instead of an exponential number. Since, by Lemmas 4 and 5, G admits a 2UBE if and only if the pertinent graph of the unique child of the root Q-node admits an embedding type of (at least) one of the 18 possible embedding types, this approach allows us to solve the 2UBE TESTING problem for G .

The computation of the set of embedding types realizable by μ is based on whether μ is an S-, P-, Q-, or an R-node. Recall that the only possible embedding choices for G happen at P- and R-nodes. While the treatment of Q- and S-nodes does not require any modification when considering the variable and the fixed embedding settings, for P- and R-nodes we will discuss how to compute the embedding types that are realizable by μ in both such settings separately. In particular, in the fixed embedding scenario the above characterization needs to additionally satisfy the constraints imposed by the fixed embedding on the skeletons of the P- and R-nodes in \mathcal{T} . Note that, a leaf Q-node only admits Type \triangleright - $\langle L, L, L \rangle$ and Type \triangleleft - $\langle R, R, R \rangle$. Also, combining 2UBEs of the two children of an S-node μ always yields a valid 2UBE of $\text{pert}(\mu)$. The next lemma shows how to compute all the embedding types realizable by μ .

Lemma 6 *Let μ be an S-node. The set of embedding types realizable by μ can be computed in $O(1)$ time, both in the fixed and in the variable embedding setting.*

Proof Let μ be an S-node with poles s_μ and t_μ . Let μ' and μ'' be the two children of μ with poles $s_{\mu'}$, $t_{\mu'}$ and $s_{\mu''}$, $t_{\mu''}$, respectively, where $s_{\mu'} = s_\mu$, $t_{\mu'} = s_{\mu''}$, and $t_{\mu''} = t_\mu$. Clearly, combining each pair of 2UBEs of the two children of μ always yields a 2UBE of $\text{pert}(\mu)$. Let $\langle x', y', z' \rangle$ and $\langle x'', y'', z'' \rangle$ be any two embedding types realizable by μ' and μ'' , respectively. The embedding type $\langle x, y, z \rangle$ of any 2UBE of $\text{pert}(\mu)$ resulting from a series combination of 2UBEs of the above types can be computed as follows. We have $x = x'$ and $z = z''$. As for y , we have that: (i) $y = N$ iff $y' = y'' = N$; (ii) $y = L$ if either $y' = L$ and $y'' \in \{L, N\}$ or $y'' = L$ and $y' \in \{L, N\}$; (iii) $y = R$ if either $y' = R$ and $y'' \in \{R, N\}$ or $y'' = R$ and $y' \in \{R, N\}$; (iv) $y = B$ if at least one of y' and y'' is B or one of them is L and the other is R . Since there is a constant number of embedding types realizable by the children of μ , the statement follows. \square

Since P- and R-nodes require a more complex treatment, in the following we only state the lemmas concerning the computation time of the set of realizable embedding types for these nodes. We defer the proofs of such lemmas to the next sections. Specifically, in Subsection 6.5 we show how to efficiently compute the set of embedding types realizable by a P-node in the fixed and variable embedding setting (Subsection 6.5.1 and Subsection 6.5.2, respectively). We have the following.

Lemma 7 *Let μ be a P-node with k children. The set of embedding types realizable by $\text{pert}(\mu)$ can be computed in $O(k)$ time in the fixed embedding setting.*

Lemma 8 *Let μ be a P-node with k children. The set of embedding types realizable by $\text{pert}(\mu)$ can be computed in $O(k^2)$ time in the variable embedding setting.*

We remark that Lemmas 6 and 7 yield a counterpart, in the fixed embedding setting, of the linear-time algorithm by Di Giacomo et al. [46] to compute 2UBEs of series-parallel graphs, which we formalize in the following statement.

Theorem 7 *There exists an $O(n)$ -time algorithm to decide whether an n -vertex series-parallel st-graph admits an embedding-preserving 2UBE.*

Finally, in Subsection 6.4 we show how to efficiently compute the set of embedding types realizable by an R-node. In particular, we prove the following.

Lemma 9 *Let μ be an R-node with k children and whose skeleton has branchwidth fi . The set of embedding types realizable by $\text{pert}(\mu)$ can be computed in $O(2^{O(\beta \log \beta)} \cdot k)$ time, both in the fixed and in the variable embedding setting, provided that a sphere-cut decomposition (T, ξ, Π) of width β of $\text{skel}^+(\mu)$ is given.*

By Lemmas 6 to 9, and since T has $O(|G|)$ size [43, 51], we get the following.

Theorem 8 *There exists an $O(2^{O(\beta \log \beta)} \cdot n + n^2 + g(n))$ -time algorithm to decide if an n -vertex planar st-graph (resp. plane st-graph) of branchwidth β admits a 2UBE (resp. an embedding-preserving 2UBE), where $g(n)$ is the computation time of a sphere-cut decomposition of an n -vertex plane graph.*

Since $g(n)$ is $O(n^3)$ by the result in [87], we get the following.

Corollary 2 *There exists an $O(2^{O(\beta \log \beta)} \cdot n + n^3)$ -time algorithm to decide if an n -vertex planar st-graph (resp. plane st-graph) of branchwidth β admits a 2UBE (resp. an embedding-preserving 2UBE).*

Furthermore, since for any n -vertex planar graph, it holds $\beta \leq 2.122\sqrt{n}$ [60], Corollary 2 implies that the 2UBE TESTING problem is solvable in sub-exponential time.

Corollary 3 *There exists an $O(2^{O(\sqrt{n} \log \sqrt{n})} + n^3)$ -time algorithm to decide if an n -vertex planar st-graph (resp. plane st-graph) admits a 2UBE (resp. an embedding-preserving 2UBE).*

6.4 Computing the Embedding Types when μ is an R-node

Let μ be an R-node with poles s_μ and t_μ ; refer to Fig. 14. We consider the graph $\text{skel}^+(\mu) = \text{skel}(\mu) \cup (s_\mu, t_\mu)$, which is a triconnected planar graph. Since $\text{skel}^+(\mu)$ is a minor of G , its branchwidth is bounded by the branchwidth β of G [86, Theorem 4.1]. Let \mathcal{E}_μ be an embedding of $\text{skel}^+(\mu)$ on the sphere and let (T, ξ, Π) be a sphere-cut decomposition of $\text{skel}^+(\mu)$ with embedding \mathcal{E}_μ of width smaller than or equal to fi ; refer to Fig. 14a.

We root T at the leaf ρ of T such that $\xi(\rho) = (s_\mu, t_\mu)$. Consider each arc $a = (p_a, c_a)$ of T connecting the parent node p_a with the child node c_a in T . The noose O_a

of a divides the sphere into two closed regions whose boundary is O_a . We associate a with the subgraph skel_a of $\text{skel}^+(\mu)$ lying in the region not containing the edge (s_μ, t_μ) , and we refer to the face of skel_a that contains such an edge as the *outer face* of skel_a . We denote by pert_a the subgraph of $\text{pert}(\mu)$ obtained by replacing, in skel_a , each virtual edge e with the pertinent graph of the node corresponding to e in the SPQR-tree \mathcal{T} of G .

Intuitively, our strategy to compute the embedding types admitted by $\text{pert}(\mu)$ is to visit the nodes of T in post order. When a node c_a incident to the edge $a = (p_a, c_a)$ is visited, we compute a succinct description of size $O(\text{fi})$ of the properties of the noose O_a in a 2UBE of pert_a . Note that the arc a^* incident to the root ρ is such that skel_{a^*} (resp. pert_{a^*}) coincides with $\text{skel}(\mu)$ (resp. with $\text{pert}(\mu)$). When we reach a^* , we use the computed properties for O_{a^*} to determine which embedding types are realizable by $\text{pert}(\mu)$. The fixed and the variable embedding setting are treated analogously. In the fixed embedding setting, we only consider the embedding of $\text{skel}^+(\mu)$ in which the embedding of $\text{skel}(\mu)$ is inherited by the (fixed) embedding of G . In the variable embedding setting, instead, we consider each of the two embeddings of $\text{skel}^+(\mu)$ obtained by flipping the embedding of $\text{skel}(\mu)$ at its poles.

Let a be an arc of T and let p_0, p_1, \dots, p_k be the sequence of maximal (upward or downward) directed paths traversed in a clockwise visit of the outer face of skel_a . The next lemma shows that $k \in O(\beta)$.

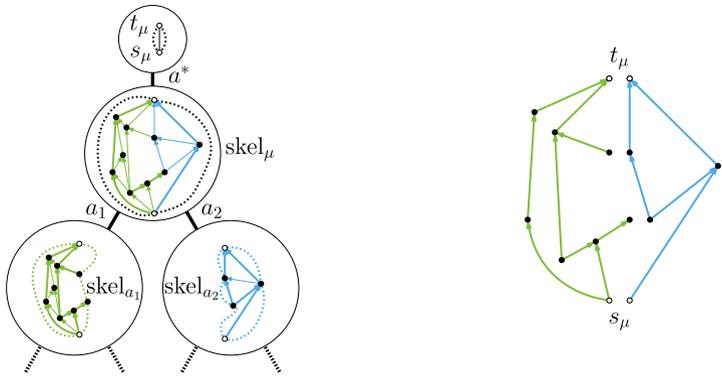
Lemma 10 *For each arc $a \in T$, by clockwise visiting the outer face of skel_a , we traverse at most $O(\text{fi})$ maximal (upward or downward) directed paths.*

Proof Let $p_{u,v}$ be the path along the outer face of skel_a connecting two vertices of $\text{mid}(a)$ that are clockwise consecutive in the circular ordering $\pi_a \in \Pi$ of $\text{mid}(a)$. By Property 3, $\text{pert}(\mu)$ is an $s_\mu t_\mu$ -graph, and thus its faces have a single source and a single sink. If $p_{u,v}$ was composed of more than three maximal (upward or downward) directed paths, then the outer face of skel_a would contain either two sources or two sinks belonging to $p_{u,v}$. The statement then follows from the fact that $|\text{mid}(a)| \leq \text{fi}$. \square

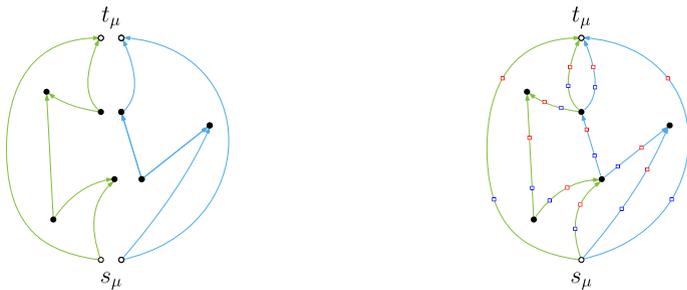
Let u and v be two clockwise consecutive vertices along the outer face of skel_a that are connected by an (upward or downward) directed path p_i . Given a 2UBE $\langle \pi_a, \sigma_a \rangle$ of pert_a , we associate with p_i an *outer-visibility triple* $t_i = \langle u_vis, spine_vis, v_vis \rangle$ of p_i in $\langle \pi_a, \sigma_a \rangle$, for $i = 1, \dots, k$, which encodes the information about the visibility of the spine between u and v from the outer face of $\Gamma(\pi_a, \sigma_a)$. In particular, we have:

- u_vis = true** if and only if there is a portion of the spine incident to u and between u and v that is visible from the outer face of $\Gamma(\pi_a, \sigma_a)$.
- v_vis = true** if and only if there is a portion of the spine incident to v and between u and v that is visible from the outer face of $\Gamma(\pi_a, \sigma_a)$.
- spine_vis = true** if and only if there is a portion of the spine between u and v that is visible from the outer face of $\Gamma(\pi_a, \sigma_a)$.

Observe that when $u_vis = true$ or $v_vis = true$, then necessarily $spine_vis = true$. Hence, we have five possible types for the outer-visibility triples $\langle u_vis, spine_vis, v_vis \rangle$.



(a) A sphere-cut decomposition of $\text{skel}(\mu) \cup (s_\mu, t_\mu)$ (b) Outer faces of skel_{a_1} and skel_{a_2}



(c) Directed paths of skel_{a_1} and skel_{a_2} (d) Graph $A = A_1 \cup A_2$

Fig. 14 Illustrations for the R-node case. **a** A partial sphere-cut decomposition of the graph $\text{skel}(\mu) \cup (s_\mu, t_\mu)$ rooted at the edge (s_μ, t_μ) , where μ is an R-node. The nooses are dotted curves. **b** Outer faces of graphs skel_{a_1} and skel_{a_2} . **c** The graphs defined by the directed paths on the outer face of skel_{a_1} and skel_{a_2} . **d** The auxiliary graph A for skel_a

We denote by X_a the set of all the endpoints of the paths p_1, \dots, p_k . Let $\langle \pi_\mu, \sigma_\mu \rangle$ be a 2UBE of $\text{pert}(\mu)$, let a be an arc of T , and let $\langle \pi_a, \sigma_a \rangle$ be the restriction of $\langle \pi_\mu, \sigma_\mu \rangle$ to pert_a . Further, let $\langle \pi'_a, \sigma'_a \rangle \neq \langle \pi_a, \sigma_a \rangle$ be a 2UBE of pert_a . The next lemma can be proved with the same strategy as in the proof of Lemma 5.

Lemma 11 *Graph $\text{pert}(\mu)$ admits a 2UBE $\langle \pi'_\mu, \sigma'_\mu \rangle$ whose restriction to pert_a is $\langle \pi'_a, \sigma'_a \rangle$ if the following two conditions hold: 1.the bottom-to-top order of the vertices in X_a is the same in π_a as in π'_a , and 2.for each $i = 1, \dots, k$, the outer-visibility triple of p_i in $\Gamma(\pi'_a, \sigma'_a)$ is the same as the outer-visibility triple of p_i in $\Gamma(\pi_a, \sigma_a)$.*

When visiting T bottom-up, we compute and store in each arc a of T those pairs $(\pi_a, \langle t_1, t_2, \dots, t_k \rangle)$, called *outer-shape pairs*, where π_a is a bottom-to-top order of the vertices in X_a and t_i is an outer-visibility triple for path p_i , for $i = 1, \dots, k$, such that there exists a 2UBE $\langle \pi, \sigma \rangle$ of pert_a satisfying the following properties: 1. ordering π_a is the restriction of π to X_a and 2. t_i is the outer-visibility triple determined by $\langle \pi, \sigma \rangle$ when clockwise traversing the outer face of $\Gamma(\pi, \sigma)$ between the endpoints of p_i . By Lemma 11, this information is enough to succinctly describe all the relevant properties

of the 2UBEs of pert_a . Moreover, by Property 3 $\text{skel}_{a^*} = \text{skel}(\mu)$ is bounded by two directed paths p_1 and p_2 whose end-vertices are s_μ and t_μ . Thus, when we reach the arc a^* , all the embedding types that are realizable by $\text{pert}(\mu)$ can be computed by inspecting the outer-visibility triples t_1 and t_2 , over all outer-shape pairs $(\pi_a, \langle t_1, t_2 \rangle)$ stored in a^* .

Next, we show how to compute all outer-shape pairs $(\pi_a, \langle t_1, t_2, \dots, t_k \rangle)$, for each arc $a \in T$.

Suppose a leads to a leaf ℓ of T such that $\xi(\ell) = (u, v)$. Assume that pert_a is oriented from u to v and that the clockwise boundary of skel_a is composed by the upward path $p_1 = u, v$ and the downward path $p_2 = v, u$, the other case being analogous. Then, all the outer-shape pairs $(\pi_a, \langle t_1, t_2 \rangle)$ for a are such that $\pi_a = u, v$, since pert_a is an uv -graph. Let μ_ℓ be the child of μ in the SPQR-tree \mathcal{T} associated with the virtual edge (u, v) . We compute the outer-visibility triples $t_1 = \langle \alpha, \varphi, \gamma \rangle$ and $t_2 = \langle \alpha', \varphi', \gamma' \rangle$ for p_1 and p_2 , respectively, starting from the embedding types realizable by μ_ℓ , as follows. For each embedding type $\langle s_vis, spine_vis, t_vis \rangle$ of the node μ_ℓ we define t_1 and t_2 as follows. If $s_vis = L$, we set $\alpha = true$ and $\alpha' = false$. If $s_vis = R$ we set $\alpha = false$ and $\alpha' = true$. If $s_vis = N$ we set both $\alpha = \alpha' = false$. Analogously, we set φ and φ' according to the value of t_vis . We set $\gamma = true$ if $spine_vis = B$ or $spine_vis = L$. Otherwise, we set $\gamma = false$. Finally, we set $\gamma' = true$ if $spine_vis = B$ or $spine_vis = R$. Otherwise, we set $\gamma' = false$.

Suppose now that a leads to a non-leaf node λ of T . Then, λ has two children λ_1 and λ_2 , reached by two arcs a_1 and a_2 of T , for which we have already computed all the outer-shape pairs. We are going to use the following observation, which follows from the planarity of $\text{skel}(\mu)$ and from the fact that the nooses are pairwise non-crossing.

Observation 1 Let $\langle \pi_\mu, \sigma_\mu \rangle$ be a 2UBE of $\text{skel}(\mu)$, and let $\langle \pi_{a_1}, \sigma_{a_1} \rangle$ and $\langle \pi_{a_2}, \sigma_{a_2} \rangle$ be the 2UBEs of skel_{a_1} and skel_{a_2} , respectively, determined by $\langle \pi_\mu, \sigma_\mu \rangle$. The following two properties hold: 1. the edges of skel_{a_1} lie in the outer face of the drawing obtained by restricting $\Gamma(\pi_\mu, \sigma_\mu)$ to skel_{a_2} , and vice versa, and 2. the edges of skel_{a_1} and skel_{a_2} do not interleave around a vertex in $\Gamma(\pi_\mu, \sigma_\mu)$.

For $i = 1, 2$, let $A_i = (V_i, E_i)$ be the auxiliary (multi)-graph of pert_{a_i} defined as follows. Initialize $V_i = X_{a_i}$ and add to E_i a directed edge (u, v) for each path p directed from u to v traversed when clockwise visiting the outer face of skel_{a_i} ; see Fig. 14c. Then, replace each directed edge (u, v) with a directed path (u, x', x'', v) , unless (u, v) is an edge in G . Clearly, graph A_i is connected. Note that, there exists a one-to-one correspondence between the directed paths along the outer face of skel_{a_i} (see Fig. 14b) and the directed paths along the outer face of A_i (see Fig. 14c). We have that each 2UBE of A_i defines a set $\langle t_1, \dots, t_{k_i} \rangle$, where t_i is an outer-visibility triple describing the visibility of each directed path along the outer face of A_i in the 2UBE. We define a new graph A as the union of A_1 and A_2 ; see Fig. 14d. The purpose of graph A is that of representing the visibility of the spine on the outer face of pert_a between pairs of vertices in X_a in a 2UBE of pert_a . In particular, by assigning to the faces of the book embedding the three edges that replace each directed path p_i along the outer face of pert_a , we are able to model the outer-visibility triples of paths p_1, \dots, p_k . We compute all outer-shape pairs $\langle \pi_a, \langle t_1, \dots, t_k \rangle \rangle$ for a using Algorithm 1.

Algorithm 1 Procedure to compute the outer-shape pairs of an arc $a \in T$.

- Step 1** Construct, by brute force, the set \mathcal{U} of all the 2UBEs $\langle \pi_A, \sigma_A \rangle$ of A .
- Step 2** Let \mathcal{U}^* be the subset of the 2UBEs $\langle \pi_A, \sigma_A \rangle \in \mathcal{U}$ such that the drawing of A_1 lies in the outer face of the drawing of A_2 in $\Gamma(\pi_A, \sigma_A)$, and vice versa, and the edges of these graphs do not interleave around a vertex.
- Step 3** For each 2UBE $\langle \pi_A, \sigma_A \rangle \in \mathcal{U}^*$ of A , perform the following operations:
- Step 3.a** Let $\langle \pi_{A_i}, \sigma_{A_i} \rangle$ be the 2UBE of A_i obtained by restricting $\langle \pi_A, \sigma_A \rangle$ to A_i . Visit clockwise the outer face of $\Gamma(\pi_{A_i}, \sigma_{A_i})$ of A_i to compute the sequence $\langle t_1^i, \dots, t_{k_i}^i \rangle$ of the outer-visibility triples of the maximal directed paths between the vertices in X_{A_i} .
 - Step 3.b** If the set of outer-shape pairs stored in a_1 contains $\langle \pi_{a_1}, \langle t_1^1, \dots, t_{k_1}^1 \rangle \rangle$ and the set of outer-shape pairs stored in a_2 contains $\langle \pi_{a_2}, \langle t_1^2, \dots, t_{k_2}^2 \rangle \rangle$, then compute a pair $\langle \pi_a, \langle t_1, \dots, t_k \rangle \rangle$ for a as follows:
 - Step 3.b.1** Set π_a as the restriction of π_A to the vertices in $X_a \subseteq V(A)$;
 - Step 3.b.2** Visit clockwise the outer face of $\Gamma(\pi_A, \sigma_A)$ to compute the sequence $\langle t_1^i, \dots, t_{k_i}^i \rangle$ of the outer-visibility triples of the maximal directed paths between the vertices in X_a .

We now prove Lemma 9, which we restate for the reader’s convenience.

Lemma 9 Let μ be an R-node with k children and whose skeleton has branchwidth fi . The set of embedding types realizable by $\text{pert}(\mu)$ can be computed in $2^{O(\text{fi} \log \text{fi})} \cdot k$ time, both in the fixed and in the variable embedding setting, provided that a sphere-cut decomposition $\langle T, \xi, \Pi \rangle$ of width β of $\text{skel}^+(\mu)$ is given.

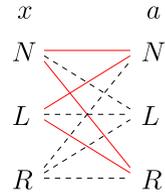
Proof We exploit Algorithm 1 on the arcs of T to compute their outer-shape pairs. As already observed, this allows us to compute the embedding types realizable by $\text{pert}(\mu)$ from the outer-shape pairs of the arc a^* incident to the root ρ of T . The correctness of Algorithm 1 descends from Lemma 4, Lemma 11 and Observation 1.

We argue about the running time of Algorithm 1 for a given arc a of T_μ . Note that the auxiliary graph A corresponding to a contains $O(\beta)$ vertices, by Lemma 10, and $O(\beta)$ edges, since it is planar. First, the running time of **Step 1** is bounded by the number of possible 2UBEs of A . By enumerating all the $|V(A)|! \in O(\beta)!$ linear orders of $V(A)$ and all the $2^{|E(A)|} \in 2^{O(\beta)}$ edge assignments for $E(A)$, we may construct $2^{O(\beta \log \beta)}$ 2UBEs of A . **Step 2** is also bounded by the numbers of 2UBEs of A , as testing whether the conditions of Observation 1 are satisfied by each such a 2UBE can be performed in time linear in the size of A , i.e., $O(\beta)$ time. Clearly, for each of the 2UBEs of A considered at **Step 3**, the remaining steps of Algorithm 1 can also be performed in $O(\beta)$ time. Thus, Algorithm 1 runs in $2^{O(\beta \log \beta)}$ time. The overall running time for computing the embedding types of an R-node μ follows from the fact that tree T contains $O(k)$ nodes and arcs, since it is a binary tree whose leaves are the k virtual edges of $\text{skel}(\mu)$, and from the time spent to compute the outer-shape pairs for each arc of T , i.e., the running time of Algorithm 1. \square

6.5 Computing the Embedding Types when μ is a P-node

We show how to efficiently compute the set of embedding types realizable by the pertinent graph of a P-node. The fixed and variable embedding settings are treated in Subsection 6.5.1 and Subsection 6.5.2, respectively.

Fig. 15 Solid lines represent rejection cases when $x \in \{N, L\}$ and $a \in \{N, R\}$



6.5.1 P-nodes: Fixed Embedding

Let μ be a P-node with poles s_μ and t_μ and let $\mu_1, \mu_2, \dots, \mu_k$ be the children of μ in the left-to-right order defined by the given embedding of G .

Let $\text{skel}^i(\mu)$ and $\text{pert}^i(\mu)$ be the subgraphs of $\text{skel}(\mu)$ and of $\text{pert}(\mu)$, respectively, determined by the children $\mu_1, \mu_2, \dots, \mu_i$ of μ ($1 \leq i \leq k$). Since $\text{pert}^1(\mu) = \text{pert}(\mu_1)$, the set X_1 of embedding types that are realizable by $\text{pert}^1(\mu)$ are those that are realizable by $\text{pert}(\mu_1)$. For $i = 2, \dots, k$, we can compute the set X_i of embedding types that are realizable by $\text{pert}^i(\mu)$ as follows. We consider each embedding type $\langle x, y, z \rangle \in X_{i-1}$ and each embedding type $\langle a, b, c \rangle$ that is realizable by $\text{pert}(\mu_i)$, and we either reject the pair, if no embedding type can be obtained by composing $\langle x, y, z \rangle$ and $\langle a, b, c \rangle$, or compute an embedding type $\langle p, q, w \rangle$ for $\text{pert}^i(\mu)$, which we add to X_i . The sequence of the two tests is described below.

Rejection:

- If $x \in \{N, L\}$ and $a \in \{N, R\}$, we reject the pair (refer to Fig. 15).
- If $z \in \{N, L\}$ and $c \in \{N, R\}$, we reject the pair.
- If $y \in \{N, L\}$, we reject the pair unless μ_i is a Q-node (which admits Type \curvearrowright - $\langle L, L, L \rangle$). In fact, if μ_i is not a Q-node, $\text{pert}(\mu_i)$ contains at least a vertex different from s_μ and from t_μ that has to be placed on the spine to the right of $\text{pert}^{i-1}(\mu)$ and between s_μ and t_μ . However, in a 2UBE of $\text{pert}^{i-1}(\mu)$ of Type $\langle x, y, z \rangle$ with $y \in \{L, N\}$, there exists no portion of the spine between s_μ and t_μ that is visible from the right page.
- If $b \in \{N, R\}$, by a similar argument as in the previous case, we reject the pair unless $i = 2$ and μ_1 is a Q-node (which admits Type \curvearrowleft - $\langle R, R, R \rangle$).

Observe that after the above rejections either μ_i or $\text{pert}^{i-1}(\mu)$ are single edges, or all the conditions below hold:

- $x = R$ or $a = L$ (refer to Fig. 15).
- $z = R$ or $c = L$.
- $y \notin \{N, L\}$.
- $b \notin \{N, R\}$.

Embedding Types for $\text{pert}^i(\mu)$:

Regarding p : If $x = a$ (i.e., both are L or both are R) then $p = x$, otherwise (either $x = N$, or $a = N$, or $x = R$ and $a = L$) $p = N$.

Regarding w : If $z = c$ (i.e., both are L or both are R) then $w = z$, otherwise (either $z = N$, or $c = N$, or $z = R$ and $c = L$) $w = N$.

Regarding q :

- If $y \in \{N, L\}$ (and, thus, μ_i is a Q-node) then $q = y$.
- If $y = R$ and $b = L$ then $q = N$. If $y = R$ and $b = B$ then $q = R$. If $y = R$ and $b = R$ ($i = 2$ and μ_1 is a Q-node, which admits Type $\mathcal{C}\text{-}\langle R, R, R \rangle$) then $q = R$.
- If $y = B$ (observe that, since $\text{pert}^{i-1}(\mu)$ is not a single edge, $b \in \{L, B\}$) then $q = b$.

Since the set of rules above defines the set X_k of embedding types realizable by $\text{pert}^k(\mu) = \text{pert}(\mu)$ and since the above computations can easily be performed in time linear in the number of children of μ , we obtain Lemma 7.

6.5.2 P-nodes: Variable Embedding

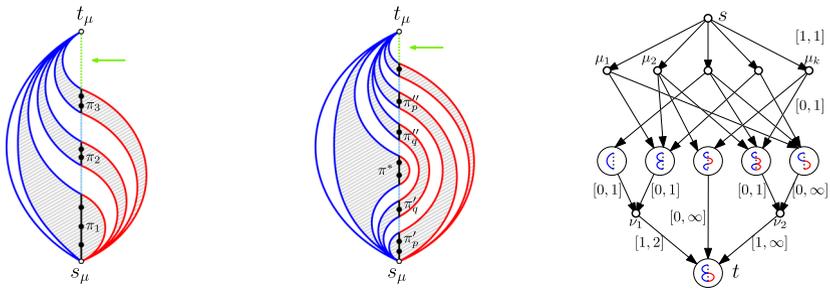
Let μ be a P-node with poles s_μ and t_μ , and with children μ_1, \dots, μ_k . A 2UBE γ of $\text{pert}(\mu)$ is obtained by choosing a permutation for the children of μ , a 2UBE γ_i for each child μ_i ($i = 1, \dots, k$), and a bottom-to-top order of the vertices of $\text{pert}(\mu)$ whose restriction to each μ_i corresponds to the bottom-to-top order of the vertices in γ_i ; the embedding types that are realizable for μ depend on the embedding types that are realizable for its children. Our approach to compute the realizable types of $\text{pert}(\mu)$ consists of considering one type at a time for μ . For each embedding type, we check whether the children of μ , together with their realizable embedding types, can be arranged in a finite number of families of permutations (which we prove to be a constant number) so to yield a 2UBE of the considered embedding type. In order to ease the following description, consider that the arrangements of the children for obtaining some embedding types can be easily derived from the arrangements to obtain the horizontally-mirrored ones by (i) reversing the left-to-right sequence of the children in the construction and (ii) by taking, for each child, the horizontally-mirrored embedding type; for instance, the arrangements to construct a 2UBE of Type $\mathcal{S}\text{-}\langle L, L, N \rangle$ can be obtained from the ones to construct a 2UBE of Type $\mathcal{P}\text{-}\langle R, R, N \rangle$, and vice versa. Moreover, two embedding types, namely Type $\mathcal{S}\text{-}\langle N, B, N \rangle$ and Type $\mathcal{S}\text{-}\langle N, N, N \rangle$, are (horizontally) self-symmetric. As a consequence, in order to consider all the embedding types that are realizable by $\text{pert}(\mu)$ we describe how to obtain only ten “relevant” embedding types (enclosed by a solid polygon in Fig. 12): $\mathcal{S}\text{-}\langle N, R, R \rangle$, $\mathcal{S}\text{-}\langle L, B, R \rangle$, $\mathcal{P}\text{-}\langle R, B, R \rangle$, $\mathcal{S}\text{-}\langle R, R, R \rangle$, $\mathcal{S}\text{-}\langle N, B, R \rangle$, $\mathcal{P}\text{-}\langle R, B, N \rangle$, $\mathcal{S}\text{-}\langle N, B, N \rangle$, $\mathcal{P}\text{-}\langle R, R, N \rangle$, $\mathcal{S}\text{-}\langle N, R, N \rangle$, and $\mathcal{S}\text{-}\langle N, N, N \rangle$.

In this section, we give necessary and sufficient conditions under which the pertinent graph of a P-node admits Type $\mathcal{S}\text{-}\langle N, R, R \rangle$. Then, we show how to test these conditions efficiently by exploiting a suitably defined flow network. The conditions for the remaining relevant embedding types, which can be tested similarly to Type $\mathcal{S}\text{-}\langle N, R, R \rangle$, are reported in Appendix A.

Conditions for Type $\mathcal{S}\text{-}\langle N, R, R \rangle$. Let μ be a P-node. We define the following two cases based on the existence of a child of μ that admits Type $\mathcal{S}\text{-}\langle N, R, R \rangle$:

NR1.1 The children of μ can be partitioned into two sets: The first set consists of a (possibly missing) Q-node child and a child that admits Type $\mathcal{S}\text{-}\langle N, R, R \rangle$. The second set consists of any number (possibly zero) of children that admit Type $\mathcal{S}\text{-}\langle L, B, R \rangle$.

NR1.2 The children of μ can be partitioned into three sets: The first set consists either of a Q-node child, or of a non-Q-node child that admits Type $\mathcal{S}\text{-}\langle R, R, R \rangle$, or



(a) $\mathfrak{S}\text{-}\langle N, R, R \rangle$; Case **NRR.1** (b) $\mathfrak{S}\text{-}\langle N, R, R \rangle$; Case **NRR.2** (c) Network \mathcal{N} for Case **NRR.2**

Fig. 16 **a–b** Cases **NRR.1** and **NRR.2** of Lemma 12 for a P-node μ that admits Type $\mathfrak{S}\text{-}\langle N, R, R \rangle$. The spine is colored either green, blue, or black. The green part is the portion of the spine that is visible from the right, the black parts correspond to the bottom-to-top sequences of the vertices of $\text{pert}(\mu)$ inherited from the 2UBEs of the children of μ (restricted to the vertices different from s_μ and from t_μ), the blue parts join sequences inherited from different children. **c** Capacitated flow network \mathcal{N} with edge demands corresponding to illustration **b**

both. The second set consists of any number (possibly zero) of children that admit Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$. The third set consists of any positive number of children that admit Type $\mathfrak{S}\text{-}\langle N, B, R \rangle$ or Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$, with at most one admitting only Type $\mathfrak{S}\text{-}\langle N, B, R \rangle$.

Lemma 12 *A P-node μ admits Type $\mathfrak{S}\text{-}\langle N, R, R \rangle$ in the variable embedding setting if and only if Case **NRR.1** or Case **NRR.2** holds.*

Proof Let μ_1, \dots, μ_k be the children of μ . About the necessity, suppose that μ admits a 2UBE γ of Type $\mathfrak{S}\text{-}\langle N, R, R \rangle$ and denote by γ_i the restriction of γ to μ_i . At most one γ_i can be of Type $\langle N, \cdot, \cdot \rangle$ and all of them must be of Type $\langle \cdot, \cdot, R \rangle$. Therefore, all of the γ_i must be of Type $\langle R, R, R \rangle$, or $\langle N, R, R \rangle$ or $\langle \cdot, B, R \rangle$. It follows that either one of the γ_i has Type $\mathfrak{S}\text{-}\langle N, R, R \rangle$ (Case **NRR.1**) or one has Type $\mathfrak{S}\text{-}\langle N, B, R \rangle$, or none of the two (Case **NRR.2**). In Case **NRR.1**, there is nothing else to be proved for the necessity. In Case **NRR.2**, one of γ_i corresponds to a Q-node and has Type $\mathfrak{Q}\text{-}\langle R, R, R \rangle$, or it corresponds to a non-Q-node and has Type $\mathfrak{F}\text{-}\langle R, R, R \rangle$, or both because otherwise γ would be of Type $\langle \cdot, B, R \rangle$; further one of the γ_i must have Type $\mathfrak{S}\text{-}\langle N, B, R \rangle$ or $\mathfrak{S}\text{-}\langle L, B, R \rangle$, because otherwise γ would be of Type $\langle R, \cdot, R \rangle$.

For the sufficiency, we consider the two cases separately. In Case **NRR.1** we order the children of μ so that from left to right we have the possible Q-node child, if any, followed by the unique child that admits Type $\mathfrak{S}\text{-}\langle N, R, R \rangle$, followed by the children that admit Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$ in any order; see Fig. 16a. In Case **NRR.2** we order the children of μ so that from left to right we have the possible Q-node child, if any, followed by the possible non-Q-node child that admits Type $\mathfrak{F}\text{-}\langle R, R, R \rangle$, followed by the children that admit Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$ in any order, followed by the possible child that admits only Type $\mathfrak{S}\text{-}\langle N, B, R \rangle$, if any, followed by the children that admit Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$ in any order, if any; see Fig. 16b.

Assume that $\mu_1, \mu_2, \dots, \mu_k$ are ordered as defined above, according to the two cases **NRR.1** and **NRR.2**. We show that such an ordering of the children of μ can be used to construct a 2UBE of $\text{pert}(\mu)$ of Type $\mathfrak{S}\text{-}\langle N, R, R \rangle$. For each μ_i let $\langle \pi_i, \sigma_i \rangle$ be

a 2UBE of $\text{pert}(\mu_i)$ of the type of interest (i.e. the one considered in the description of the ordering); if μ_1 is a Q-node, then we consider $\langle \overline{\pi}_1, \sigma_1 \rangle$ of Type $\mathbb{C}\text{-}\langle R, R, R \rangle$, i.e., $\sigma_1((s_\mu, t_\mu)) = L$. We denote by π_i the restriction of $\overline{\pi}_i$ to the vertices of $\text{pert}(\mu)$ different from s_μ and from t_μ . We now describe how to construct a 2UBE $\langle \pi, \sigma \rangle$ of $\text{pert}(\mu)$. In both cases **NRR.1** and **NRR.2**, the assignment σ of the edges of $\text{pert}(\mu)$ to the two pages is the one determined by the assignments σ_i , that is, $\sigma(e) = \sigma_i(e)$ if $e \in \text{pert}(\mu_i)$. The bottom-to-top ordering π of the vertices of $\text{pert}(\mu)$ is constructed in a different way in the two cases **NRR.1** (see Fig. 16a) and **NRR.2** (see Fig. 16b).

In Case **NRR.1**, π is constructed as $s_\mu, \pi_1, \pi_2, \dots, \pi_k, t_\mu$. In Case **NRR.2**, the children of μ whose 2UBEs have Type $\mathbb{S}\text{-}\langle R, B, R \rangle$ or Type $\mathbb{S}\text{-}\langle N, B, R \rangle$ are consecutive in the left-to-right sequence of the children of μ . Let μ_p, \dots, μ_q be the subsequence formed by these children, with $2 \leq p \leq q \leq k$. For each child $\mu_i, i = p, \dots, q$, by the definition of Type $\mathbb{S}\text{-}\langle R, B, R \rangle$ and Type $\mathbb{S}\text{-}\langle N, B, R \rangle$, there exist two consecutive vertices v' and v'' in π_i such that the portion of the spine between them is visible from the left. Therefore, we can split π_i into two subsequences π'_i and π''_i where $v' \in \pi'_i$ and $v'' \in \pi''_i$. Let μ^* be the possible child of μ whose 2UBE has Type $\mathbb{C}\text{-}\langle R, R, R \rangle$. Ordering π is constructed as $s_\mu, \pi'_p, \pi''_{p+1}, \dots, \pi'_q, \pi^*, \pi''_q, \pi''_{q-1}, \dots, \pi''_p, \pi_{q+1}, \dots, \pi_k, t_\mu$, where $\pi^* = \emptyset$ if μ^* does not exist. In both Case **NRR.1** and Case **NRR.2**, $\langle \pi, \sigma \rangle$ is of a 2UBE of $\text{pert}(\mu)$ of Type $\mathbb{S}\text{-}\langle N, R, R \rangle$. □

To test whether $\text{pert}(\mu)$ admits a 2UBE of Type $\mathbb{S}\text{-}\langle N, R, R \rangle$, we show that deciding if one of Case **NRR.1** or Case **NRR.2** of Lemma 12 applies can be reduced to a flow problem on a network \mathcal{N} with edge demands.

We describe the construction of \mathcal{N} for Case **NRR.2** (refer to Figs. 16b and 16c); the construction for Case **NRR.1** is similar. Network \mathcal{N} is a capacitated flow network where each arc e has a label $[l, u]$, where l and u are the lower and upper capacity of arc e . In particular, \mathcal{N} contains a source s ; a sink t ; a node μ_i for each child μ_i of μ ; a node $\tau \in \{\mathbb{C}\text{-}\langle R, R, R \rangle, \mathbb{S}\text{-}\langle R, R, R \rangle, \mathbb{S}\text{-}\langle R, B, R \rangle, \mathbb{S}\text{-}\langle N, B, R \rangle, \mathbb{S}\text{-}\langle L, B, R \rangle\}$, representing each embedding type used in constructing the sequence of Case **NRR.2**; and two special nodes v_1 and v_2 . Network \mathcal{N} contains the following arcs: (s, μ_i) with label $[1, 1]$, for $i = 1, \dots, k$; an arc (μ_i, τ) with label $[0, 1]$, for each type τ for which $\text{pert}(\mu_i)$ admits a 2UBE of Type τ ; an arc $(\mathbb{C}\text{-}\langle R, R, R \rangle, v_1)$ with label $[0, 1]$; an arc $(\mathbb{S}\text{-}\langle R, R, R \rangle, v_1)$ with label $[0, 1]$; an arc (v_1, t) with label $[1, 2]$; an arc $(\mathbb{S}\text{-}\langle R, B, R \rangle, t)$ with label $[0, \infty]$; an arc $(\mathbb{S}\text{-}\langle R, B, R \rangle, v_2)$ with label $[0, 1]$; an arc $(\mathbb{S}\text{-}\langle L, B, R \rangle, v_2)$ with label $[0, \infty]$; and an arc (v_2, t) with label $[1, \infty]$. The size of \mathcal{N} is linear in the number k of children of μ , because the outdegree of nodes μ_i is bounded by the number of embedding types, s has outdegree k , and the remaining nodes have outdegree at most 1.

Clearly, a sequence corresponding to Case **NRR.2** exists if and only if \mathcal{N} admits a feasible flow, which has value k . Testing the existence of a feasible flow can be reduced in linear time to a max-flow problem [74]. We solve this problem by applying the standard max-flow algorithm by Ford-Fulkerson, which runs in $O(|E(\mathcal{N})| \times f)$, where f is the value of the maximum flow. Since $|E(\mathcal{N})| \in O(k)$ and $f = k$, deciding the existence of a 2UBE of Type $\mathbb{S}\text{-}\langle N, R, R \rangle$ of $\text{pert}(\mu)$ takes $O(k^2)$, which proves Lemma 8 for the case $\mathbb{S}\text{-}\langle N, R, R \rangle$.

7 Conclusion and Open Problems

Our results provide significant advances on the complexity of the k UBE TESTING problem, i.e., deciding whether a directed graph admits a k -page book embedding. For $k \geq 3$ we showed that the problem is NP-hard; for $k = 2$ we described polynomial-time and FPT algorithms for notable families of planar st -graphs. We conclude with some open problems.

- OP1. The digraphs in our NP-completeness proof are not upward planar. Since there are upward planar digraphs that do not admit a 3UBE [72], one can ask whether the k UBE TESTING problem remains NP-complete for $k \geq 3$ and upward planar digraphs. Observe that this question has been solved for $k = 2$, as it is known that 2UBE TESTING is NP-complete for planar st -graphs, which are upward planar [42].
- OP2. A natural extension of our work is to investigate other families of planar digraphs for which a 2UBE always exists or for which polynomial-time testing algorithms can be devised.
- OP3. The complexity of our algorithm to test the existence of a 2UBE for graphs with bounded branchwidth is dominated by the time complexity of computing a sphere-cut decomposition. This motivates the research of more efficient algorithms to compute a sphere-cut decomposition.

Acknowledgements The research of this work began at the Bertinoro Workshop on Graph Drawing 2018. The research has been partially supported by the following projects: (i) “Algoritmi e sistemi di analisi visuale di reti complesse e di grandi dimensioni” – Ricerca di Base 2018, Dipartimento di Ingegneria dell’Università degli Studi di Perugia. (ii) MIUR Project “AHeAD” under PRIN 20174L.F3T8, (iii) MIUR Project “MODE” under PRIN 20157EFM5C, (iv) MIUR-DAAD JMP N° 34120, (v) H2020-MSCA-RISE project 734922 – “CONNECT”, (vi) Roma Tre University Azione 4 Project “GeoView”.

Funding Open access funding provided by Università degli Studi di Perugia within the CRUI-CARE Agreement.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A P-nodes: Variable Embedding - Missing Cases from Sub-section 6.5.2

In this section, we give necessary and sufficient conditions under which a P-node admits one of the relevant embedding types $\mathfrak{S}\text{-}\langle L, B, R \rangle$, $\mathfrak{S}\text{-}\langle R, B, R \rangle$, $\mathfrak{S}\text{-}\langle R, R, R \rangle$, $\mathfrak{S}\text{-}\langle N, B, R \rangle$, $\mathfrak{S}\text{-}\langle R, B, N \rangle$, $\mathfrak{S}\text{-}\langle N, B, N \rangle$, $\mathfrak{S}\text{-}\langle R, R, N \rangle$, $\mathfrak{S}\text{-}\langle N, R, N \rangle$,

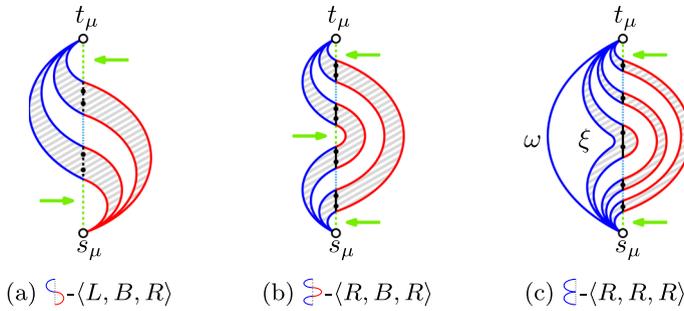


Fig. 17 Constructions for P-nodes of Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$, $\mathfrak{S}\text{-}\langle R, B, R \rangle$, and $\mathfrak{E}\text{-}\langle R, R, R \rangle$

and $\mathfrak{S}\text{-}\langle N, N, N \rangle$. The conditions for $\mathfrak{S}\text{-}\langle N, R, R \rangle$ are discussed in Subsection 6.5.2. The mentioned embedding types are enclosed by a solid polygon in Fig. 12.

Conditions for Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$, Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$, and Type $\mathfrak{E}\text{-}\langle R, R, R \rangle$.

Lemma 13 (Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$) *A P-node μ admits Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$ in the variable embedding setting if and only if all of its children admit Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$.*

Proof For the necessity, let γ be any 2UBE of $\text{pert}(\mu)$ of Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$. Let μ^* be any child of μ and let γ^* be the 2UBE of $\text{pert}(\mu^*)$ obtained by restricting γ to $\text{pert}(\mu^*)$. Since γ is of Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$, there exist in γ^* : (i) a portion of the spine incident to s_μ and between s_μ and t_μ that is visible from the left page, and (ii) a portion of the spine incident to t_μ and between s_μ and t_μ that is visible from the right page. Thus, γ^* has Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$, i.e., μ^* admits Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$. For the sufficiency, we order the children of μ arbitrarily; see Fig. 17a for an example. Clearly, the resulting 2UBE of $\text{pert}(\mu)$ is of Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$. □

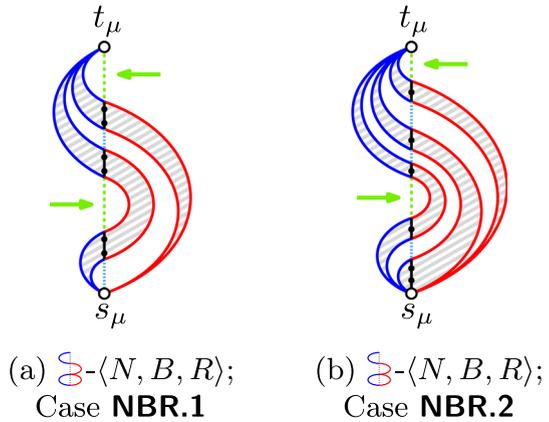
The next lemma is proved similarly to Lemma 13; see Fig. 17b for an example.

Lemma 14 (Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$) *A P-node μ admits Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$ in the variable embedding setting if and only if all of its children admit Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$.*

Lemma 15 (Type $\mathfrak{E}\text{-}\langle R, R, R \rangle$) *A P-node μ admits Type $\mathfrak{E}\text{-}\langle R, R, R \rangle$ in the variable embedding setting if and only if it is possible to partition the children of μ into two sets as follows: The first set consists either of a Q-node child ω or of a child ξ that admits Type $\mathfrak{E}\text{-}\langle R, R, R \rangle$, or both; the second set consists of any number, possibly zero, of children admitting Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$.*

Proof The necessity is easy to see. In fact, for any child μ_i of μ , the restriction of a 2UBE of $\text{pert}(\mu)$ of Type $\mathfrak{E}\text{-}\langle R, R, R \rangle$ to $\text{pert}(\mu_i)$ is either of Type $\mathfrak{E}\text{-}\langle R, R, R \rangle$ or of Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$. Also, a 2UBE of Type $\mathfrak{E}\text{-}\langle R, R, R \rangle$ for $\text{pert}(\mu)$ is incompatible with more than one 2UBE of Type $\mathfrak{E}\text{-}\langle R, R, R \rangle$ for its children, unless one of them is a single edge. If no child of μ is such that the 2UBE of $\text{pert}(\mu)$ is of Type $\mathfrak{E}\text{-}\langle R, R, R \rangle$, then the 2UBE of $\text{pert}(\mu)$ would have Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$.

Fig. 18 Constructions for P-nodes of Type $\mathfrak{S}\text{-}\langle N, B, R \rangle$



About the sufficiency: we order the children of μ so that from left to right we have: the Q-node ω , if any, with Type $\mathfrak{C}\text{-}\langle R, R, R \rangle$, followed by ξ , if any, followed by the remaining children of μ in any order; see Fig. 17c. □

Conditions for Type $\mathfrak{S}\text{-}\langle N, B, R \rangle$. Let μ be a P-node, we define two cases:

NBR.1 All children of μ admit either Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$ or Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$, and there exist at least a child that admits Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$ and a child that admits Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$.

NBR.2 There exists a child of μ that admits Type $\mathfrak{S}\text{-}\langle N, B, R \rangle$ and all other children admit either Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$ or Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$.

Lemma 16 (Type $\mathfrak{S}\text{-}\langle N, B, R \rangle$) A P-node μ admits Type $\mathfrak{S}\text{-}\langle N, B, R \rangle$ in the variable embedding setting if and only if Case **NBR.1** or **NBR.2** holds.

Proof It is easy to see that these conditions are necessary. In fact, for any child μ_i of μ , the restriction of a 2UBE of $\text{pert}(\mu)$ of Type $\mathfrak{S}\text{-}\langle N, B, R \rangle$ to $\text{pert}(\mu_i)$ is clearly of Type $\langle X, B, R \rangle$, with $X \in \{R, L, N\}$, since a 2UBE of $\text{pert}(\mu)$ of type $\langle S, B, T \rangle$ implies that the embeddings of the 2UBEs of the pertinent graphs of all children of μ must be of type $\langle S', B, T' \rangle$, with $S, S', T, T' \in \{R, L, N\}$.

For the sufficiency, both in Case **NBR.1** and in Case **NBR.2** we order the children from left to right in such a way that children with 2UBEs of Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$ precede all children with 2UBEs of Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$; see Figs.18a,18b. In Case **NBR.2**, the child with 2UBE of Type $\mathfrak{S}\text{-}\langle N, B, R \rangle$ is placed between the last child with 2UBE of Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$, if any, and the first child with 2UBE of Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$, if any; see Fig. 18b. □

Conditions for Type $\mathfrak{P}\text{-}\langle R, B, N \rangle$. Let μ be a P-node, we define two cases:

RBN.1 All children of μ admit either Type $\mathfrak{P}\text{-}\langle R, B, L \rangle$ or Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$, and there exist at least a child that admits Type $\mathfrak{P}\text{-}\langle R, B, L \rangle$ and a child that admits Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$.

RBN.2 There exists a child of μ that admits Type $\mathfrak{S}\text{-}\langle R, B, N \rangle$ and all other children admit either Type $\mathfrak{P}\text{-}\langle R, B, L \rangle$ or Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$.

Lemma 17 (Type $\mathfrak{S}\text{-}\langle R, B, N \rangle$) *A P-node μ admits Type $\mathfrak{S}\text{-}\langle R, B, N \rangle$ in the variable embedding setting if and only if Case **RBN.1** or **RBN.2** holds.*

Proof Cases **RBN.1** and **RBN.2** are symmetric to Cases **NBR.1** and **NBR.2**, i.e., they can be obtained from Cases **NBR.1** and **NBR.2** after a vertical flip of the realizable embedding types of the children of μ . Thus, the necessity and sufficiency proofs are symmetric to those presented in Lemma 16. □

Conditions for Type $\mathfrak{S}\text{-}\langle N, B, N \rangle$. Let μ be a P-node. We define six cases:

NBN.1 There exists a child of μ that admits Type $\mathfrak{S}\text{-}\langle N, B, N \rangle$ and the other children admit either Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$ or Type $\mathfrak{S}\text{-}\langle L, B, L \rangle$.

NBN.2 It is possible to partition the children of μ into three sets as follows. The first set consists of at least one child that admits Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$ or Type $\mathfrak{S}\text{-}\langle N, B, R \rangle$, with at most one child that admits only Type $\mathfrak{S}\text{-}\langle N, B, R \rangle$. The second set consists of any number, even zero, of children that admit Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$. The third set consists of at least one child that admits Type $\mathfrak{S}\text{-}\langle L, B, N \rangle$ or Type $\mathfrak{P}\text{-}\langle L, B, L \rangle$, with at most one child that admits only Type $\mathfrak{S}\text{-}\langle L, B, N \rangle$.

NBN.3 This case is obtained from Case **NBN.2** by considering the types admitted by its children after a vertical flip.

NBN.4–NBN.6 These cases are obtained from Cases **NBN.1–NBN.3**, respectively, by considering the horizontally-mirrored embedding types for the children of μ .

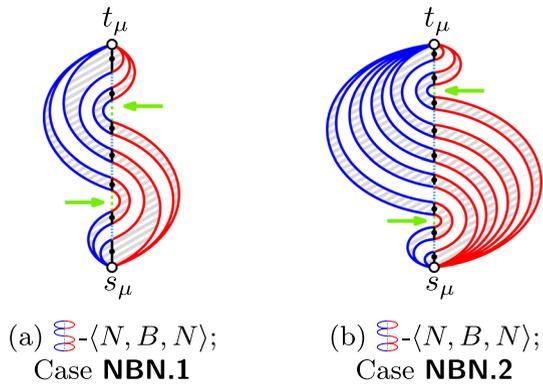
Lemma 18 (Type $\mathfrak{S}\text{-}\langle N, B, N \rangle$) *A P-node μ admits Type $\mathfrak{S}\text{-}\langle N, B, N \rangle$ in the variable embedding setting if and only if at least one of Cases **NBN.1–NBN.6** holds.*

Proof Since Type $\mathfrak{S}\text{-}\langle N, B, N \rangle$ is self-symmetric, we have that, for each case **NBN.1–NBN.3**, we can obtain a symmetric one (i.e., cases **NBN.4–NBN.6**) by reversing the left-to-right sequence of the children in the construction and by taking, for each child, the horizontally-mirrored embedding type. Thus, for each pair of symmetric cases, we only describe one of them.

For the sufficiency, we consider the three cases separately. In Case **NBN.1** we order the children of μ so that from left to right we have the children that admit Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$ (if any), in any order, followed by the unique child that admits Type $\mathfrak{S}\text{-}\langle N, B, N \rangle$, followed by the children that admit Type $\mathfrak{S}\text{-}\langle L, B, L \rangle$ (if any), in any order; see Fig. 19a. In Case **NBN.2** we order the children of μ so that from left to right we have the possible children that admit Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$ (if any), in any order, followed by the possible child that admits only Type $\mathfrak{S}\text{-}\langle N, B, R \rangle$ (if any), followed by the possible children that admit Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$ (if any), in any order, followed by the possible child that admits only Type $\mathfrak{S}\text{-}\langle L, B, N \rangle$ (if any), followed by the possible children that admit Type $\mathfrak{P}\text{-}\langle L, B, L \rangle$ (if any), in any order (see Fig. 19b). In Case **NBN.3** we proceed as in Case **NBN.2** by considering the types after a vertical flip. Clearly, these orderings allow for a 2UBE of $\text{pert}(\mu)$ of Type $\mathfrak{S}\text{-}\langle N, B, N \rangle$.

The necessity can be argued by considering that in any 2UBE γ of $\text{pert}(\mu)$ of Type $\mathfrak{S}\text{-}\langle N, B, N \rangle$ the following hold: at most one of the restrictions of γ to the children

Fig. 19 Constructions for P-nodes of Type $\mathfrak{S}\text{-}\langle N, B, N \rangle$



of μ is of Type $\langle N, \cdot, \cdot \rangle$; at most one of the restriction of γ to the children of μ is of Type $\langle \cdot, \cdot, N \rangle$; and all the restrictions of γ to the children of μ are of Type $\langle \cdot, B, \cdot \rangle$. \square

Conditions for Type $\mathfrak{P}\text{-}\langle R, R, N \rangle$. Let μ be a P-node. We define two cases:

RRN.1 The children of μ can be partitioned into two sets: The first set consists either of a Q-node child or of a child that admits Type $\mathfrak{P}\text{-}\langle R, R, N \rangle$, or both. The second set consists of any number, even zero, of children that admit Type $\mathfrak{P}\text{-}\langle R, B, L \rangle$.

RRN.2 The children of μ can be partitioned into three sets: The first set consists either of a Q-node child or of a non-Q-node child that admits Type $\mathfrak{S}\text{-}\langle R, R, R \rangle$, or both. The second part consists of any number, even zero, of children that admit Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$. The third set consists of at least one child that admits Type $\mathfrak{S}\text{-}\langle R, B, N \rangle$ or Type $\mathfrak{P}\text{-}\langle R, B, L \rangle$, with at most one child that admits only Type $\mathfrak{S}\text{-}\langle R, B, N \rangle$.

Lemma 19 (Type $\mathfrak{P}\text{-}\langle R, R, N \rangle$) A P-node μ admits Type $\mathfrak{P}\text{-}\langle R, R, N \rangle$ in the variable embedding setting if and only if Case **RRN.1** or Case **RRN.2** holds.

Proof Observe that Cases **RRN.1** and **RRN.2** are symmetric to Cases **NRR.1** and **NRR.2** (see Lemma 12), i.e., they can be obtained from Cases **NRR.1** and **NRR.2** after a vertical flip of the types of the children. The necessity and sufficiency proofs are symmetric to those presented in Lemma 12. \square

Conditions for Type $\mathfrak{S}\text{-}\langle N, R, N \rangle$. Let μ be a P-node, we define six cases:

NRN.1 There exists a child that admits Type $\mathfrak{S}\text{-}\langle N, R, N \rangle$ and the other non-Q-node children admit Type $\mathfrak{P}\text{-}\langle L, B, L \rangle$.

NRN.2 It is possible to partition the children of μ into three sets as follows. The first set consists either of a Q-node child or of a child that admits Type $\mathfrak{P}\text{-}\langle R, R, N \rangle$, or both. The second set consists of any number, even zero, of children that admit Type $\mathfrak{P}\text{-}\langle R, B, L \rangle$. The third set consists of at least one child that admits Type $\mathfrak{S}\text{-}\langle N, B, L \rangle$ or Type $\mathfrak{P}\text{-}\langle L, B, L \rangle$, with at most one child that admits only Type $\mathfrak{S}\text{-}\langle N, B, L \rangle$.

NRN.3 This case is obtained from Case **NRN.2** by considering the types of the children of μ after a vertical flip.

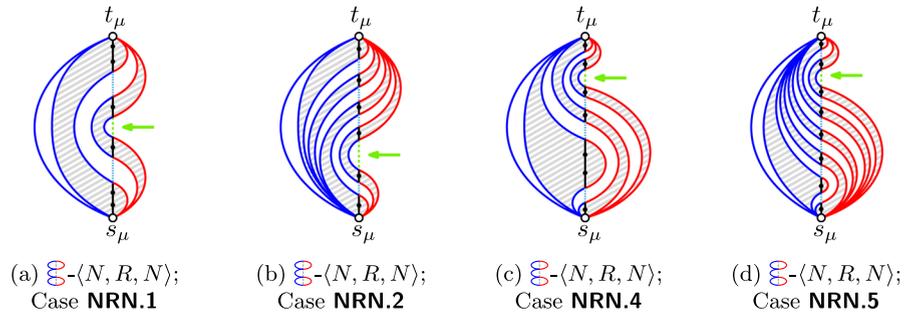


Fig. 20 Constructions for P-nodes of Type $\mathfrak{S}\text{-}\langle N, R, N \rangle$

NRN.4 It is possible to partition the children of μ into two sets as follows. The first set consists either of a Q-node child or of a non-Q-node child that admits Type $\mathfrak{S}\text{-}\langle R, R, R \rangle$, or both. The second set consists of at least one child that admits Type $\mathfrak{S}\text{-}\langle N, B, N \rangle$ or Type $\mathfrak{S}\text{-}\langle L, B, L \rangle$, with at most one child that admits only Type $\mathfrak{S}\text{-}\langle N, B, N \rangle$.

NRN.5 It is possible to partition the children of μ into three sets as follows. The first set consists either of a Q-node child or of a non-Q-node child that admits Type $\mathfrak{S}\text{-}\langle R, R, R \rangle$, or both. The second set consists of any number, even zero, of children that admit Type $\mathfrak{S}\text{-}\langle N, B, R \rangle$ or Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$, with at most one child that admits only Type $\mathfrak{S}\text{-}\langle N, B, R \rangle$. The third set consists of at least one child that admits Type $\mathfrak{S}\text{-}\langle L, B, N \rangle$ or Type $\mathfrak{S}\text{-}\langle L, B, L \rangle$, with at most one child that admits only Type $\mathfrak{S}\text{-}\langle L, B, N \rangle$.

NRN.6 This case is obtained from Case **NRN.5** by considering the types of the children of μ after a vertical flip.

Lemma 20 (Type $\mathfrak{S}\text{-}\langle N, R, N \rangle$) A P-node μ admits Type $\mathfrak{S}\text{-}\langle N, R, N \rangle$ in the variable embedding setting if and only if one of Cases **NRN.1–NRN.6** holds.

Proof For the sufficiency, we consider the six cases separately. In Case **NRN.1** we order the children of μ so that from left to right we have the Q-node child (if any), with Type $\mathfrak{S}\text{-}\langle R, R, R \rangle$, followed by the unique Type $\mathfrak{S}\text{-}\langle N, R, N \rangle$ child, followed by the Type $\mathfrak{S}\text{-}\langle L, B, L \rangle$ children (if any), in any order; see Fig. 20a.

In Case **NRN.2** we order the children of μ so that from left to right we have the Q-node child (if any), with Type $\mathfrak{S}\text{-}\langle R, R, R \rangle$, followed by the Type $\mathfrak{S}\text{-}\langle R, R, N \rangle$ child (if any), followed by the Type $\mathfrak{S}\text{-}\langle R, B, L \rangle$ children (if any), in any order, followed by the unique Type $\mathfrak{S}\text{-}\langle N, B, L \rangle$ child (if any), followed by the Type $\mathfrak{S}\text{-}\langle L, B, L \rangle$ children (if any), in any order; see 20b. In Case **NRN.3** we proceed as in Case **NRN.2** by considering the types after a vertical flip. In Case **NRN.4** we order the children of μ so that from left to right we have the Q-node child (if any), with Type $\mathfrak{S}\text{-}\langle R, R, R \rangle$, followed by the Type $\mathfrak{S}\text{-}\langle R, R, R \rangle$ non-Q-node child (if any), followed by the unique Type $\mathfrak{S}\text{-}\langle N, B, N \rangle$ child (if any), followed by the Type $\mathfrak{S}\text{-}\langle L, B, L \rangle$ children (if any), in any order; see Fig. 20c. In Case **NRN.5** we order the children of μ so that from left to right we have the Q-node child (if any), with Type $\mathfrak{S}\text{-}\langle R, R, R \rangle$, followed by the Type $\mathfrak{S}\text{-}\langle R, R, R \rangle$ non-Q-node child (if any), followed by the unique Type

$\mathfrak{B}\text{-}\langle N, B, R \rangle$ child (if any), followed by the Type $\mathfrak{S}\text{-}\langle L, B, R \rangle$ children (if any), in any order, followed by the unique Type $\mathfrak{S}\text{-}\langle L, B, N \rangle$ child (if any), followed by the Type $\mathfrak{S}\text{-}\langle L, B, L \rangle$ children (if any), in any order; see Fig. 20d. In Case **NRN.6** we proceed as in Case **NRN.5** by considering the types after a vertical flip.

The necessity can be argued by considering that in any 2UBE γ of $\text{pert}(\mu)$ of Type $\mathfrak{B}\text{-}\langle N, R, N \rangle$ we have the following: at most one of the restrictions of γ to the children of μ is of Type $\langle N, \cdot, \cdot \rangle$; at most one of the restrictions of γ to the children of μ is of Type $\langle \cdot, \cdot, N \rangle$; and that all the restrictions of γ on the children of μ must be of Type $\langle \cdot, R, \cdot \rangle$ or of Type $\langle \cdot, B, \cdot \rangle$. However, not all the restrictions of γ to the children of μ may be of Type $\langle \cdot, B, \cdot \rangle$ as otherwise the resulting embeddings would be of Type $\langle \cdot, B, \cdot \rangle$. \square

Conditions for Type $\mathfrak{B}\text{-}\langle N, N, N \rangle$. Let μ be a P-node. In order to describe when μ is of Type $\mathfrak{B}\text{-}\langle N, N, N \rangle$, we need to define twenty-six cases. We assume that Case **NNN.i** applies only if Cases **NNN.k**, with $1 \leq k \leq i - 1$, do not apply.

NNN.1 There exist two children, one of which is a Q-node and the other admits Type $\mathfrak{B}\text{-}\langle N, N, N \rangle$.

NNN.2 It is possible to partition the children of μ into three sets as follows. The first set consists of a child that admits Type $\mathfrak{B}\text{-}\langle N, R, N \rangle$; the second set consists either of a Q-node child or of a non-Q-node child that admits Type $\mathfrak{S}\text{-}\langle R, R, R \rangle$, or both. The third set consists of any number, even zero, of children that admit Type $\mathfrak{S}\text{-}\langle L, B, L \rangle$.

NNN.3 It is possible to partition the children of μ into five sets as follows. The first set consists either of a Q-node child or of a non-Q-node child that admits Type $\mathfrak{S}\text{-}\langle R, R, R \rangle$, or both. The second set consists of any number of children that admit Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$. The third set consists of a child that admits Type $\mathfrak{B}\text{-}\langle N, B, N \rangle$. The fourth set consists of any number of children that admit Type $\mathfrak{S}\text{-}\langle L, B, L \rangle$. The fifth set consists either of a Q-node child or of a non-Q-node child that admits Type $\mathfrak{B}\text{-}\langle L, L, L \rangle$, or both.

NNN.4 It is possible to partition the children of μ into four sets as follows. The first set consists of one child that admits Type $\mathfrak{P}\text{-}\langle R, R, N \rangle$. The second set consists of any number of children that admit Type $\mathfrak{P}\text{-}\langle R, B, L \rangle$ or Type $\mathfrak{S}\text{-}\langle L, B, L \rangle$. The third set consists of at most one child that admits only Type $\mathfrak{B}\text{-}\langle N, B, L \rangle$. The fourth set consists either of a Q-node child or of a non-Q-node child that admits Type $\mathfrak{B}\text{-}\langle L, L, L \rangle$, or both.

NNN.5 This case is obtained from Case **NNN.4** by considering the types of the children of μ after a vertical flip.

NNN.6 There exists a child that admits Type $\mathfrak{P}\text{-}\langle R, R, N \rangle$, any number of children that admit Type $\mathfrak{P}\text{-}\langle R, B, L \rangle$, at most one child that admits only Type $\mathfrak{B}\text{-}\langle N, L, L \rangle$, and at least one child that is either a Q-node or admits Type $\mathfrak{B}\text{-}\langle N, L, L \rangle$.

NNN.7 This case is obtained from Case **NNN.6** by considering the types of the children of μ after a vertical flip.

NNN.8 There exists a child that admits Type $\mathfrak{S}\text{-}\langle R, B, N \rangle$, at most one child that admits only Type $\mathfrak{B}\text{-}\langle N, B, L \rangle$, any number of children that admit Type $\mathfrak{S}\text{-}\langle R, B, R \rangle$, Type $\mathfrak{P}\text{-}\langle R, B, L \rangle$, or Type $\mathfrak{S}\text{-}\langle L, B, L \rangle$. Furthermore, there exist either a non-Q-node child

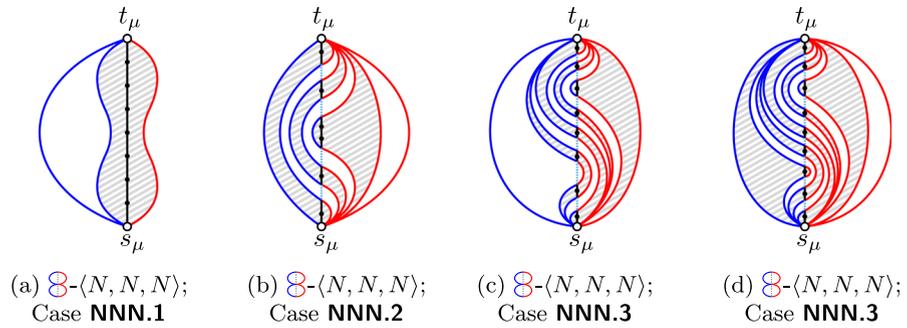


Fig. 21 Constructions for P-nodes of Type \mathcal{B} - $\langle N, N, N \rangle$. Subfigures (c) and (d) show the constructions for Case **NNN.3** when there exists no Type \mathcal{E} - $\langle R, R, R \rangle$ non-Q-node child and when there exists such a child, respectively

that admits Type \mathcal{E} - $\langle R, R, R \rangle$ or a non-Q-node child that admits Type \mathcal{B} - $\langle L, L, L \rangle$, or both; if only one of such children exists then μ has also a Q-node child.

NNN.9 This case is obtained from Case **NNN.8** by considering the types of the children of μ after a vertical flip.

NNN.10 There exists a child that admits Type \mathcal{B} - $\langle R, B, N \rangle$, and any number of children that admit Type \mathcal{B} - $\langle R, B, R \rangle$ or Type \mathcal{P} - $\langle R, B, L \rangle$. Furthermore, there exist either a non-Q-node child that admits Type \mathcal{E} - $\langle R, R, R \rangle$ or a child that admits Type \mathcal{B} - $\langle N, L, L \rangle$, or both; if only one of such children exists then μ has also a Q-node child.

NNN.11 This case is obtained from Case **NNN.10** by considering the types of the children of μ after a vertical flip.

NNN.12 There exists a non-Q-node child that admits Type \mathcal{E} - $\langle R, R, R \rangle$, any number of children that admit Type \mathcal{B} - $\langle R, B, R \rangle$ or Type \mathcal{S} - $\langle L, B, R \rangle$, and either a Q-node child or a non-Q-node child that admits Type \mathcal{B} - $\langle L, L, L \rangle$, or both.

NNN.13 This case is obtained from Case **NNN.12** by considering the types of the children of μ after a vertical flip.

NNN.14–NNN.26 These cases are obtained from Cases **NNN.1–NNN.13**, respectively, by considering the types of the children of μ after a horizontal mirroring.

Lemma 21 (Type \mathcal{B} - $\langle N, N, N \rangle$) *A P-node μ admits Type \mathcal{B} - $\langle N, N, N \rangle$ in the variable embedding setting if and only if one of Cases **NNN.1–NNN.26** holds.*

Proof Since Type \mathcal{B} - $\langle N, N, N \rangle$ is self-symmetric, we only need to consider the Cases **NNN.1–NNN.13**; the Cases **NNN.14–NNN.26** being analogous. In Case **NNN.1** we order the children of μ so that from left to right we have the Q-node child, with Type \mathcal{Q} - $\langle R, R, R \rangle$, followed by the Type \mathcal{B} - $\langle N, N, N \rangle$ child; see Fig. 21a. In Case **NNN.2** we order the children of μ so that from left to right we have the Type \mathcal{B} - $\langle N, R, N \rangle$ child, followed by the Type \mathcal{B} - $\langle L, B, L \rangle$ children (if any), in any order, followed by the Type \mathcal{B} - $\langle L, L, L \rangle$ non-Q-node child (if any), followed by Q-node child (if any), with Type \mathcal{D} - $\langle L, L, L \rangle$; see Fig 21b. In Case **NNN.3** we proceed as follows. If there

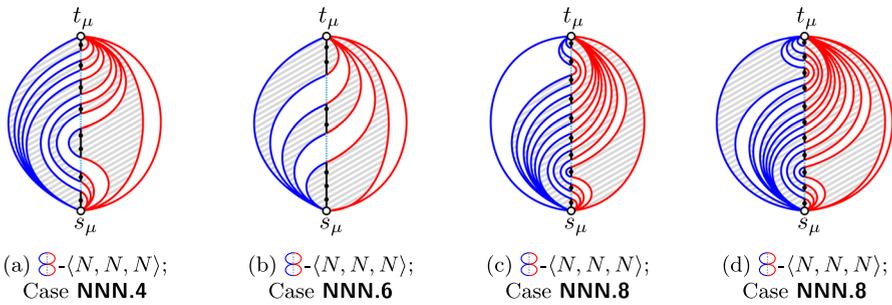


Fig. 22 Constructions for P-nodes of Type $\mathcal{S}\text{-}\langle N, N, N \rangle$. Subfigures (c) and (d) show the constructions for Case **NNN.8** when there exists no Type $\mathcal{S}\text{-}\langle R, R, R \rangle$ non-Q-node child and when there exists such a child, respectively

exists no Type $\mathcal{S}\text{-}\langle R, R, R \rangle$ non-Q-node child, we order the children of μ so that from left to right we have the Q-node child (which exists by the conditions of the case), with Type $\mathcal{C}\text{-}\langle R, R, R \rangle$, followed by the Type $\mathcal{S}\text{-}\langle R, B, R \rangle$ children (if any), in any order, followed by the Type $\mathcal{S}\text{-}\langle N, B, N \rangle$ child, followed by the Type $\mathcal{S}\text{-}\langle L, B, L \rangle$ children (if any), in any order, followed by the Type $\mathcal{B}\text{-}\langle L, L, L \rangle$ non-Q-node child (which exists by the conditions of the case). See Fig. 21c. Otherwise (there exists one Type $\mathcal{S}\text{-}\langle R, R, R \rangle$ non-Q-node child), we order the children of μ so that from left to right we have the Type $\mathcal{S}\text{-}\langle R, R, R \rangle$ non-Q-node child, followed by the Type $\mathcal{S}\text{-}\langle R, B, R \rangle$ children (if any), in any order, followed by the Type $\mathcal{S}\text{-}\langle N, B, N \rangle$ child, followed by the Type $\mathcal{S}\text{-}\langle L, B, L \rangle$ children (if any), in any order, followed by the Type $\mathcal{B}\text{-}\langle L, L, L \rangle$ non-Q-node child (if any), followed by the Type $\mathcal{D}\text{-}\langle L, L, L \rangle$ Q-node child, if any. See Fig. 21d. In Case **NNN.4** we order the children of μ so that from left to right we have the Type $\mathcal{P}\text{-}\langle R, R, N \rangle$ child, followed by the Type $\mathcal{P}\text{-}\langle R, B, L \rangle$ children (if any), in any order, followed by the Type $\mathcal{S}\text{-}\langle N, B, L \rangle$ child (if any), followed by the Type $\mathcal{S}\text{-}\langle L, B, L \rangle$ children (if any), in any order, followed by the Type $\mathcal{B}\text{-}\langle L, L, L \rangle$ non-Q-node child (if any), followed by the Type $\mathcal{D}\text{-}\langle L, L, L \rangle$ Q-node child (if any). See Fig. 22a. In Case **NNN.5** we proceed as in Case **NNN.4** by considering the types after a vertical flip. In Case **NNN.6** we order the children of μ so that from left to right we have the Type $\mathcal{P}\text{-}\langle R, R, N \rangle$ child, followed by any number of Type $\mathcal{P}\text{-}\langle R, B, L \rangle$ children, followed by the child that admits only Type $\mathcal{B}\text{-}\langle N, L, L \rangle$ (if any), followed by the Q-node child (if any), with Type $\mathcal{D}\text{-}\langle L, L, L \rangle$; see Fig. 22b. In Case **NNN.7** we proceed as in Case **NNN.6** by considering the types after a vertical flip. In Case **NNN.8** we proceed as follows. If there exists no Type $\mathcal{S}\text{-}\langle R, R, R \rangle$ non-Q-node child, we order the children of μ so that from left to right we have the Q-node child (which exists by the conditions of the case), with Type $\mathcal{C}\text{-}\langle R, R, R \rangle$, followed by the Type $\mathcal{S}\text{-}\langle R, B, R \rangle$ children (if any), in any order, followed by the Type $\mathcal{S}\text{-}\langle R, B, N \rangle$ child, followed by the Type $\mathcal{P}\text{-}\langle R, B, L \rangle$ children (if any), in any order, followed by the Type $\mathcal{S}\text{-}\langle N, B, L \rangle$ child (if any) followed by the Type $\mathcal{S}\text{-}\langle L, B, L \rangle$ children, if any, in any order, followed by the Type $\mathcal{B}\text{-}\langle L, L, L \rangle$ non-Q-node child (which exists by the conditions of the case). See Fig. 22c. Otherwise (there exists one Type $\mathcal{S}\text{-}\langle R, R, R \rangle$ non-Q-node child), we order the children of μ so that from left to right we have the Type $\mathcal{S}\text{-}\langle R, R, R \rangle$ non-Q-node child, followed by the Type $\mathcal{S}\text{-}\langle R, B, R \rangle$ children (if

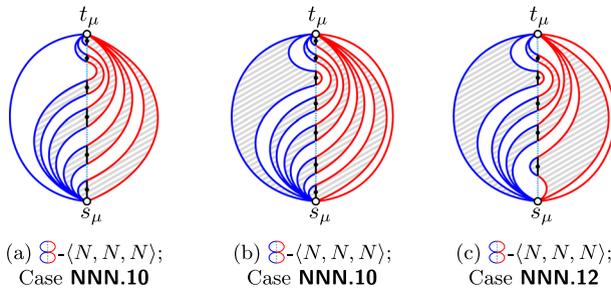


Fig. 23 Constructions for P-nodes of Type $\mathfrak{B}\text{-}\langle N, N, N \rangle$. Subfigures (a) and (b) show the constructions for Case **NNN.10** when there exists no Type $\mathfrak{E}\text{-}\langle R, R, R \rangle$ non-Q-node child and when there exists such a child, respectively

any), in any order, followed by the Type $\mathfrak{B}\text{-}\langle R, B, N \rangle$ child, followed by the Type $\mathfrak{P}\text{-}\langle R, B, L \rangle$ children (if any), in any order, followed by the Type $\mathfrak{B}\text{-}\langle N, B, L \rangle$ child (if any), followed by the Type $\mathfrak{B}\text{-}\langle L, B, L \rangle$ children (if any), in any order, followed by the Type $\mathfrak{B}\text{-}\langle L, L, L \rangle$ non-Q-node child (if any), followed by the Q-node child (if any), with Type $\mathfrak{D}\text{-}\langle L, L, L \rangle$; see Fig. 22d. In Case **NNN.9** we proceed as in Case **NNN.8** by considering the types after a vertical flip. In Case **NNN.10** we proceed as follows. If there exists no Type $\mathfrak{E}\text{-}\langle R, R, R \rangle$ non-Q-node child, we order the children of μ so that from left to right we have the Q-node child (which exists by the conditions of the case), with Type $\mathfrak{C}\text{-}\langle R, R, R \rangle$, followed by the Type $\mathfrak{B}\text{-}\langle R, B, R \rangle$ children (if any), in any order, followed by the Type $\mathfrak{B}\text{-}\langle R, B, N \rangle$ child, followed by the Type $\mathfrak{P}\text{-}\langle R, B, L \rangle$ children (if any), in any order, followed by the Type $\mathfrak{B}\text{-}\langle N, L, L \rangle$ child (which exists by the conditions of the case). See Fig. 23a. Otherwise (there exists one Type $\mathfrak{E}\text{-}\langle R, R, R \rangle$ non-Q-node child), we order the children of μ so that from left to right we have the Type $\mathfrak{E}\text{-}\langle R, R, R \rangle$ non-Q-node child, followed by the Type $\mathfrak{B}\text{-}\langle R, B, R \rangle$ children (if any), in any order, followed by the Type $\mathfrak{B}\text{-}\langle R, B, N \rangle$ child, followed by the Type $\mathfrak{P}\text{-}\langle R, B, L \rangle$ children (if any), in any order, followed by the Type $\mathfrak{B}\text{-}\langle N, L, L \rangle$ child (if any), followed by the Q-node child (if any), with Type $\mathfrak{D}\text{-}\langle L, L, L \rangle$. See Fig. 23b. In Case **NNN.11** we proceed as in Case **NNN.10** by considering the types after a vertical flip. In Case **NNN.12** we order the children of μ so that from left to right we have the non-Q-node Type $\mathfrak{E}\text{-}\langle R, R, R \rangle$ child, followed by the Type $\mathfrak{B}\text{-}\langle R, B, R \rangle$ children (if any), in any order, followed by the Type $\mathfrak{B}\text{-}\langle L, B, R \rangle$ children (if any), in any order, followed by the Type $\mathfrak{B}\text{-}\langle L, L, L \rangle$ non-Q-node child (if any), followed by the Q-node child (if any), with Type $\mathfrak{D}\text{-}\langle L, L, L \rangle$; see Fig. 23c. In Case **NNN.13** we proceed as in Case **NNN.12** by considering the types after a vertical flip.

The necessity can be argued by considering that in any 2UBE γ of $\text{pert}(\mu)$ of Type $\mathfrak{B}\text{-}\langle N, N, N \rangle$ we have the following: at most one of the restriction of γ to the children of μ is of Type $\langle N, \cdot, \cdot \rangle$ and at most one of the restrictions of γ to the children of μ is of Type $\langle \cdot, \cdot, N \rangle$ (notice that these two children may coincide in a single child of Type $\langle N, \cdot, N \rangle$); at most one of the restriction of γ to the non-Q-node children of μ is of Type $\langle L, L, L \rangle$ and at most one of the restriction of γ to the non-Q-node children of μ is of Type $\langle R, R, R \rangle$. □

References

1. Ábrego, B.M., Aichholzer, O., Fernández-Merchant, S., Ramos, P., Salazar, G.: The 2-page crossing number of K_n . In: T. K. Dey and S. Whitesides, editors, Symposium on Computational Geometry 2012, SoCG '12, pp. 397–404. ACM, (2012)
2. Akitaya, H.A., Demaine, E.D., Hesterberg, A., Liu, Q.C.: Upward partitioned book embeddings. In: F. Frati and K. Ma, editors, GD 2017, volume 10692 of LNCS, pp. 210–223. Springer, (2017)
3. Alam, M.J., Bekos, M.A., Dujmovic, V., Gronemann, M., Kaufmann, M., Pupyrev, S.: On dispersable book embeddings. *Theor. Comput. Sci.* **861**, 1–22 (2021)
4. Alhashem, M., Jourdan, G., Zaguia, N.: On the book embedding of ordered sets. *Ars Combin.* **119**, 47–64 (2015)
5. Alzohairi, M., Rival, I.: Series-parallel planar ordered sets have pagenumber two. In: S. C. North, editor, Graph Drawing, GD '96, volume 1190 of LNCS, pp. 11–24. Springer, (1996)
6. Angelini, P., Bekos, M.A., Didimo, W., Grilli, L., Kindermann, P., Mchedlidze, T., Prutkin, R., Symvonis, A., Tappini, A.: Greedy rectilinear drawings. In: Biedl, T., Kerren, A. (eds.) GD 2018, volume 11282 of LNCS. Springer, (2018)
7. Angelini, P., Chaplick, S., Cornelsen, S., Da Lozzo, G.: On upward-planar L-drawings of graphs. In: Szeider, S., Ganian, R., Silva, A. (eds.) 47th International Symposium on Mathematical Foundations of Computer Science, MFCS 2022, August 22–26, 2022, Vienna, Austria, volume 241 of LIPIcs, pages 10:1–10:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, (2022)
8. Angelini, P., Da Lozzo, G., Di Battista, G., Donato, V.D., Kindermann, P., Rote, G., Rutter, I.: Windrose planarity: embedding graphs with direction-constrained edges. *AMS Trans. Algorithms* **14**(4), 54:1–54:24 (2018)
9. Angelini, P., Da Lozzo, G., Di Battista, G., Frati, F.: Strip planarity testing for embedded planar graphs. *Algorithmica* **77**(4), 1022–1059 (2017)
10. Angelini, P., Da Lozzo, G., Di Battista, G., Frati, F., Patrignani, M.: 2-Level quasi-planarity or how caterpillars climb (SPQR)-trees. In: Marx, D. editor, Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms. SODA 2021, Virtual Conference **2021**, 2779–2798 (2021). (**SIAM**, **10** - **13**)
11. Angelini, P., Da Lozzo, G., Di Battista, G., Frati, F., Patrignani, M., Rutter, I.: Intersection-link representations of graphs. *J. Graph Algorithms Appl.* **21**(4), 731–755 (2017)
12. Angelini, P., Da Lozzo, G., Neuwirth, D.: Advancements on SEFE and partitioned book embedding problems. *Theor. Comput. Sci.* **575**, 71–89 (2015)
13. Angelini, P., Di Bartolomeo, M., Di Battista, G.: Implementing a partitioned 2-page book embedding testing algorithm. In: Graph Drawing, volume of LNCS **7704**, 79–89 (2012). (**Springer**)
14. Angelini, P., Di Battista, G., Frati, F., Patrignani, M., Rutter, I.: Testing the simultaneous embeddability of two graphs whose intersection is a biconnected or a connected graph. *J. Dis. Algorithms* **14**, 150–172 (2012)
15. Angelini, P., Eppstein, D., Frati, F., Kaufmann, M., Lazard, S., Mchedlidze, T., Teillaud, M., Wolff, A.: Universal point sets for drawing planar graphs with circular arcs. *J. Graph Algorithms Appl.* **18**(3), 313–324 (2014)
16. Badent, M., Di Giacomo, E., Liotta, G.: Drawing colored graphs on colored points. *Theor. Comput. Sci.* **408**(2–3), 129–142 (2008)
17. Bannister, M.J., Eppstein, D.: Crossing minimization for 1-page and 2-page drawings of graphs with bounded treewidth. In: Duncan, C.A., Symvonis, A. (eds.) GD 2014 **8871**, 210–221 (2014). (**LNCS, Springer**)
18. Bekos, M.A., Bruckdorfer, T., Kaufmann, M., Raftopoulou, C.N.: The book thickness of 1-planar graphs is constant. *Algorithmica* **79**(2), 444–465 (2017)
19. Bekos, M.A., Da Lozzo, G., Frati, F., Gronemann, M., Mchedlidze, T., Raftopoulou, C.N.: Recognizing dags with page-number 2 is NP-complete. *Theor. Comput. Sci.* **946**, 113689 (2023)
20. Bekos, M.A., Da Lozzo, G., Frati, F., Gronemann, M., Mchedlidze, T., Raftopoulou, C.N.: Recognizing dags with page-number 2 is NP-complete. In: Angelini, P., von Hanxleden, R. (eds.) Graph Drawing and Network Visualization, pp. pp. 361–370. Springer International Publishing, Cham (2023)
21. Bekos, M.A., Da Lozzo, G., Griesbach, S., Gronemann, M., Montecchiani, F., Raftopoulou, C.N.: Book embeddings of nonplanar graphs with small faces in few pages. In: Cabello, S., Chen, D.Z., (eds.) 36th International Symposium on Computational Geometry, SoCG 2020, June 23–26, 2020, Zürich, Switzerland, volume 164 of LIPIcs, pages 16:1–16:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, (2020)

22. Bekos, M.A., Di Giacomo, E., Didimo, W., Liotta, G., Montecchiani, F., Raftopoulou, C.N.: Edge partitions of optimal 2-plane and 3-plane graphs. In: Brandstädt, A., Köhler, E., Meer, K., (eds.) Graph-Theoretic Concepts in Computer Science, WG 2018, volume 11159 of LNCS, pages 27–39. Springer, (2018)
23. Bekos, M.A., Gronemann, M., Raftopoulou, C.N.: Two-page book embeddings of 4-planar graphs. *Algorithmica* **75**(1), 158–185 (2016)
24. Bekos, M.A., Kaufmann, M., Klute, F., Pupyrev, S., Raftopoulou, C.N., Ueckerdt, T.: Four pages are indeed necessary for planar graphs. *J. Comput. Geom.* **11**(1), 332–353 (2020)
25. Bernhart, F., Kainen, P.C.: The book thickness of a graph. *J. Combinat. Theory, Series B* **27**(3), 320–331 (1979)
26. Bertolazzi, P., Di Battista, G., Didimo, W.: Quasi-upward planarity. *Algorithmica* **32**(3), 474–506 (2002)
27. Bertolazzi, P., Di Battista, G., Mannino, C., Tamassia, R.: Optimal upward planarity testing of single-source digraphs. *SIAM J. Comput.* **27**(1), 132–169 (1998)
28. Bhole, S., Da Lozzo, G., Montecchiani, F., Nöllenburg, M.: On the upward book thickness problem: combinatorial and complexity results. *Eur. J. Comb.* **110**, 103662 (2023)
29. Biedl, T.C., Shermer, T.C., Whitesides, S., Wismath, S.K.: Bounds for orthogonal 3-d graph drawing. *J. Graph Algorithms Appl.* **3**(4), 63–79 (1999)
30. Binucci, C., Da Lozzo, G., Giacomo, E.D., Didimo, W., Mchedlidze, T., Patrignani, M.: Upward book embeddings of st-graphs. In: Barequet, G., Wang, Y., (eds.) 35th International Symposium on Computational Geometry, SoCG 2019, June 18–21, 2019, Portland, Oregon, USA, volume 129 of LIPIcs, pages 13:1–13:22. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, (2019)
31. Binucci, C., Di Giacomo, E., Hossain, M.I., Liotta, G.: 1-page and 2-page drawings with bounded number of crossings per edge. *Eur. J. Comb.* **68**, 24–37 (2018)
32. Binucci, C., Didimo, W.: Computing quasi-upward planar drawings of mixed graphs. *Comput. J.* **59**(1), 133–150 (2016)
33. Binucci, C., Didimo, W., Patrignani, M.: Upward and quasi-upward planarity testing of embedded mixed graphs. *Theor. Comput. Sci.* **526**, 75–89 (2014)
34. Brandenburg, F.: Upward planar drawings on the standing and the rolling cylinders. *Comput. Geom.* **47**(1), 25–41 (2014)
35. Cardinal, J., Hoffmann, M., Kusters, V., Tóth, C.D., Wettstein, M.: Arc diagrams, flip distances, and hamiltonian triangulations. *Comput. Geom.* **68**, 206–225 (2018)
36. Chaplick, S., Chimani, M., Cornelsen, S., Da Lozzo, G., Nöllenburg, M., Patrignani, M., Tollis, I.G., Wolff, A.: Planar L-drawings of directed graphs. In: Frati, F., Ma, K., (eds.) GD 2017, volume 10692 of LNCS, pages 465–478. Springer, (2017)
37. Chung, F.R.K., Leighton, F.T., Rosenberg, A.L.: Embedding graphs in books: a layout problem with applications to VLSI design. *SIAM J. Algebraic Dis. Methods* **8**(1), 33–58 (1987)
38. Cimikowski, R.J.: An analysis of some linear graph layout heuristics. *J. Heuristics* **12**(3), 143–153 (2006)
39. Da Lozzo, G., Di Battista, G., Frati, F.: Extending upward planar graph drawings. *Comput. Geom.* **91**, 101668 (2020)
40. Da Lozzo, G., Di Battista, G., Frati, F., Patrignani, M., Roselli, V.: Upward planar morphs. *Algorithmica* **82**(10), 2985–3017 (2020)
41. Di Battista, G., Eades, P., Tamassia, R., Tollis, I.G.: *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice Hall PTR, USA (1998)
42. Di Battista, G., Tamassia, R.: Algorithms for plane representations of acyclic digraphs. *Theor. Comput. Sci.* **61**, 175–198 (1988)
43. Di Battista, G., Tamassia, R.: On-line planarity testing. *SIAM J. Comput.* **25**(5), 956–997 (1996)
44. Di Battista, G., Tamassia, R., Tollis, I.G.: Area requirement and symmetry display of planar upward drawings. *Dis. Comput. Geomet.* **7**, 381–401 (1992)
45. Di Giacomo, E., Didimo, W., Liotta, G., Wismath, S.K.: Curve-constrained drawings of planar graphs. *Comput. Geom.* **30**(1), 1–23 (2005)
46. Di Giacomo, E., Didimo, W., Liotta, G., Wismath, S.K.: Book embeddability of series-parallel digraphs. *Algorithmica* **45**(4), 531–547 (2006)
47. Di Giacomo, E., Giordano, F., Liotta, G.: Upward topological book embeddings of dags. *SIAM J. Dis. Math.* **25**(2), 479–489 (2011)

48. Di Giacomo, E., Liotta, G., Trotta, F.: On embedding a graph on two sets of points. *Int. J. Found. Comput. Sci.* **17**(5), 1071–1094 (2006)
49. Di Giacomo, E., Liotta, G., Trotta, F.: Drawing colored graphs with constrained vertex positions and few bends per edge. *Algorithmica* **57**(4), 796–818 (2010)
50. Didimo, W.: Upward graph drawing. In: *Encyclopedia of Algorithms*, pages 2308–2312. Springer, (2016)
51. Didimo, W., Giordano, F., Liotta, G.: Upward spirality and upward planarity testing. *SIAM J. Dis. Math.* **23**(4), 1842–1899 (2009)
52. Dorn, F., Penninx, E., Bodlaender, H.L., Fomin, F.V.: Efficient exact algorithms on planar graphs: exploiting sphere cut decompositions. *Algorithmica* **58**(3), 790–810 (2010)
53. Dujmovic, V., Eppstein, D., Hickingbotham, R., Morin, P., Wood, D.R.: Stack-number is not bounded by queue-number. *CoRR*, abs/2011.04195, (2020)
54. Dujmović, V., Wood, D.R.: On linear layouts of graphs. *Dis. Math. Theor. Comput. Sci.* **6**(2), 339–358 (2004)
55. Dujmović, V., Wood, D.R.: Stacks, queues and tracks: layouts of graph subdivisions. *Dis. Math. Theor. Comput. Sci.* **7**(1), 155–202 (2005)
56. Enomoto, H., Miyauchi, M.S.: Embedding graphs into a three page book with $O(m \log n)$ crossings of edges over the spine. *SIAM J. Dis. Math.* **12**(3), 337–341 (1999)
57. Enomoto, H., Miyauchi, M.S., Ota, K.: Lower bounds for the number of edge-crossings over the spine in a topological book embedding of a graph. *Discret. Appl. Math.* **92**(2–3), 149–155 (1999)
58. Enomoto, H., Nakamigawa, T., Ota, K.: On the pagenumber of complete bipartite graphs. *J. Combinat. Theory, Series B* **71**(1), 111–120 (1997)
59. Everett, H., Lazard, S., Liotta, G., Wismath, S.K.: Universal sets of n points for one-bend drawings of planar graphs with n vertices. *Dis. Comput. Geomet.* **43**(2), 272–288 (2010)
60. Fomin, F.V., Thilikos, D.M.: New upper bounds on the decomposability of planar graphs. *J. Graph Theory* **51**(1), 53–81 (2006)
61. Frati, F., Fulek, R., Ruiz-Vargas, A.J.: On the page number of upward planar directed acyclic graphs. *J. Graph Algorithms and Appl.* **17**(3), 221–244 (2013)
62. Ganley, J.L., Heath, L.S.: The pagenumber of k -trees is $O(k)$. *Discret. Appl. Math.* **109**(3), 215–221 (2001)
63. Garg, A., Tamassia, R.: On the computational complexity of upward and rectilinear planarity testing. *SIAM J. Comput.* **31**(2), 601–625 (2001)
64. Giordano, F., Liotta, G., Mchedlidze, T., Symvonis, A., Whitesides, S.: Computing upward topological book embeddings of upward planar digraphs. *J. Dis. Algorithms* **30**, 45–69 (2015)
65. Gutwenger, C., Mutzel, P.: A linear time implementation of SPQR-trees. In: Marks, J. (ed) *Graph Drawing, GD 2000*, volume 1984 of LNCS, pages 77–90. Springer, (2000)
66. Heath, L., Leighton, F., Rosenberg, A.: Comparing queues and stacks as mechanisms for laying out graphs. *SIAM J. Dis. Math.* **5**(3), 398–412 (1992)
67. Heath, L.S., Pemmaraju, S.V.: Stack and queue layouts of posets. *SIAM J. Dis. Math.* **10**(4), 599–625 (1997)
68. Heath, L.S., Pemmaraju, S.V.: Stack and queue layouts of directed acyclic graphs: Part II. *SIAM J. Comput.* **28**(5), 1588–1626 (1999)
69. Heath, L.S., Pemmaraju, S.V., Trenk, A.N.: Stack and queue layouts of directed acyclic graphs: Part I. *SIAM J. Comput.* **28**(4), 1510–1539 (1999)
70. Hong, S., Nagamochi, H.: Simpler algorithms for testing two-page book embedding of partitioned graphs. *Theor. Comput. Sci.* **725**, 79–98 (2018)
71. Hopcroft, J.E., Tarjan, R.E.: Dividing a graph into triconnected components. *SIAM J. Comput.* **2**(3), 135–158 (1973)
72. Hung, L.T.Q.: A Planar Poset which Requires 4 Pages. PhD thesis, Institute of Computer Science, University of Wrocław, (1989)
73. Jungeblut, P., Merker, L., Ueckerdt, T.: A sublinear bound on the page number of upward planar graphs. In: J. S. Naor and N. Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9–12, 2022*, pages 963–978. SIAM, (2022)
74. Kleinberg, J.M., Tardos, É.: *Algorithm design*. Addison-Wesley, (2006)
75. Löffler, M., Tóth, C.D.: Linear-size universal point sets for one-bend drawings. In: *Graph Drawing*, volume 9411 of LNCS, pages 423–429. Springer, (2015)

76. Malitz, S.M.: Genus g graphs have pagenumber $O(\sqrt{g})$. *J. Algorithms* **17**(1), 85–109 (1994)
77. Malitz, S.M.: Graphs with E edges have pagenumber $O(\sqrt{E})$. *J. Algorithms* **17**(1), 71–84 (1994)
78. Masuda, S., Nakajima, K., Kashiwabara, T., Fujisawa, T.: Crossing minimization in linear embeddings of graphs. *IEEE Trans. Comput.* **39**(1), 124–127 (1990)
79. Mchedlidze, T., Symvonis, A.: Crossing-free acyclic hamiltonian path completion for planar st -digraphs. In: Dong, Y., Du, D., Ibarra, O.H., (eds.) *Algorithms and Computation, ISAAC 2009*, volume 5878 of LNCS, pages 882–891. Springer, (2009)
80. Mchedlidze, T., Symvonis, A.: Unilateral orientation of mixed graphs. In: *SOFSEM 2010*, volume 5901 of LNCS, pages 588–599. Springer, (2010)
81. Mchedlidze, T., Symvonis, A.: Crossing-optimal acyclic HP-completion for outerplanar st -digraphs. *J. Graph Algorithms and Appl.* **15**(3), 373–415 (2011)
82. Nowakowski, R., Parker, A.: Ordered sets, pagenumbers and planarity. *Order* **6**(3), 209–218 (1989)
83. Opatrny, J.: Total ordering problem. *SIAM J. Comput.* **8**(1), 111–114 (1979)
84. Pemmaraju, S.V.: Exploring the Powers of Stacks and Queues via Graph Layouts. PhD thesis, Virginia Polytechnic Institute and State University at Blacksburg, Virginia, (1992)
85. Rextin, A., Healy, P.: Dynamic upward planarity testing of single source embedded digraphs. *Comput. J.* **60**(1), 45–59 (2017)
86. Robertson, N., Seymour, P.D.: Graph minors. x. obstructions to tree-decomposition. *J. Comb. Theory, Ser. B*, **52**(2):153–190, (1991)
87. Seymour, P.D., Thomas, R.: Call routing and the ratcatcher. *Combinatorica* **14**(2), 217–241 (1994)
88. Syslo, M.M.: Bounds to the page number of partially ordered sets. In: Nagl, M., (ed.) *Graph-Theoretic Concepts in Computer Science, WG '89*, volume 411 of LNCS, pp. 181–195. Springer, (1989)
89. Unger, W.: On the k -colouring of circle-graphs. In: Cori, R., Wirsing, M., (eds.) *STACS 88*, volume 294 of LNCS, pp. 61–72. Springer, (1988)
90. Unger, W.: The complexity of colouring circle graphs (extended abstract). In: Finkel, A., Jantzen, M., (eds.) *STACS 92*, volume 577 of LNCS, pp. 389–400. Springer, (1992)
91. Wigderson, A.: The complexity of the Hamiltonian circuit problem for maximal planar graphs. Technical report, 298, EECS Department, Princeton University, (1982)
92. Wood, D.R.: Bounded degree book embeddings and three-dimensional orthogonal graph drawing. In: *Graph Drawing*, volume 2265 of LNCS, pp. 312–327. Springer, (2001)
93. Yannakakis, M.: Embedding planar graphs in four pages. *J. Comput. Syst. Sci.* **38**(1), 36–67 (1989)
94. Yannakakis, M.: Planar graphs that need four pages. *J. Comb. Theory, Ser. B*, **145**:241–263, (2020)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Carla Binucci¹ · Giordano Da Lozzo² · Emilio Di Giacomo¹ ·
Walter Didimo¹  · Tamara Mchedlidze³ · Maurizio Patrignani²

Carla Binucci
carla.binucci@unipg.it

Giordano Da Lozzo
giordano.dallozzo@uniroma3.it

Emilio Di Giacomo
emilio.digiacoemo@unipg.it

Tamara Mchedlidze
t.msentiintze@uu.nl

Maurizio Patrignani
maurizio.patrignani@uniroma3.it

- ¹ Department of Engineering, University of Perugia, Perugia, Italy
- ² Department of Engineering, Roma Tre University, Rome, Italy
- ³ Utrecht University, Utrecht, The Netherlands