Theoretical Computer Science 998 (2024) 114541

Contents lists available at ScienceDirect

Theoretical Computer Science

journal homepage: www.elsevier.com/locate/tcs

Upward planarity testing of biconnected outerplanar DAGs solves partition ☆,☆☆

Fabrizio Frati

Roma Tre University, Rome, Italy

ARTICLE INFO

Communicated by D.-Z. Du

Keywords:
Upward planarity
Graph drawing
Outerplanar graph
Partition

ABSTRACT

We show an $\mathcal{O}(n)$-time reduction from the problem of testing whether a multiset of positive integers can be partitioned into two multisets so that the sum of the integers in each multiset is equal to $n/2$ to the problem of testing whether an n-vertex biconnected outerplanar DAG admits an upward planar drawing. This constitutes the first barrier to the existence of efficient algorithms for testing the upward planarity of DAGs with no large triconnected minor.

We also show a result in the opposite direction. Suppose that partitioning a multiset of positive integers into two multisets so that the sum of the integers in each multiset is $n/2$ can be solved in $f(n)$ time. Let G be an n-vertex biconnected outerplanar DAG and e be an edge incident to the outer face of an outerplanar drawing of G. Then it can be tested in $\mathcal{O}(f(n))$ time whether G admits an upward planar drawing with e on the outer face.

1. Introduction

An *upward planar drawing* of a directed acyclic graph (DAG, for short) maps each vertex to a point of the plane and each edge to a Jordan arc monotonically increasing from the source to the sink of the edge, so that no two edges intersect, except at common end-vertices. Upward planarity is perhaps the most natural extension of planarity to directed graphs. As such, the problem of testing whether a DAG G admits an upward planar drawing has received an enormous attention in the literature. While the problem was early shown to be NP-hard [17], it is known to be tractable in many important cases, among which: (i) if G has a fixed combinatorial embedding, that the required upward planar drawing has to respect [1]; (ii) if G has a single source [2,6,21,22] or, more in general, a bounded number of sources [8]; (iii) if G contains no large triconnected minor [9,13,30]. The parameterized complexity of the problem has also been investigated [7,8,13,18].

This paper focuses on *outerplanar DAGs* – fitting case (iii) in the above list. These are DAGs whose underlying (undirected) graph admits an *outerplanar drawing*, i.e., a planar drawing in which all the vertices are incident to the outer face. An $\mathcal{O}(n^2)$-time algorithm for testing whether an n-vertex outerplanar DAG admits a (not necessarily outerplanar) upward planar drawing is known [30]. Although we do not fully understand the details of the mentioned algorithm [30], an $\mathcal{O}(n^2)$-time upward planarity testing algorithm

☆ This article belongs to Section A: Algorithms, automata, complexity and games, Edited by Paul Spirakis.

☆☆ Research partially supported by PRIN projects no. 2022ME9Z78 "NextGRAAL: Next-generation algorithms for constrained GRAph visuALization" and no. 2022TS4Y3N "EXPAND: scalable algorithms for EXPloratory Analyses of heterogeneous and dynamic Networked Data".

E-mail address: fabrizio.frati@uniroma3.it.

https://doi.org/10.1016/j.tcs.2024.114541
Received 25 July 2023; Received in revised form 23 December 2023; Accepted 27 March 2024

Available online 2 April 2024

is now known [9] for all the n-vertex DAGs with no K_4-minor – this is the class of *directed partial* 2*-trees*, which includes the one of the outerplanar DAGs.

Can we design an upward planarity test for outerplanar DAGs running in linear time? We do not know. However, in this paper, we link the complexity of the problem to the one of a very famous problem in complexity theory. The PARTITION problem takes as input a multiset $\mathcal{S}$ of positive integers and asks whether $\mathcal{S}$ admits a partition into two multisets such that the sum of the elements in each multiset is the same.[1] PARTITION is one of the "six basic NP-complete problems" in the notorious book [16] by Garey and Johnnson. In the same book, it is proved that the problem admits a pseudo-polynomial-time algorithm, that is, an algorithm which runs in polynomial time if each element, or equivalently the sum of the elements, is bounded by a polynomial function of the number of elements. Specifically, there is an algorithm based on dynamic programming that solves PARTITION in $\mathcal{O}(k \cdot n)$ time, if we denote by k the number of elements and by n the sum of the elements in the set.[2] As k might be in $\Omega(n)$, the running time of the algorithm is $\mathcal{O}(n^2)$, when expressed solely as a function of n. The time needed to solve PARTITION is nowadays known to be $\mathcal{O}(n \log^2 n)$, due to a recent algorithm for the more general SUBSET SUM problem [25].

Our main result is the following.

Theorem 1. *Let $\mathcal{S} = \{a_1, \dots, a_k\}$ be an instance of the* PARTITION *problem, where $\sum_{i=1}^{k} a_i = n$. It is possible to construct in $\mathcal{O}(n)$ time an $\mathcal{O}(n)$-vertex biconnected outerplanar DAG G which is upward planar if and only if $\mathcal{S}$ is a positive instance of* PARTITION. *Moreover, a partition of $\mathcal{S}$ into two sets $\mathcal{S}_1$ and $\mathcal{S}_2$ such that $\sum_{a_i \in \mathcal{S}_1} a_i = \sum_{a_i \in \mathcal{S}_2} a_i = n/2$ can be constructed in $\mathcal{O}(n)$ time from an upward plane embedding of G.*

Theorem 1 provides evidence that an $\mathcal{O}(n)$-time algorithm for testing the upward planarity of an n-vertex biconnected outerplanar DAG might be very hard to obtain, as it would imply an $\mathcal{O}(n)$-time algorithm for PARTITION, which would be a ground-breaking result. We also remark that, due to Theorem 1, a super-linear lower bound for PARTITION would imply the same lower bound for testing the upward planarity of a biconnected outerplanar DAG. The authors of this paper are unaware of any graph drawing problem that does not involve real numbers (e.g., a prescribed point set), that is not NP-hard, and whose complexity is super-linear. Hence, this would constitute an important step towards establishing polynomial lower bounds for the complexity of graph drawing problems.

We can also prove a reduction in the opposite direction, as in the following.

Theorem 2. *Suppose that an algorithm exists that solves an instance $\mathcal{S}$ of the* PARTITION *problem in $f(n)$ time, for some function $f(n) \in \Omega(n)$, where n denotes the sum of the elements in $\mathcal{S}$. Then, for any n-vertex biconnected outerplanar DAG G and for any edge e incident to the outer face of an outerplanar drawing of G, it can be tested in $\mathcal{O}(f(n))$ time whether G admits an upward planar drawing in which e is incident to the outer face.*

We make three remarks on Theorem 2. First, together with the mentioned $\mathcal{O}(n \log^2 n)$-time algorithm to solve the PARTITION problem [25], Theorem 2 implies the following.

Theorem 3. *For any n-vertex biconnected outerplanar DAG G and for any edge e incident to the outer face of an outerplanar drawing of G, there exists an algorithm that tests in $\mathcal{O}(n \log^2 n)$ time whether G admits an upward planar drawing in which e is incident to the outer face.*

Second, although its statement does not mention this explicitly, the proof of Theorem 1 shows that, given an instance $\mathcal{S} = \{a_1, \dots, a_k\}$ of PARTITION with $\sum_{i=1}^{k} a_i = n$, it is possible to construct in $\mathcal{O}(n)$ time an $\mathcal{O}(n)$-vertex biconnected outerplanar DAG G such that, *for a certain edge e incident to the outer face of an outerplanar drawing of G*, we have that G admits an upward planar drawing *in which e is incident to the outer face* if and only if $\mathcal{S}$ is a positive instance of PARTITION. This, together with Theorem 2, implies that the PARTITION problem is linear-time equivalent to the problem of testing whether a biconnected outerplanar DAG G with a given edge e incident to the outer face of an outerplanar drawing of G admits an upward planar drawing in which e is incident to the outer face.

Third, the approach of solving an algorithmic graph drawing problem by first constraining a prescribed edge to be incident to the outer face, and by only later getting rid of that constraint, nicely interfaces with the decomposition of a biconnected planar graph into its triconnected components. This decomposition is usually expressed by a tree, namely the SPQR-tree for a biconnected planar graph, or the SPQ-tree for a biconnected partial 2-tree, or the (extended) dual tree for a biconnected outerplanar graph; having a prescribed edge incident to the outer face corresponds to rooting the decomposition tree at some node, which then allows for a bottom-up computation on the tree. This approach has been successfully used for problems related to ours, namely testing the upward planarity of directed partial 2-trees [9] and testing the rectilinear planarity of outerplanar graphs [15]; for these two problems, removing the constraint about the prescribed edge on the outer face is doable without increasing the algorithm's running time, by using a strategy devised in [14].

Related results. Fast approximation and exact algorithms for PARTITION are known [23,24,27], running in time exponential in k [20,26,28,33]. See the PhD thesis by Schreiber [32] for more references.

---

1 In the rest of the paper, we use the term "set" in place of "multiset", implicitly assuming that sets might contain multiple occurrences of the same element.

2 In the literature, the number of elements is denoted by n and their sum by B, however as will be evident soon, our notation is more convenient for this paper.

PARTITION is a special case of the SUBSET SUM problem, whose input consists of a set $\mathcal{S}$ of k positive integers that sum up to n, and of a target integer t. The problem asks whether there exists a subset of $\mathcal{S}$ whose elements sum up to t; note that, when $t = n/2$, this is exactly the PARTITION problem. A plethora of results about the SUBSET SUM problem are known. Here we mention only a few of them. Let k' and m be the number of distinct elements in $\mathcal{S}$ and the value of the maximum element in $\mathcal{S}$, respectively. Pisinger [31] presented an $\mathcal{O}(nm)$-time algorithm for the SUBSET SUM problem. Koiliaris and Xu [25] presented an algorithm running in $\mathcal{O}\left(\min\left\{\sqrt{k'}m\log^{5/2} m, m^{4/3}\log^2 m, n\log n\log(k'\log m)\right\}\right)$ time; note that the last of the three terms in the previous minimum gives an $\mathcal{O}(n\log^2 n)$ upper bound, as $k' \in \mathcal{O}(n)$ and $m \in \mathcal{O}(n)$; this is the previously mentioned best-known upper bound for the PARTITION problem. Bringmann [5] presented a randomized algorithm running in $\mathcal{O}(k + t\log t\log^3(k/\delta)\log k)$ time for computing a set that contains every value $t' \leq t$ with probability larger than $1-\delta$.

The complexity of upward planarity testing is intertwined with the one of *rectilinear planarity testing*, which asks whether a graph admits a planar drawing in which each edge is horizontal or vertical. The reductions proving NP-hardness for rectilinear planarity testing and upward planarity testing are similar [17] and, like upward planarity testing, rectilinear planarity testing is tractable when the graph has a fixed combinatorial embedding [4,34] and when it contains no large triconnected minor [3,11,12]. Our results highlight a possible divergence between upward and rectilinear planarity testing, as the rectilinear planarity of an outerplanar graph can be tested in linear time [15].

2. Preliminaries

A *plane embedding* of a biconnected planar graph G is an equivalence class of planar drawings of G, where two drawings Γ_1 and Γ_2 are equivalent if the clockwise order of the edges incident to each vertex is the same in Γ_1 and Γ_2, and the clockwise order of the vertices along the cycle delimiting the outer face is the same in Γ_1 and Γ_2. We often talk about *faces of a plane embedding*, meaning faces of any planar drawing in that equivalence class. If the drawings in the class are outerplanar, we have an *outerplane embedding*; this is actually unique, up to an inversion of all the orders defining it, for a biconnected outerplanar graph. We denote by $f_{\mathcal{E}}$ the outer face of a plane embedding $\mathcal{E}$.

Let G be a DAG. A *switch* of G is a vertex that is a source or a sink. An upward planar drawing of G determines an assignment of labels to the angles of the corresponding plane embedding, where an *angle* α at a vertex u in a face f of a plane embedding represents an incidence of u on f. In particular, α is *flat* and has label 0 if the edges delimiting it are one incoming into u and one outgoing from u. Otherwise, α is a *switch* angle and it is *small* (with label -1) or *large* (with label 1) depending on whether the (geometric) angle at f representing α is smaller or larger than 180°, respectively. An *upward plane embedding* of G is an equivalence class of upward planar drawings of G, where two drawings are equivalent if they determine the same plane embedding $\mathcal{E}$ for G and the same label assignment for the angles of $\mathcal{E}$. The label assignments that enhance a plane embedding to an upward plane embedding are characterized as follows.

Theorem 4 *([1,13]). Let G be a digraph with plane embedding $\mathcal{E}$, and λ be a label assignment for the angles of $\mathcal{E}$. Then $\mathcal{E}$ and λ define an upward plane embedding of G if and only if the following hold:*

(UP0) *If α is a switch angle then α is small or large, otherwise it is flat.*
(UP1) *If v is a switch vertex, the number of small, flat and large angles incident to v is equal to $\deg(v)-1$, 0, and 1, respectively.*
(UP2) *If v is a non-switch vertex, the number of small, flat and large angles incident to v is equal to $\deg(v)-2$, 2, and 0, respectively.*
(UP3) *If f is an internal face (the outer face) of $\mathcal{E}$, the number of small angles in f is equal to the number of large angles in f plus 2 (resp. minus 2).*

Let $(\mathcal{E},\lambda)$ be an upward plane embedding and let u be a vertex incident to $f_{\mathcal{E}}$. We often talk about *the angle at u inside $\mathcal{E}$*, referring to the complement of the angle at u in $f_{\mathcal{E}}$. Then the angle at u inside $\mathcal{E}$ is small, flat, or large if and only if the angle at u in $f_{\mathcal{E}}$ is large, flat, or small, respectively.

The UPWARD PLANARITY TESTING problem asks whether a given DAG admits an upward planar drawing. Because of Theorem 4, this is equivalent to deciding whether the DAG admits a plane embedding $\mathcal{E}$ and a label assignment for the angles of $\mathcal{E}$ satisfying Conditions (UP0)–(UP3).

3. Solving partition via upward planarity

In this section we prove Theorem 1.

Let $\mathcal{S} = \{a_1,\dots,a_k\}$ be an instance of the PARTITION problem, where $\sum_{i=1}^{k} a_i = n$. We start by describing the construction of an $\mathcal{O}(n)$-vertex biconnected outerplanar DAG G from $\mathcal{S}$. Refer to Fig. 1. The structure of G is simple, indeed it consists of a cycle C and of some other cycles, each sharing an edge with C and disjoint with each other. That is, the weak dual of an outerplane embedding $\mathcal{O}$ of G is a star, where the *weak dual* of $\mathcal{O}$ is the graph that has a vertex for each internal face of $\mathcal{O}$ and an edge between two vertices if the corresponding faces share an edge.

We initialize G to a cycle $C = (u_0, u_1, \dots, u_{6n+2k+1})$. Let $\mathcal{P}$ be the path $(u_1, u_2, \dots, u_{6n+2k})$ in C. Let $\mathcal{P}_1$ be the path that comprises the first $6a_1+2$ vertices $(u_1,\dots,u_{6a_1+2})$ of $\mathcal{P}$, let $\mathcal{P}_2$ be the path that comprises the next $6a_2+2$ vertices $(u_{6a_1+3},\dots,u_{6a_1+6a_2+4})$ of $\mathcal{P}$, and so on till $\mathcal{P}_k$. For each edge (u_j,u_{j+1}) of C that belongs to a path $\mathcal{P}_i$, we introduce in G two vertices v_j and w_j and the edges of a path $\mathcal{Q}_j := (u_j, v_j, w_j, u_{j+1})$; we denote by C_j the cycle $\mathcal{Q}_j \cup (u_j,u_{j+1})$. Finally, we direct the edges of G as follows.

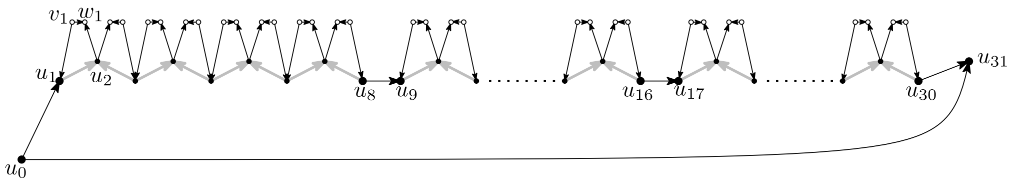

Fig. 1. Construction of G from $\mathcal{S} = \{1, 1, 2\}$; then $n = 4$ and $k = 3$. Vertices of C are filled disks, while other vertices are empty disks. Larger disks represent u_0, $u_{6n+2k+1}$, and the initial and final vertices of the paths $\mathcal{P}_1, \dots, \mathcal{P}_k$, which are represented by thicker gray lines.

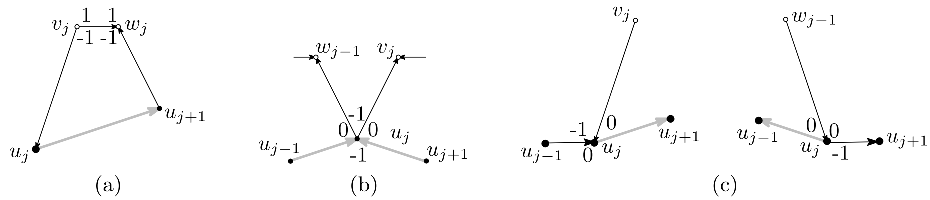

Fig. 2. Labels for some angles of $\mathcal{E}$. (a) Switch vertices v_j and w_j. (b) Vertices u_j that are internal to some path $\mathcal{P}_i$. (c) Vertices u_j that are the first or last vertex of some path $\mathcal{P}_i$.

- First, for $i = 1, \dots, k$, each edge (u_j, u_{j+1}) of $\mathcal{P}_i$ is outgoing from the vertex with odd index.
- Second, for $i = 1, \dots, k-1$, the edge connecting the last vertex of $\mathcal{P}_i$ to the first vertex of $\mathcal{P}_{i+1}$ is outgoing from the former vertex.
- Third, both the edges incident to u_0 are outgoing from u_0 and both the edges incident to $u_{6n+2k+1}$ are incoming into $u_{6n+2k+1}$.
- Finally, the edges of the path $\mathcal{Q}_j$ are outgoing from v_j and incoming into w_j if j is odd, while they are incoming into v_j and outgoing from w_j if j is even.

Clearly, G is an $\mathcal{O}(n)$-vertex biconnected outerplanar DAG and can be constructed in $\mathcal{O}(n)$ time. Next, we prove that $\mathcal{S}$ is a positive instance of PARTITION if and only if G is a positive instance of the UPWARD PLANARITY TESTING problem. For $i = 1, \dots, k$, let $\mathcal{V}_i$ be the set that comprises the vertices of $\mathcal{P}_i$ and the vertices of the paths $\mathcal{Q}_j$ such that (u_j, u_{j+1}) belongs to $\mathcal{P}_i$.

($\Longrightarrow$) Suppose first that there exists a partition of $\mathcal{S}$ into two sets $\mathcal{S}_1$ and $\mathcal{S}_2$ such that $\sum_{a_i \in \mathcal{S}_1} a_i = \sum_{a_i \in \mathcal{S}_2} a_i = n/2$. We construct a plane embedding $\mathcal{E}$ of G as follows. Starting from any plane embedding of C, we embed inside C any path $\mathcal{Q}_j$ such that the edge (u_j, u_{j+1}) belongs to a path $\mathcal{P}_i$ with $a_i \in \mathcal{S}_1$; also, we embed outside C any path $\mathcal{Q}_j$ such that the edge (u_j, u_{j+1}) belongs to a path $\mathcal{P}_i$ with $a_i \in \mathcal{S}_2$, so that C_j delimits an internal face of $\mathcal{E}$. Next, we specify a label assignment λ for the angles of $\mathcal{E}$.

- First, consider any degree-2 switch vertex v_j (resp. w_j); see Fig. 2(a). We label -1 the angle at v_j (resp. at w_j) in the internal face of $\mathcal{E}$ delimited by C_j, and $+1$ the other angle v_j (resp. w_j) is incident to; this assignment respects Conditions (UP0), (UP1), and, vacuously, (UP2) of Theorem 4.
- Second, consider any degree-4 non-switch vertex u_j that is an internal vertex of a path $\mathcal{P}_i$; see Fig. 2(b). We label 0 the angles at u_j in the internal faces of $\mathcal{E}$ delimited by C_j and C_{j+1}, and -1 the other two angles u_j is incident to; this assignment respects Conditions (UP0), (UP2), and, vacuously, (UP1) of Theorem 4.
- Third, consider any degree-3 non-switch vertex u_j that is the first (resp. last) vertex of a path $\mathcal{P}_i$; see Fig. 2(c). We label 0 the angle at u_j in the internal face of $\mathcal{E}$ delimited by C_j (resp. by C_{j-1}). If C_j (resp. C_{j-1}) is embedded inside C, then we label 0 the angle at u_j in the outer face of $\mathcal{E}$ and -1 the remaining angle u_j is incident to, otherwise we label -1 the angle at u_j in the outer face of $\mathcal{E}$ and 0 the remaining angle u_j is incident to; this assignment respects Conditions (UP0), (UP2), and, vacuously, (UP1) of Theorem 4.
- Finally, consider the degree-2 switch vertex u_0 (resp. $u_{6n+2k+1}$). We label $+1$ the angle at u_0 (resp. $u_{6n+2k+1}$) in the outer face of $\mathcal{E}$, and -1 the other angle u_0 (resp. $u_{6n+2k+1}$) is incident to; this assignment respects Conditions (UP0), (UP1), and, vacuously, (UP2) of Theorem 4.

In order to prove that $(\mathcal{E}, \lambda)$ is an upward plane embedding of G, it remains to prove that λ satisfies Condition (UP3) of Theorem 4. Consider any internal face delimited by a cycle C_j. Two angles in such a face, those at v_j and w_j, are labeled -1, while the other two angles, those at u_j and u_{j+1}, are labeled 0, hence λ respects Condition (UP3) for the considered face. Two faces remain to be considered, namely the outer face $f_{\mathcal{E}}$ and the internal face $g_{\mathcal{E}}$ incident to u_0. In order to deal with these faces, we need to observe the following facts.

- If a_i belongs to $\mathcal{S}_1$, the number of angles of $f_{\mathcal{E}}$ labeled -1 at the vertices in $\mathcal{V}_i$ is $6a_i$, while the number of angles of $f_{\mathcal{E}}$ labeled $+1$ at the vertices in $\mathcal{V}_i$ is 0. Indeed, all the angles of $f_{\mathcal{E}}$ at the vertices of $\mathcal{P}_i$ are labeled -1, except for the angles at the end-vertices of $\mathcal{P}_i$, which are labeled 0. The vertices v_j and w_j in $\mathcal{V}_i$ are not incident to $f_{\mathcal{E}}$. Analogously, if a_i belongs to $\mathcal{S}_2$, the number of angles of $g_{\mathcal{E}}$ labeled -1 at the vertices in $\mathcal{V}_i$ is $6a_i$, while the number of angles of $g_{\mathcal{E}}$ labeled $+1$ at the vertices in $\mathcal{V}_i$ is 0.

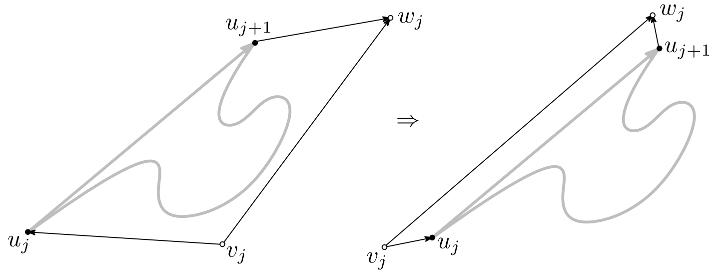

Fig. 3. The embedding $\mathcal{E}$ can be modified by re-embedding a single path $\mathcal{Q}_j$ so that every cycle C_j delimits an internal face of $\mathcal{E}$.

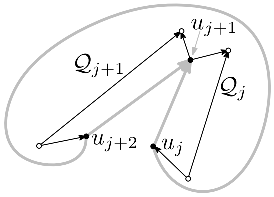

Fig. 4. If $\mathcal{Q}_j$ lies inside C (gray in the figure) in $\mathcal{E}$, then $\mathcal{Q}_{j+1}$ also lies inside C in $\mathcal{E}$.

- If a_i belongs to $\mathcal{S}_1$, the number of angles of $g_\mathcal{E}$ labeled -1 at the vertices in $\mathcal{V}_i$ is $6a_i + 2$, while the number of angles of $g_\mathcal{E}$ labeled $+1$ at the vertices in $\mathcal{V}_i$ is $12a_i + 2$. Indeed, all the angles of $g_\mathcal{E}$ at the vertices of $\mathcal{P}_i$ are labeled -1 and all the angles of $g_\mathcal{E}$ at the vertices v_j and w_j in $\mathcal{V}_i$ are labeled $+1$ (note that there are $6a_i + 1$ cycles C_j such that the edge (u_j, u_{j+1}) belongs to $\mathcal{P}_i$). Analogously, if a_i belongs to $\mathcal{S}_2$, then the number of angles of $f_\mathcal{E}$ labeled -1 at the vertices in $\mathcal{V}_i$ is $6a_i + 2$, while the number of angles of $f_\mathcal{E}$ labeled $+1$ at the vertices in $\mathcal{V}_i$ is $12a_i + 2$.

It follows that the number of angles of $f_\mathcal{E}$ labeled $+1$ is $2 + \sum_{a_i \in \mathcal{S}_2}(12a_i + 2)$, where the first term of the sum comes from the angles at u_0 and $u_{6n+2k+1}$, while the number of angles of $f_\mathcal{E}$ labeled -1 is $\sum_{a_i \in \mathcal{S}_1}(6a_i) + \sum_{a_i \in \mathcal{S}_2}(6a_i + 2)$. Thus, the number of angles of $f_\mathcal{E}$ labeled $+1$ minus the number of angles of $f_\mathcal{E}$ labeled -1 is $2 + \sum_{a_i \in \mathcal{S}_2}(6a_i) - \sum_{a_i \in \mathcal{S}_1}(6a_i) = 2$, where the last equality follows from $\sum_{a_i \in \mathcal{S}_1} a_i = \sum_{a_i \in \mathcal{S}_2} a_i$. Hence, λ respects Condition (UP3) of Theorem 4 for $f_\mathcal{E}$. Analogously, the number of angles of $g_\mathcal{E}$ labeled $+1$ is $\sum_{a_i \in \mathcal{S}_1}(12a_i + 2)$, while the number of angles of $g_\mathcal{E}$ labeled -1 is $2 + \sum_{a_i \in \mathcal{S}_1}(6a_i + 2) + \sum_{a_i \in \mathcal{S}_2}(6a_i)$, where the first term of the sum comes from the angles at u_0 and $u_{6n+2k+1}$. Thus, the number of angles of $g_\mathcal{E}$ labeled -1 minus the number of angles of $g_\mathcal{E}$ labeled $+1$ is $2 - \sum_{a_i \in \mathcal{S}_1}(6a_i) + \sum_{a_i \in \mathcal{S}_2}(6a_i) = 2$, where the last equality follows from $\sum_{a_i \in \mathcal{S}_1} a_i = \sum_{a_i \in \mathcal{S}_2} a_i$. Hence, λ respects Condition (UP3) of Theorem 4 for $g_\mathcal{E}$.

($\Longleftarrow$) Suppose now that G admits an upward plane embedding $(\mathcal{E}, \lambda)$. We construct a partition of $\mathcal{S}$ into two sets $\mathcal{S}_1$ and $\mathcal{S}_2$ such that $\sum_{a_i \in \mathcal{S}_1} a_i = \sum_{a_i \in \mathcal{S}_2} a_i = n/2$ as follows. For $i = 1, \dots, k$, denote by $\mathcal{Q}^*(i)$ the path $\mathcal{Q}_j$ that forms a cycle with the first edge (u_j, u_{j+1}) of $\mathcal{P}_i$. If $\mathcal{Q}^*(i)$ lies inside C in $\mathcal{E}$, then we put a_i into $\mathcal{S}_1$, otherwise into $\mathcal{S}_2$. Clearly, the partition can be constructed in $\mathcal{O}(n)$ time. Observe that every cycle C_j delimits an internal face of $\mathcal{E}$, except, possibly, for a single cycle C_j which might delimit the outer face of $\mathcal{E}$. If that is the case, we re-embed the path $\mathcal{Q}_j$ so that every cycle C_j delimits an internal face of $\mathcal{E}$. This modification does not alter whether any path $\mathcal{Q}_j$ is embedded inside or outside C; see Fig. 3.

We now prove that $\sum_{a_i \in \mathcal{S}_1} a_i = n/2$, which implies that $\sum_{a_i \in \mathcal{S}_2} a_i = n/2$. As before, let $f_\mathcal{E}$ be the outer face of $\mathcal{E}$ and $g_\mathcal{E}$ be the internal face of $\mathcal{E}$ incident to u_0. We are going to use the following lemmata.

Lemma 1. *For any $i \in \{1, \dots, k\}$, either all the paths $\mathcal{Q}_j$ such that (u_j, u_{j+1}) belongs to $\mathcal{P}_i$ lie inside C, or they all lie outside C in $\mathcal{E}$.*

Proof. Consider any index j such that $\mathcal{Q}_j$ and $\mathcal{Q}_{j+1}$ both exist. We prove that, if $\mathcal{Q}_j$ lies inside C in $\mathcal{E}$, then $\mathcal{Q}_{j+1}$ also lies inside C in $\mathcal{E}$; refer to Fig. 4. An analogous proof shows that, if $\mathcal{Q}_j$ lies outside C in $\mathcal{E}$, then $\mathcal{Q}_{j+1}$ also lies outside C in $\mathcal{E}$. Note that this claim implies the statement of the lemma. Let $(\mathcal{E}_C, \lambda_C)$ be the upward plane embedding of C which is the restriction of $(\mathcal{E}, \lambda)$ to C. Since u_{j+1} is a non-switch vertex in C_j and in C_{j+1}, the angles at u_{j+1} in the faces delimited by C_j and by C_{j+1} are both flat in $(\mathcal{E}, \lambda)$, by Condition (UP0) of Theorem 4. Since u_{j+1} is a switch vertex in C, the planarity of $\mathcal{E}$ implies that the angles at u_{j+1} inside and outside C in $(\mathcal{E}_C, \lambda_C)$ are large and small, respectively. Then the planarity of $\mathcal{E}$ implies that C_{j+1} lies inside C. □

Lemma 2. *For any $i \in \{1, \dots, k\}$, suppose that $\mathcal{Q}^*(i)$ lies inside C in $\mathcal{E}$. Then:*

- *the number of angles of $f_\mathcal{E}$ labeled -1 at the vertices in $\mathcal{V}_i$ is $6a_i$;*
- *the number of angles of $f_\mathcal{E}$ labeled $+1$ at the vertices in $\mathcal{V}_i$ is 0;*
- *the number of angles of $g_\mathcal{E}$ labeled -1 at the vertices in $\mathcal{V}_i$ is $6a_i + 2$;*
- *the number of angles of $g_\mathcal{E}$ labeled $+1$ at the vertices in $\mathcal{V}_i$ is $12a_i + 2$.*

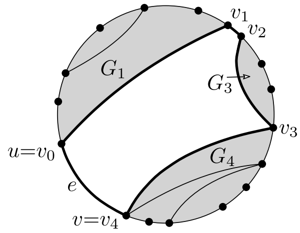

Fig. 5. A biconnected outerplanar DAG G. Edge directions are not shown. A cycle C is thick and the internal faces of the subgraphs G_i of G are colored gray.

Proof. By Lemma 1, all the paths $\mathcal{Q}_j$ such that the edge (u_j, u_{j+1}) belongs to $\mathcal{P}_i$ lie inside C in $\mathcal{E}$. We first argue about the angles at the vertices in $\mathcal{V}_i$ in $(\mathcal{E}, \lambda)$. Every cycle C_j such that the edge (u_j, u_{j+1}) belongs to $\mathcal{P}_i$ delimits an internal face f_j of $\mathcal{E}$. Since u_j and u_{j+1} are degree-2 non-switch vertices of C_j, by Condition (UP0) of Theorem 4, the angles at u_j and u_{j+1} in f_j are both flat in $(\mathcal{E}, \lambda)$. By Condition (UP3), the angles at v_j and w_j in f_j are both small in $(\mathcal{E}, \lambda)$, and thus, by Condition (UP1), the angles at v_j and w_j in $g_{\mathcal{E}}$ are both large in $(\mathcal{E}, \lambda)$. Every vertex u_j in $\mathcal{V}_i$ which is neither the first nor the last vertex of $\mathcal{P}_i$ has two incident flat angles in f_{j-1} and f_j, hence by Condition (UP2), its incident angles in $f_{\mathcal{E}}$ and $g_{\mathcal{E}}$ are both small. Finally, the first (resp. last) vertex of $\mathcal{P}_i$ has one incident flat angle in f_j (resp. by f_{j-1}) and one incident flat angle in $f_{\mathcal{E}}$, hence, by Condition (UP1), its incident angle in $g_{\mathcal{E}}$ is small.

We are now ready to count. All the angles in $f_{\mathcal{E}}$ at the vertices of $\mathcal{P}_i$ are labeled -1, except for the angles at the end-vertices of $\mathcal{P}_i$, which are labeled 0; these are the $6a_i$ angles in $f_{\mathcal{E}}$ labeled -1. The vertices v_j and w_j in $\mathcal{V}_i$ are not incident to $f_{\mathcal{E}}$. Thus, no angle in $f_{\mathcal{E}}$ is labeled $+1$. Further, all the angles in $g_{\mathcal{E}}$ at the vertices of $\mathcal{P}_i$ are labeled -1; these are the $6a_i + 2$ angles in $g_{\mathcal{E}}$ labeled -1. Finally, all the angles in $g_{\mathcal{E}}$ at the vertices v_j and w_j in $\mathcal{V}_i$ are labeled $+1$; these are the $12a_i + 2$ angles in $g_{\mathcal{E}}$ labeled $+1$, as there are $6a_i + 1$ cycles C_j such that the edge (u_j, u_{j+1}) belongs to $\mathcal{P}_i$. □

The proof of the following is analogous to the one of Lemma 2.

Lemma 3. *For any $i \in \{1, \dots, k\}$, suppose that $\mathcal{Q}^*(i)$ lies outside C in $\mathcal{E}$. Then:*

- *the number of angles of $f_{\mathcal{E}}$ labeled -1 at the vertices in $\mathcal{V}_i$ is $6a_i + 2$;*
- *the number of angles of $f_{\mathcal{E}}$ labeled $+1$ at the vertices in $\mathcal{V}_i$ is $12a_i + 2$;*
- *the number of angles of $g_{\mathcal{E}}$ labeled -1 at the vertices in $\mathcal{V}_i$ is $6a_i$;*
- *the number of angles of $g_{\mathcal{E}}$ labeled $+1$ at the vertices in $\mathcal{V}_i$ is 0.*

Recall that, if $\mathcal{Q}^*(i)$ lies inside (outside) C, then $a_i \in \mathcal{S}_1$ (resp. $a_i \in \mathcal{S}_2$). Thus, by Lemmata 2 and 3, we have that, over all the vertices in $\mathcal{V}_1 \cup \dots \cup \mathcal{V}_k$, the face $f_{\mathcal{E}}$ is incident to $\sum_{a_i \in \mathcal{S}_1}(6a_i) + \sum_{a_i \in \mathcal{S}_2}(6a_i + 2)$ angles labeled -1 and to $\sum_{a_i \in \mathcal{S}_2}(12a_i + 2)$ angles labeled $+1$, hence the difference δ between the number of angles labeled $+1$ in $f_{\mathcal{E}}$ and the number of angles labeled -1 in $f_{\mathcal{E}}$ at the vertices in $\mathcal{V}_1 \cup \dots \cup \mathcal{V}_k$ is equal to $6(\sum_{a_i \in \mathcal{S}_2} a_i - \sum_{a_i \in \mathcal{S}_1} a_i)$. Assume, for a contradiction, that $\sum_{a_i \in \mathcal{S}_1} a_i \neq \sum_{a_i \in \mathcal{S}_2} a_i$. This implies that δ is larger than or equal to 6, or smaller than or equal to -6. Apart from $\mathcal{V}_1 \cup \dots \cup \mathcal{V}_k$, the vertex set of G only includes two more vertices, namely u_0 and $u_{u_{6n+2k+1}}$. Thus, the difference between the number of angles labeled $+1$ in $f_{\mathcal{E}}$ and the number of angles labeled -1 in $f_{\mathcal{E}}$ is larger than or equal to 4, or smaller than or equal to -4. However, by Condition (UP3) of Theorem 4, such a difference is equal to 2, a contradiction which proves that $\sum_{a_i \in \mathcal{S}_1} a_i = \sum_{a_i \in \mathcal{S}_2} a_i$ and concludes the proof of Theorem 1.

4. Solving upward planarity via partition

In this section we prove Theorem 2. Our algorithm borrows some ideas from the algorithm for testing the rectilinear planarity of an outerplanar graph [15]. Throughout the section, we assume that there exists an algorithm with $f(n)$ running time that decides whether a set of positive integers can be partitioned into two sets, so that the sum of the elements in each set is equal to $n/2$.

Refer to Fig. 5. Let G be an n-vertex biconnected outerplanar DAG, let $\mathcal{O}$ be any of its two outerplane embeddings, let $e = (u, v)$ be a prescribed edge incident to $f_{\mathcal{O}}$, and let $C = (v_0 = u, v_1, \dots, v_k = v)$ be the cycle bounding the internal face of $\mathcal{O}$ the edge e is incident to. For any $i \in \{1, \dots, k\}$ such that the edge (v_{i-1}, v_i) is not incident to $f_{\mathcal{O}}$, let $\mathcal{P}_i$ be the path along $f_{\mathcal{O}}$ that connects v_{i-1} and v_i and that does not comprise vertices of C other than v_{i-1} and v_i. We denote by G_i the subgraph of G induced by the vertices of $\mathcal{P}_i$.

We want to test in $\mathcal{O}(f(n))$ time whether G admits an upward plane embedding in which e is incident to the outer face. Our algorithm not only decides whether G admits such an embedding, but it actually computes the *feasible set* $\mathcal{F}$ of G, that is, the set of all pairs (μ, ν) with $\mu, \nu \in \{-1, 0, 1\}$ such that G admits an upward plane embedding $(\mathcal{E}, \lambda)$ in which e is incident to $f_{\mathcal{E}}$ and the angles at u and v inside $\mathcal{E}$ are μ and ν, respectively; we call this a (μ, ν)-*embedding*. As the computation of the feasible sets will eventually be done recursively, we assume to have the feasible set $\mathcal{F}_i$ of each subgraph G_i of G; note that, in a (μ_i, ν_i)-embedding of G_i, the edge (v_{i-1}, v_i) is incident to the outer face. Since $\mathcal{F}$ might comprise at most 9 pairs, we consider each pair (μ, ν) individually and show how to test whether $(\mu, \nu) \in \mathcal{F}$ in $\mathcal{O}(f(k))$ time. The basis for our algorithm is the following structural lemma.

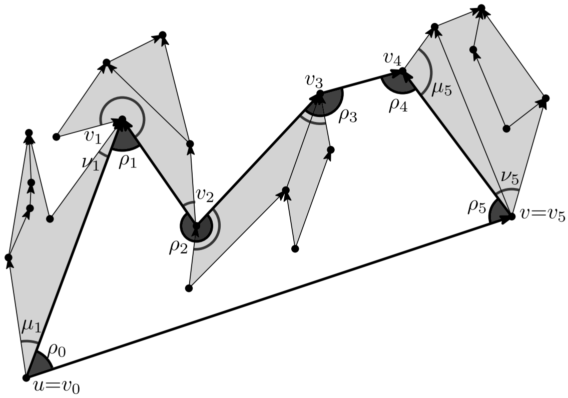

Fig. 6. Illustration for Lemma 4. The cycle C is thick, the internal faces of the subgraphs G_i are gray, the angles ρ_i are black opaque, and the angles μ_i and ν_i are not filled.

Lemma 4. *For any $\mu, \nu \in \{-1, 0, 1\}$, we have that G admits a (μ, ν)-embedding if and only if there exist values $\rho_0, \dots, \rho_k, \mu_1, \dots, \mu_k, \nu_1, \dots, \nu_k$ in $\{-1, 0, 1\}$ and there exists an* in-out assignment *$\mathcal{A}$, that is, an assignment of each subgraph G_i of G to the inside or to the outside of C, such that (refer to Fig. 6):*

- *$\mathcal{P}1$: For $i = 0, \dots, k$, if v_i is a switch vertex in C, then $\rho_i \in \{-1, 1\}$, else $\rho_i = 0$.*
- *$\mathcal{P}2$: $\sum_{i=0}^{k} \rho_i = -2$.*
- *$\mathcal{P}3$: For $i = 1, \dots, k$, if G_i is defined, then $(\mu_i, \nu_i) \in \mathcal{F}_i$.*
- *$\mathcal{P}4$: For $i = 1, \dots, k$, if G_i is defined and $\mathcal{A}(G_i) =$ in, then $\mu_i \leq \rho_{i-1}$ and $\nu_i \leq \rho_i$; also, if G_i is defined and $\mathcal{A}(G_i) =$ out, then $\mu_i \leq -\rho_{i-1}$ and $\nu_i \leq -\rho_i$.*
- *$\mathcal{P}5$: For $i = 1, \dots, k-1$, we have $\nu_i + \mu_{i+1} \leq 0$.*
- *$\mathcal{P}6$: For $i = 1, \dots, k-1$, if $\nu_i = \mu_{i+1} = 0$, if G_i and G_{i+1} are defined, and if $\mathcal{A}(G_i) = \mathcal{A}(G_{i+1}) =$ in, then $\rho_i = 1$; also, if $\nu_i = \mu_{i+1} = 0$, if G_i and G_{i+1} are defined, and if $\mathcal{A}(G_i) = \mathcal{A}(G_{i+1}) =$ out, then $\rho_i = -1$.*
- *$\mathcal{P}7$: If G_1 is undefined or $\mathcal{A}(G_1) =$ in, then $\mu = \rho_0$, otherwise $\mu = \rho_0 + \mu_1 + 1$. If G_k is undefined or $\mathcal{A}(G_k) =$ in, then $\nu = \rho_k$, otherwise $\nu = \rho_k + \nu_k + 1$.*

Proof. ($\Longrightarrow$) Let $(\mathcal{E}, \lambda)$ be a (μ, ν)-embedding of G. For each $i \in \{1, \dots, k\}$ such that G_i is defined, let $\mathcal{A}(G_i) =$ in if G_i lies inside C in $\mathcal{E}$, and let $\mathcal{A}(G_i) =$ out otherwise. For $i = 0, \dots, k$, let ρ_i be the angle at v_i inside the upward plane embedding $\mathcal{E}_C$ of C induced by $\mathcal{E}$. Also, for $i = 1, \dots, k$, if G_i is defined, let μ_i and ν_i be the angles at v_{i-1} and v_i, respectively, inside the embedding $\mathcal{E}_i$ of G_i induced by $\mathcal{E}$, otherwise let $\mu_i = \nu_i = -1$.

Properties $\mathcal{P}_1$ and $\mathcal{P}_2$ then follow from Properties (UP0) and (UP3) of Theorem 4, respectively, applied to $\mathcal{E}_C$. Property $\mathcal{P}_3$ follows from the definition of μ_i and ν_i. Property $\mathcal{P}_4$ follows from the fact that, if G_i is defined and $\mathcal{A}(G_i) =$ in, then G_i lies inside C in $\mathcal{E}$, thus the angles at v_{i-1} and v_i inside $\mathcal{E}_i$, which are equal to μ_i and ν_i, are respectively smaller than or equal to the angles at v_{i-1} and v_i inside $\mathcal{E}_C$, which are equal to ρ_{i-1} and ρ_i; the case in which G_i is defined and $\mathcal{A}(G_i) =$ out can be discussed similarly. Property $\mathcal{P}_5$ is trivial if G_i or G_{i+1} is undefined, as in this case one of ν_i and μ_{i+1} is equal to -1 and the other one does not exceed 1; otherwise, it comes from Properties (UP1) and (UP2) of Theorem 4 – by restricting $\mathcal{E}$ to the cycles bounding the outer faces of $\mathcal{E}_i$ and $\mathcal{E}_{i+1}$. Property $\mathcal{P}_6$ comes from the fact that an angle comprising two flat angles is large – note that $\mathcal{E}_i$ and $\mathcal{E}_{i+1}$ are one outside the other, since e is incident to $f_{\mathcal{E}}$. Finally, Property $\mathcal{P}_7$ comes from the fact that the angles at u and v inside $\mathcal{E}$ are μ and ν, respectively. If G_1 is undefined or if it lies inside C in $\mathcal{E}$, then the angle at u inside $\mathcal{E}$, which is μ, coincides with the angle at u inside $\mathcal{E}_C$, which is ρ_0. If G_1 is defined and lies outside C in $\mathcal{E}$, then the angles at u inside $\mathcal{E}_1$ and inside $\mathcal{E}_C$ "sum up" to compose the angle at u inside $\mathcal{E}$, however the numerical definition of such angles needs a $+1$ in the sum; for example, if the angles at u inside $\mathcal{E}_1$ and inside $\mathcal{E}_C$ are both equal to -1, then the angle at u inside $\mathcal{E}$ is also equal to -1; or if they are both equal to 0, then the angle at u inside $\mathcal{E}$ is equal to 1. A similar argument applies for the angles at v.

($\Longleftarrow$) Assume now that there exist values $\rho_0, \dots, \rho_k, \mu_1, \dots, \mu_k, \nu_1, \dots, \nu_k$ in $\{-1, 0, 1\}$ and there exists an in-out assignment $\mathcal{A}$ satisfying Properties $\mathcal{P}_1$–$\mathcal{P}_7$. We construct a (μ, ν)-embedding of G as follows. Let $\mathcal{E}_C$ be any of the two plane embeddings of C. For $i = 0, \dots, k$, we label the angle at v_i inside $\mathcal{E}_C$ as ρ_i and the angle at v_i in $f_{\mathcal{E}_C}$ as $-\rho_i$. By Properties $\mathcal{P}_1$ and $\mathcal{P}_2$, the described labeling λ_C, together with $\mathcal{E}_C$, defines an upward plane embedding of C. For $i = 1, \dots, k$, if G_i is defined, then we let $(\mathcal{E}_i, \lambda_i)$ be a (μ_i, ν_i)-embedding of G_i. This exists by Property $\mathcal{P}_3$; we embed $\mathcal{E}_i$ inside C if $\mathcal{A}(G_i) =$ in and we embed $\mathcal{E}_i$ outside C if $\mathcal{A}(G_i) =$ out. This results in a plane embedding $\mathcal{E}$ of G. The assignment λ of labels to the angles in $\mathcal{E}$ is the one obtained by combining $(\mathcal{E}_C, \lambda_C)$ with the embeddings $(\mathcal{E}_i, \lambda_i)$. In particular, let $g_{\mathcal{E}}$ be the internal face of $\mathcal{E}$ incident to e. Then the angle at v_i in $g_{\mathcal{E}}$ is obtained from ρ_i by subtracting $\nu_i + 1$ if G_i is defined and $\mathcal{A}(G_i) =$ in, and by subtracting $\mu_{i+1} + 1$ if G_{i+1} is defined and $\mathcal{A}(G_{i+1}) =$ in. The angle at v_i in $f_{\mathcal{E}}$ is determined analogously. Properties $\mathcal{P}_4$–$\mathcal{P}_6$ ensure that $(\mathcal{E}, \lambda)$ is an upward plane embedding. In particular, Property $\mathcal{P}_4$ ensures that the angles at v_{i-1} and v_i inside $\mathcal{E}_i$ "fit" inside $\mathcal{E}_C$ (if $\mathcal{A}(G_i) =$ in) or outside $\mathcal{E}_C$ (if $\mathcal{A}(G_i) =$ out). Furthermore, Property $\mathcal{P}_5$ ensures that the angles at v_i inside $\mathcal{E}_i$ and $\mathcal{E}_{i+1}$ fit together in $\mathcal{E}$ (avoiding the possibility that they are one large and one flat, or both large). Moreover, Property $\mathcal{P}_6$ ensures that the angles at v_i inside $\mathcal{E}_i$ and $\mathcal{E}_{i+1}$ simultaneously fit inside $\mathcal{E}_C$ (if G_i and G_{i+1} are defined and $\mathcal{A}(G_i) = \mathcal{A}(G_{i+1}) = \text{in}$) or outside $\mathcal{E}_C$ (if G_i and G_{i+1} are defined and $\mathcal{A}(G_i) = \mathcal{A}(G_{i+1}) = \text{out}$). Finally, Property $\mathcal{P}_7$ ensures that the upward plane embedding $(\mathcal{E}, \lambda)$ is a (μ, ν)-embedding, by requiring that the angles at u and v inside $\mathcal{E}$ are μ and ν, respectively. □

By Lemma 4, in order to test whether a pair (μ, ν) with $\mu, \nu \in \{-1, 0, 1\}$ belongs to the feasible set $\mathcal{F}$ of G, our algorithm will decide whether there exist values $\rho_0, \dots, \rho_k$, $\mu_1, \dots, \mu_k$, $\nu_1, \dots, \nu_k$ in $\{-1, 0, 1\}$ and an in-out assignment $\mathcal{A}$ satisfying Properties $\mathcal{P}_1$–$\mathcal{P}_7$.

The overview of our algorithm is as follows. First, we fix $\mathcal{O}(1)$ values among the ones that we have to decide, in all possible ways. This ensures that things are "fine" close to u and v, at the expense of performing the next steps $\mathcal{O}(1)$ times, namely once for each way to fix the above $\mathcal{O}(1)$ values. Then we show that the remaining values μ_i and ν_i can be decided in $\mathcal{O}(k)$ time without loss of generality by looking at the graph structure. Finally, we have to decide the values ρ_i, as well as the in-out assignment $\mathcal{A}$. The previous decisions on the values μ_i and ν_i bind together some subgraphs G_i; assigning such bound subgraphs to the inside or to the outside of C gives a contribution to the sum of the angles inside or outside C. Balancing such contributions can be done via PARTITION.

Recall that we have, for each subgraph G_i of G, its feasible set $\mathcal{F}_i$. We can assume that $\mathcal{F}_i \neq \emptyset$, as otherwise we can conclude that G admits no (μ, ν)-embedding, by Property $\mathcal{P}_3$ of Lemma 4. The following lemma will be useful.

Lemma 5. *A feasible set $\mathcal{F}_i$ does not contain two pairs $(-1, \cdot)$ and $(0, \cdot)$, and does not contain two pairs $(\cdot, -1)$ and $(\cdot, 0)$.*

Proof. If $\mathcal{F}_i$ contains a pair $(-1, \cdot)$, then v_{i-1} is a switch vertex in G_i, while if it contains a pair $(0, \cdot)$, then v_{i-1} is a non-switch vertex in G_i, hence $\mathcal{F}_i$ cannot contain both pairs. The proof for the pairs $(\cdot, -1)$ and $(\cdot, 0)$ is analogous. □

We start by fixing ρ_0 and ρ_k in every possible way – each of them has a value in $\{-1, 0, 1\}$. Also, if G_1 is defined, we fix μ_1 and $\mathcal{A}(G_1)$ in every possible way – note that $\mu_1 \in \{-1, 0, 1\}$ and $\mathcal{A}(G_1) \in \{\text{in}, \text{out}\}$; otherwise, we let $\mu_1 = -1$. Analogously, if G_k is defined, we fix μ_k, ν_k, and $\mathcal{A}(G_k)$ in every possible way, otherwise we let $\mu_k = \nu_k = -1$. This results in $\mathcal{O}(1)$ tuples of values $\langle \rho_0, \rho_k, \mu_1, \mathcal{A}(G_1), \mu_k, \nu_k, \mathcal{A}(G_k) \rangle$. For each of them, we test in $\mathcal{O}(1)$ time whether it satisfies Properties $\mathcal{P}_1$ (this concerns the values ρ_0 and ρ_k), $\mathcal{P}_3$ (this concerns the values μ_k and ν_k), $\mathcal{P}_4$ (this concerns the values ρ_0, μ_1 and $\mathcal{A}(G_1)$, and the values ρ_k, ν_k and $\mathcal{A}(G_k)$), and $\mathcal{P}_7$ (this concerns the values ρ_0, μ_1 and $\mathcal{A}(G_1)$, and the values ρ_k, ν_k and $\mathcal{A}(G_k)$) of Lemma 4. If not, we discard the tuple. Otherwise, we say it is a *candidate tuple*. We are going to independently test in $\mathcal{O}(f(k))$ time whether each candidate tuple is *extensible*, i.e., whether there exist values for $\rho_1, \dots, \rho_{k-1}, \mu_2, \dots, \mu_{k-1}, \nu_1, \dots, \nu_{k-1}$ and assignments for the subgraphs $G_2, \dots, G_{k-1}$ that are defined that, together with the values in the candidate tuple, satisfy Properties $\mathcal{P}_1$–$\mathcal{P}_7$ of Lemma 4.

The following lemmata prove that the values $\mu_1, \nu_1, \mu_2, \dots, \nu_{k-1}$ can be decided *without loss of generality*, in this order; this means that, if we decided the values $\mu_1, \nu_1, \mu_2, \nu_2, \dots, \mu_i$ (or $\nu_1, \mu_2, \nu_2, \dots, \nu_i$), then the candidate tuple under consideration is extensible if and only if an extension exists with the chosen values $\nu_1, \mu_2, \nu_2, \dots, \mu_i$ (resp., $\nu_1, \mu_2, \nu_2, \dots, \nu_i$).

Lemma 6. *Suppose that, for some $i \in \{1, \dots, k-1\}$, the values $\mu_1, \nu_1, \dots, \mu_i$ have been decided without loss of generality. Then ν_i can also be decided without loss of generality in $\mathcal{O}(1)$ time, as follows. If G_i is undefined, we let $\nu_i = -1$. Otherwise, if $\mathcal{F}_i$ does not contain any pair $(\mu_i, \cdot)$, then it is concluded that the current candidate tuple is not extensible. Otherwise, we let ν_i be the smallest value β such that $(\mu_i, \beta) \in \mathcal{F}_i$.*

Proof. First, deciding ν_i takes $\mathcal{O}(1)$ time as the size of $\mathcal{F}_i$ is $\mathcal{O}(1)$. In order to prove that ν_i is decided without loss of generality, we only discuss the case in which G_i is defined, as the case in which it is not is simpler.

Since $\mu_1, \nu_1, \dots, \mu_i$ have been decided without loss of generality, if the candidate tuple $\langle \rho_0, \rho_k, \mu_1, \mathcal{A}(G_1), \nu_k, \mathcal{A}(G_k) \rangle$ is extensible, an extension exists with the chosen values $\nu_1, \mu_2, \nu_2, \dots, \mu_i$. If $\mathcal{F}_i$ does not contain any pair $(\mu_i, \cdot)$, then because the value μ_i has been decided without loss of generality, we can conclude that the current candidate tuple is not extensible. Otherwise, suppose that an extension exists and that, in such an extension, ν_i is not the smallest value β such that $(\mu_i, \beta) \in \mathcal{F}_i$, say $\nu_i = \beta' \neq \beta$. By Property $\mathcal{P}_3$ and by definition of feasible set, we have $(\mu_i, \beta') \in \mathcal{F}_i$, which implies that $\beta' > \beta$. Then, in the extension of $\langle \rho_0, \rho_k, \mu_1, \mathcal{A}(G_1), \nu_k, \mathcal{A}(G_k) \rangle$, we can replace the value $\nu_i = \beta'$ with $\nu_i = \beta$, thus decreasing such a value and leaving unaltered all other values and the in-out assignment. The resulting set of values, together with the in-out assignment, still satisfies Properties $\mathcal{P}_1$–$\mathcal{P}_7$ of Lemma 4. Indeed, Properties $\mathcal{P}_1$, $\mathcal{P}_2$, and $\mathcal{P}_7$ are not affected by the change of value for ν_i (note that $i \leq k-1$). Property $\mathcal{P}_3$ is satisfied as $(\mu_i, \beta) \in \mathcal{F}_i$. The inequalities $\nu_i \leq \rho_i$, $\nu_i \leq -\rho_i$, and $\nu_i + \mu_{i+1} \leq 0$ of Properties $\mathcal{P}_4$ and $\mathcal{P}_5$ are satisfied with $\nu_i = \beta'$, hence they are also satisfied with $\nu_i = \beta < \beta'$. Finally, it cannot be that $\nu_i = \beta = \mu_{i+1} = 0$, as this would imply that $\beta' = 1$, and hence that Property $\mathcal{P}_5$ is not satisfied in the given extension; it follows that Property $\mathcal{P}_6$ is also satisfied. □

Lemma 7. *Suppose that, for some $i \in \{1, \dots, k-2\}$, the values $\mu_1, \nu_1, \dots, \nu_i$ have been decided without loss of generality. Then μ_{i+1} can also be decided without loss of generality in $\mathcal{O}(1)$ time as follows. If G_{i+1} is undefined, we let $\mu_{i+1} = -1$. Otherwise, we distinguish four cases.*

- *If $\nu_i = 1$, then $\mu_{i+1} = -1$; if $\mathcal{F}_{i+1}$ does not contain any pair $(-1, \cdot)$, then it is concluded that the current candidate tuple is not extensible.*
- *If $\nu_i = 0$ and $\mathcal{F}_{i+1}$ contains a pair $(-1, \cdot)$ or contains a pair $(0, \cdot)$, then we let $\mu_{i+1} = -1$ or $\mu_{i+1} = 0$, respectively. If $\nu_i = 0$ and $\mathcal{F}_{i+1}$ contains neither a pair $(-1, \cdot)$ nor a pair $(0, \cdot)$, then it is concluded that the current candidate tuple is not extensible.*

- *Suppose that $\nu_i = -1$ and v_i is a non-switch vertex in C. If $\mathcal{F}_{i+1}$ contains a pair $(-1,\cdot)$ or contains a pair $(0,\cdot)$, then we let $\mu_{i+1} = -1$ or $\mu_{i+1} = 0$, respectively. If $\mathcal{F}_{i+1}$ contains neither a pair $(-1,\cdot)$ nor a pair $(0,\cdot)$, then it is concluded that the current candidate tuple is not extensible.*
- *Suppose finally that $\nu_i = -1$ and v_i is a switch vertex in C.*
  - *If $\mathcal{F}_{i+1}$ contains a pair (α,β) with $\{\alpha,\beta\} \subseteq \{-1,0\}$, we let $\mu_{i+1} = \alpha$.*
  - *Otherwise, if $\mathcal{F}_{i+1}$ contains the pair $(1,-1)$, we let $\mu_{i+1} = 1$.*
  - *Otherwise, if $\mathcal{F}_{i+1}$ contains the pair $(1,0)$ and does not contain the pair $(-1,1)$, we let $\mu_{i+1} = 1$.*
  - *Otherwise, if $\mathcal{F}_{i+1}$ contains the pair $(\alpha,1)$ with $\alpha \in \{-1,0\}$ and does not contain the pair $(1,0)$, we let $\mu_{i+1} = \alpha$.*
  - *Otherwise, if $\mathcal{F}_{i+1}$ contains the pairs $(1,0)$ and $(-1,1)$, we look at v_{i+1} and G_{i+2}. If v_{i+1} is a non-switch vertex in C, we let $\mu_{i+1} = 1$. Otherwise, if G_{i+2} is defined and $\mathcal{F}_{i+2}$ contains a pair $(0,\cdot)$, we let $\mu_{i+1} = 1$. Otherwise, we let $\mu_{i+1} = -1$.*
  - *Otherwise, $\mathcal{F}_{i+1} = \{(1,1)\}$, and we let $\mu_{i+1} = 1$.*

Proof. First, deciding μ_{i+1} takes $\mathcal{O}(1)$ time as the size of $\mathcal{F}_{i+1}$ is $\mathcal{O}(1)$. In order to prove that μ_{i+1} is decided without loss of generality, we only discuss the case in which G_{i+1} is defined, as the case in which it is not is simpler. Since $\mu_1, \nu_1, \dots, \nu_i$ have been decided without loss of generality, if there exists an extension of the current candidate tuple, there exists an extension φ in which $\mu_1, \nu_1, \dots, \nu_i$ have the decided values.

Suppose first that $\nu_i = 1$. Property $\mathcal{P}_5$ then implies that $\mu_{i+1} = -1$ in φ, which coincides with our decision. Property $\mathcal{P}_3$, together with the fact that $\mu_1, \nu_1, \dots, \nu_i$ have been decided without loss of generality, implies that the current candidate tuple is not extensible if $\mathcal{F}_{i+1}$ does not contain any pair $(-1,\cdot)$.

Suppose next that $\nu_i = 0$. Property $\mathcal{P}_5$ then implies that $\mu_{i+1} = -1$ or $\mu_{i+1} = 0$. By Lemma 5, we have that $\mathcal{F}_{i+1}$ cannot contain both a pair $(-1,\cdot)$ and a pair $(0,\cdot)$. Hence, also in this case, our decision for the value of μ_{i+1} coincides with the one in φ. Property $\mathcal{P}_3$, together with the fact that $\mu_1, \nu_1, \dots, \nu_i$ have been decided without loss of generality, implies that the current candidate tuple is not extensible if $\mathcal{F}_{i+1}$ does not contain any pair $(-1,\cdot)$ or $(0,\cdot)$.

Suppose next that $\nu_i = -1$ and v_i is a non-switch vertex in C. By Property $\mathcal{P}_1$, we have $\rho_i = 0$. By Property $\mathcal{P}_4$, we have $\mu_{i+1} \leq 0$, hence $\mu_{i+1} = -1$ or $\mu_{i+1} = 0$. By Lemma 5, we have that $\mathcal{F}_{i+1}$ cannot contain both a pair $(-1,\cdot)$ and a pair $(0,\cdot)$. Hence, also in this case, our decision for the value of μ_{i+1} coincides with the one in φ. Property $\mathcal{P}_3$, together with the fact that $\mu_1, \nu_1, \dots, \nu_i$ have been decided without loss of generality, implies that the current candidate tuple is not extensible if $\mathcal{F}_{i+1}$ does not contain any pair $(-1,\cdot)$ or $(0,\cdot)$.

Suppose finally that $\nu_i = -1$ and v_i is a switch vertex in C. By Property $\mathcal{P}_1$, we have $\rho_i = -1$ or $\rho_i = 1$ in φ; suppose the former, as the other case is analogous. Also, let α' and β' be the values of μ_{i+1} and ν_{i+1} in φ. We distinguish some cases.

- Suppose that $\mathcal{F}_{i+1}$ contains a pair (α,β) with $\{\alpha,\beta\} \subseteq \{-1,0\}$. By Lemma 5, we have that $\mathcal{F}_{i+1}$ contains only one such a pair, and that $\alpha \leq \alpha'$ and $\beta \leq \beta'$. Then, in φ, we can replace the value $\mu_{i+1} = \alpha'$ and $\nu_{i+1} = \beta'$ with $\mu_{i+1} = \alpha$ and $\nu_{i+1} = \beta$, leaving unaltered all other values and the in-out assignment. The resulting set of values, together with the in-out assignment, still satisfies Properties $\mathcal{P}_1$–$\mathcal{P}_7$ of Lemma 4. Indeed, Properties $\mathcal{P}_1$, $\mathcal{P}_2$, and $\mathcal{P}_7$ are not affected by the change of value for μ_{i+1} and ν_{i+1} (note that $i \leq k-2$). Property $\mathcal{P}_3$ is satisfied as $(\alpha,\beta) \in \mathcal{F}_{i+1}$. The inequalities $\mu_{i+1} \leq \rho_i$, $\nu_{i+1} \leq \rho_{i+1}$, $\mu_{i+1} \leq -\rho_i$, $\nu_{i+1} \leq -\rho_{i+1}$, $\nu_i + \mu_{i+1} \leq 0$, and $\nu_{i+1} + \mu_{i+2} \leq 0$ of Properties $\mathcal{P}_4$ and $\mathcal{P}_5$ are satisfied with $\mu_{i+1} = \alpha'$ and $\nu_{i+1} = \beta'$, hence they are also satisfied with $\mu_{i+1} = \alpha \leq \alpha'$ and $\nu_{i+1} = \beta \leq \beta'$. Finally, if $\beta = \nu_{i+1} = \mu_{i+2} = 0$, then Property $\mathcal{P}_5$ and $\beta \leq \beta'$ imply that β' is also equal to 0, hence Property $\mathcal{P}_6$ is satisfied by the new values as it is satisfied by φ.
- Suppose that the previous case does not apply and that $\mathcal{F}_{i+1}$ contains the pair $(1,-1)$. Then, in φ, we can replace the value $\mu_{i+1} = \alpha'$ and $\nu_{i+1} = \beta'$ with $\mu_{i+1} = 1$ and $\nu_{i+1} = -1$, leaving unaltered all other values and the in-out assignment, except for G_{i+1} which is now assigned to out – previously it might have been assigned to in or to out. The resulting set of values, together with the in-out assignment, satisfies Properties $\mathcal{P}_1$–$\mathcal{P}_7$ of Lemma 4. In particular, Properties $\mathcal{P}_1$, $\mathcal{P}_2$, and $\mathcal{P}_7$ are not affected by the change of value for μ_{i+1} and ν_{i+1}, while Property $\mathcal{P}_3$ is satisfied as $(1,-1) \in \mathcal{F}_{i+1}$. Property $\mathcal{P}_4$ is satisfied since we have $\mu_{i+1} = 1 = -\rho_i$ and $\nu_{i+1} = -1 \leq -\rho_{i+1}$. Property $\mathcal{P}_5$ is satisfied since we have $\nu_i + \mu_{i+1} = -1 + 1 = 0$ and $\nu_{i+1} + \mu_{i+2} = -1 + \mu_{i+2} \leq 0$. Finally, Property $\mathcal{P}_6$ is vacuously satisfied, since $\mu_{i+1} \neq 0$ and $\nu_{i+1} \neq 0$.
- Suppose that the previous cases do not apply, that $\mathcal{F}_{i+1}$ contains the pair $(1,0)$, and that $\mathcal{F}_{i+1}$ does not contain the pair $(-1,1)$. Then, in φ, we can replace the value $\mu_{i+1} = \alpha'$ and $\nu_{i+1} = \beta'$ with $\mu_{i+1} = 1$ and $\nu_{i+1} = 0$, leaving unaltered all other values and the in-out assignment. The resulting set of values, together with the in-out assignment, satisfies Properties $\mathcal{P}_1$–$\mathcal{P}_7$ of Lemma 4. In particular, Properties $\mathcal{P}_1$, $\mathcal{P}_2$, and $\mathcal{P}_7$ are not affected by the change of value for μ_{i+1} and ν_{i+1}, while Property $\mathcal{P}_3$ is satisfied as $(1,0) \in \mathcal{F}_{i+1}$. We now discuss Property $\mathcal{P}_4$. Because of the assumptions of this case, we have that $\mathcal{F}_{i+1}$ contains $(1,0)$, might contain $(0,1)$ and $(1,1)$, and does not contain any other pair. Since φ satisfies Property $\mathcal{P}_3$, we have $\alpha' \geq 0$ and $\beta' \geq 0$. Since φ satisfies Property $\mathcal{P}_4$, we have $\mathcal{A}(G_{i+1}) = \text{out}$, hence $\mu_{i+1} = 1 = -\rho_i$ and $\nu_{i+1} = 0 \leq \beta' \leq -\rho_{i+1}$. Property $\mathcal{P}_5$ is satisfied since we have $\nu_i + \mu_{i+1} = -1 + 1 = 0$ and $\nu_{i+1} + \mu_{i+2} = 0 + \mu_{i+2} \leq \beta' + \mu_{i+2} \leq 0$, given that φ satisfies Property $\mathcal{P}_5$. Finally, $\mu_{i+1} \neq 0$; further, if $\beta = \nu_{i+1} = \mu_{i+2} = 0$, then Property $\mathcal{P}_5$ and $\beta \leq \beta'$ imply that β' is also equal to 0, hence Property $\mathcal{P}_6$ is satisfied as it is satisfied by φ.
- Suppose that the previous cases do not apply, that $\mathcal{F}_{i+1}$ contains the pair $(\alpha,1)$, with $\alpha \in \{-1,0\}$, and that $\mathcal{F}_{i+1}$ does not contain the pair $(1,0)$. By Lemma 5, we have that $\mathcal{F}_{i+1}$ contains only one such a pair, and that $\alpha \leq \alpha'$. Because of the assumptions of this case, we have that $\mathcal{F}_{i+1}$ contains $(\alpha,1)$, might contain $(1,1)$, and does not contain any other pair. Hence, we have $\beta' = 1$. Then, in φ, we can replace the value $\mu_{i+1} = \alpha'$ with $\mu_{i+1} = \alpha$, leaving unaltered all other values and the in-out assignment. Since μ_{i+1}

does not increase, the proof that the resulting set of values, together with the in-out assignment, satisfies Properties $\mathcal{P}_1$–$\mathcal{P}_7$ of Lemma 4 is the same as in the first case.

- Suppose that the previous cases do not apply, and that $\mathcal{F}_{i+1}$ contains the pairs $(1, 0)$ and $(-1, 1)$. Because of the assumptions of this case, we have that $\mathcal{F}_{i+1}$ might contain $(0, 1)$, $(1, 1)$, and does not contain any other pair. Hence, $\beta' \geq 0$. We further distinguish three sub-cases.
  - First, suppose that v_{i+1} is a non-switch vertex in $\mathcal{C}$. By Property $\mathcal{P}_1$, we have $\rho_{i+1} = 0$ in φ. By Property $\mathcal{P}_4$, we have that $\nu_{i+1} \leq 0$. Hence, we have $\alpha' = 1$ and $\beta' = 0$, thus we can pick the value 1 for μ_{i+1} without loss of generality.
  - Second, suppose that v_{i+1} is a switch vertex in $\mathcal{C}$, that G_{i+2} is defined, and that $\mathcal{F}_{i+2}$ contains a pair $(0, \cdot)$. By Lemma 5, we have that $\mathcal{F}_{i+2}$ does not contain a pair $(-1, \cdot)$, hence $\mu_{i+2} \geq 0$ in φ. By Property $\mathcal{P}_5$, we have that $\nu_{i+1} \leq 0$. Hence, we have $\alpha' = 1$ and $\beta' = 0$, thus we can pick the value 1 for μ_{i+1} without loss of generality.
  - Third, suppose that v_{i+1} is a switch vertex in $\mathcal{C}$, and that G_{i+2} is undefined or that $\mathcal{F}_{i+2}$ does not contain a pair $(0, \cdot)$. First, we observe that $\mu_{i+2} = -1$ in φ; indeed, $\mu_{i+2} = 1$, together with $\nu_{i+1} = \beta' \geq 0$, would violate Property $\mathcal{P}_5$ for φ. Then, in φ, we can replace the value $\mu_{i+1} = \alpha'$ and $\nu_{i+1} = \beta'$ with $\mu_{i+1} = -1$ and $\nu_{i+1} = 1$, leaving unaltered all other values and the in-out assignment. The resulting set of values, together with the in-out assignment, satisfies Properties $\mathcal{P}_1$–$\mathcal{P}_7$ of Lemma 4. In particular, Properties $\mathcal{P}_1$, $\mathcal{P}_2$, and $\mathcal{P}_7$ are not affected by the change of value for μ_{i+1} and ν_{i+1}, while Property $\mathcal{P}_3$ is satisfied as $(-1, 1) \in \mathcal{F}_{i+1}$. We now discuss Property $\mathcal{P}_4$. Since $\beta' \geq 0$, we have that $\rho_{i+1} = 1$ if $\mathcal{A}(G_{i+1}) = \text{in}$ and $\rho_{i+1} = -1$ if $\mathcal{A}(G_{i+1}) = \text{out}$; since $\mathcal{A}(G_{i+1})$ does not change, we have that the new value $\nu_{i+1} = 1$ is equal to ρ_{i+1} if $\mathcal{A}(G_{i+1}) = \text{in}$ and to $-\rho_{i+1}$ if $\mathcal{A}(G_{i+1}) = \text{out}$. Property $\mathcal{P}_5$ is satisfied since we have $\nu_i + \mu_{i+1} = -1 - 1 < 0$ and $\nu_{i+1} + \mu_{i+2} = 1 - 1 = 0$. Finally, Property $\mathcal{P}_6$ is vacuously satisfied, since $\mu_{i+1} \neq 0$ and $\nu_{i+1} \neq 0$.
- Finally, if neither of the previous cases applies, $\mathcal{F}_{i+1}$ can only contain the pair $(1, 1)$, and in fact it has to contain it, given that $\mathcal{F}_{i+1}$ is not empty. By Property $\mathcal{P}_3$, we can pick the value 1 for μ_{i+1} without loss of generality.

This concludes the proof of the lemma. □

By means of Lemmata 6 and 7, we can assume to have decided without loss of generality the values $\mu_1, \ldots, \mu_k, \nu_1, \ldots, \nu_k$. Actually, when the value of ν_{k-1} is decided, we also need to check whether Property $\mathcal{P}_5$ of Lemma 4 is satisfied by ν_{k-1} and μ_k, as the value of μ_k was decided when fixing the current candidate tuple. The procedure described in Lemmata 6 and 7 takes $\mathcal{O}(k)$ time, namely $\mathcal{O}(1)$ time per fixed value. By Property $\mathcal{P}_1$, we can also fix the value $\rho_i = 0$ for every vertex v_i which is a non-switch vertex in $\mathcal{C}$. It remains to fix the value ρ_i for every vertex v_i which is a switch vertex in $\mathcal{C}$ (to -1 or $+1$, again by Property $\mathcal{P}_1$), and to decide the in-out assignment $\mathcal{A}$. The in-out assignments of some components are bound together, as in the following lemma.

Lemma 8. *Suppose that, for some $i \in \{1, \ldots, k-1\}$, the values ν_i and μ_{i+1} have been both fixed to 0. Then, if v_i is a switch vertex in $\mathcal{C}$ we have $\mathcal{A}(G_i) = \mathcal{A}(G_{i+1})$, otherwise we have $\mathcal{A}(G_i) \neq \mathcal{A}(G_{i+1})$.*

Proof. The first part of the statement comes from the fact that if v_i is a switch vertex in $\mathcal{C}$, then $\{\rho_i, -\rho_i\} = \{-1, 1\}$, hence Property $\mathcal{P}_4$ of Lemma 4 is violated if G_i and G_{i+1} are assigned to different sides of $\mathcal{C}$. The second part of the statement comes from the fact that if v_i is a non-switch vertex in $\mathcal{C}$, then $\rho_i = -\rho_i = 0$, hence Property $\mathcal{P}_6$ is violated if G_i and G_{i+1} are assigned to the same side of $\mathcal{C}$. □

If we consider a maximal sequence of components $G_p, G_{p+1}, \ldots, G_q$ such that any two consecutive components are bound together, deciding whether G_p is embedded inside or outside $\mathcal{C}$ decides whether each of the components $G_p, G_{p+1}, \ldots, G_q$ is embedded inside or outside $\mathcal{C}$, via repeated applications of Lemma 8.

Furthermore, the values $\rho_p, \rho_{p+1}, \ldots, \rho_{q-1}$ are also determined by the same decision, by Property $\mathcal{P}_6$: If G_i and G_{i+1} both lie inside $\mathcal{C}$, then $\rho_i = 1$, while if they both lie outside $\mathcal{C}$, then $\rho_i = -1$ (while if they lie one inside and one outside $\mathcal{C}$, then v_i is a switch vertex in $\mathcal{C}$ and ρ_i has been already fixed to 0). The values ρ_{p-1} and ρ_q might also be determined by the choice of whether G_p is embedded inside or outside $\mathcal{C}$, if $\mu_p \geq 0$ and if $\nu_q \geq 0$, respectively, by Property $\mathcal{P}_4$. Thus, each maximal sequence of components $G_p, G_{p+1}, \ldots, G_q$ such that any two consecutive components are bound together gives us a (not necessarily positive) integer, which is the sum of all the ρ_i's that are determined by the choice of whether G_p is embedded inside or outside $\mathcal{C}$. Note that distinct maximal sequences of components are independently embedded inside or outside $\mathcal{C}$. By Property $\mathcal{P}_2$, the sum of the positive integers that are obtained in this way has to coincide with the sum of the negative integers, up to a constant additive term. This is where Partition shows up to help.

We now refine this argument; see Fig. 7. We are going to define a multiset $\mathcal{S}$ of positive integers. Initially, we have $\mathcal{S} = \emptyset$. Consider every maximal sequence of components $G_p, G_{p+1}, \ldots, G_q$ such that, for $i = p, p+1, \ldots, q-1$, the values ν_i and μ_{i+1} have both been fixed to 0. Initialize an integer a_p to 0 and a variable c to in; this variable only assumes values in $\{\text{in}, \text{out}\}$. If $\mu_p >= 0$ and v_{p-1} is a switch vertex in $\mathcal{C}$, then increase a_p by 1 and *mark* v_{p-1} – thus taking note of the fact that the value ρ_{p-1} has already contributed to generate some integer. Then, for $i = p, \ldots, q-1$, proceed as follows. If v_i is a non-switch vertex in $\mathcal{C}$, then change the value of c. If v_i is a switch vertex in $\mathcal{C}$, then increase the value of a_p by 1 if $c = \text{in}$ and decrease the value of a_p by 1 if $c = \text{out}$; in both cases, mark v_i. Finally, if $\nu_q \geq 0$ and v_q is a switch vertex in $\mathcal{C}$, then increase the value of a_p by 1 if $c = \text{in}$ and decrease the value of a_p by 1 if $c = \text{out}$; in both cases, mark v_q. Insert into $\mathcal{S}$ the absolute value of a_p. Henceforth, this process inserts a positive integer into $\mathcal{S}$ for every maximal sequence of components $G_p, G_{p+1}, \ldots, G_q$ that are bound together – this is true also for sequences consisting

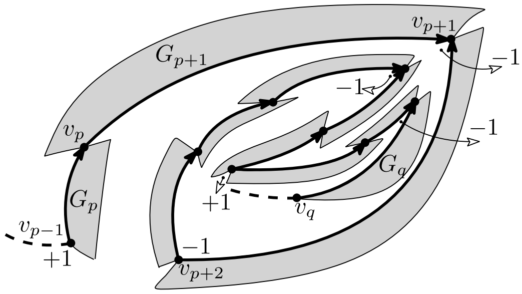

Fig. 7. Generation of an integer a_p from components $G_p, G_{p+1}, \dots, G_q$ that are bound together. The thick line is part of the cycle C. Subgraphs G_i are schematized. Contributions of the values ρ_i to a_p are shown. In this example, $a_p = -2$. Note that, in this example, the value of a_p is increased by one unit by ρ_{p-1}, since v_{p-1} is a switch vertex in C and $\mu_p \geq 0$, while the value of a_p is not altered by ρ_q, since $v_q = -1$.

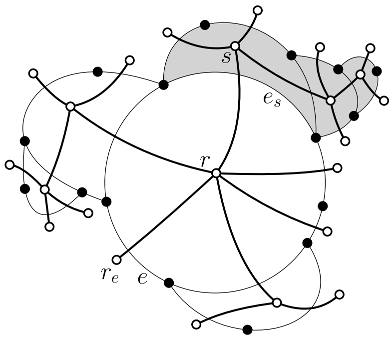

Fig. 8. The extended dual tree $\mathcal{T}$ of an outerplane embedding $\mathcal{O}$ of G; the tree $\mathcal{T}$ is represented by white disks and thick lines. The subgraph G_s of G corresponding to the subtree $\mathcal{T}_s$ of $\mathcal{T}$ rooted at a node s has its internal faces colored gray.

of a single component G_p. We also insert into $\mathcal{S}$ a number 1 for each unmarked switch vertex in C. In the resulting multiset $\mathcal{S}$ of integers, there are up to four *special* integers, the ones corresponding to v_0, G_1, v_k, and G_k. These are special because, in the current candidate tuple, the values of ρ_0 and ρ_k are already decided, as well as the in-out assignment for G_1 and G_k. The special integers might be less than four. For example, if v_0 is a non-switch vertex in C, it does not generate any integer; as another example, if v_0 is marked when considering G_1, then their corresponding integer is the same.

Let $2\ell \leq k$ be the number of switch vertices in C. Property $\mathcal{P}_2$ of Lemma 4 requires that the sum of the ρ_i's is equal to -2. Because of the above construction, this is possible if and only if the integers in $\mathcal{S}$ can be partitioned into two sets $\mathcal{S}_{\text{in}}$ and $\mathcal{S}_{\text{out}}$ such that the sum of the elements in $\mathcal{S}_{\text{in}}$ is $\ell - 1$ and the sum of the elements in $\mathcal{S}_{\text{out}}$ is $\ell + 1$. Special integers are removed from $\mathcal{S}$ and used to decrease the value $\ell - 1$ or $\ell + 1$. For example, if $\rho_0 = 1$ in the current candidate tuple and v_0 is not marked, then the integer $a_0 = 1$ is required to be assigned to $\mathcal{S}_{\text{in}}$, hence we delete it from $\mathcal{S}$ and we decrease the target sum for the elements in $\mathcal{S}_{\text{in}}$ to $\ell - 2$. After the $\mathcal{O}(1)$ special integers are considered, we are left with a set of integers $\mathcal{S}$ which needs to be partitioned into two sets $\mathcal{S}_{\text{in}}$ and $\mathcal{S}_{\text{out}}$ whose elements need to sum up to some values ℓ_{in} and ℓ_{out}, respectively. If $\ell_{\text{in}} \neq \ell_{\text{out}}$, this is not exactly an instance of PARTITION, hence consider a new integer $a^* = |\ell_{\text{in}} - \ell_{\text{out}}|$ and we aim to partition $\mathcal{S} \cup \{a^*\}$ into two sets so that the sum of the elements in each set is the same. Clearly, this is possible if and only if $\mathcal{S}$ can be partitioned into two sets $\mathcal{S}_{\text{in}}$ and $\mathcal{S}_{\text{out}}$ such that the sum of the elements in $\mathcal{S}_{\text{in}}$ is ℓ_{in} and the sum of the elements in $\mathcal{S}_{\text{out}}$ is ℓ_{out}. Thus, we can solve the problem using the $\mathcal{O}(f(\ell)) \in \mathcal{O}(f(k))$ algorithm for PARTITION.

The algorithm described in this section gives us the following.

Lemma 9. *Assume that the feasible set $\mathcal{F}_i$ of each subgraph G_i of G is known. Then the feasible set $\mathcal{F}$ of G can be computed in $\mathcal{O}(f(k))$ time.*

With Lemma 9 at hand, the algorithm to test whether G admits an upward plane embedding with a given edge e on the outer face, where e is incident to $f_{\mathcal{O}}$, follows a natural recursive strategy described below.

First, we compute the outerplane embedding $\mathcal{O}$ of G in $\mathcal{O}(n)$ time [10,19,29,35]. Then we construct the *extended dual tree* $\mathcal{T}$ of $\mathcal{O}$; see Fig. 8. This has a vertex s for each cycle C_s delimiting an internal face of $\mathcal{O}$ and has a vertex for each edge incident to $f_{\mathcal{O}}$. Two vertices s and t of $\mathcal{T}$ corresponding to cycles C_s and C_t are adjacent in $\mathcal{T}$ if C_s and C_t share an edge, and a vertex of $\mathcal{T}$ corresponding to an edge x is only adjacent to the vertex of $\mathcal{T}$ corresponding to the unique cycle delimiting an internal face of $\mathcal{O}$ the edge x belongs to. We root $\mathcal{T}$ at the leaf r_e corresponding to e. Then the subtree $\mathcal{T}_s$ of $\mathcal{T}$ rooted at an internal node s of $\mathcal{T}$ corresponds to a subgraph G_s of G, which is the union of the cycles corresponding to the vertices in $\mathcal{T}_s$.

We now perform a bottom-up visit of $\mathcal{T}$ which ends after the child r of r_e is visited. When processing an internal node s of $\mathcal{T}$ with parent p, we apply Lemma 9 in order to compute the feasible set of G_s. We denote this feasible set by $\mathcal{F}_s$ and remark that the (μ, ν)-embeddings of G_s are such that the edge e_s shared by C_s and C_p – this is the edge e if $s = r$ – is incident to the outer face. By Lemma 9, it takes $\mathcal{O}(f(k_s))$ time to process s, where k_s is the number of vertices of C_s. This sums up to $\mathcal{O}(f(n))$ time over the entire bottom-up visit. If some computed feasible set $\mathcal{F}_s$ is empty, then by Property $\mathcal{P}_3$ of Lemma 4 the feasible set of every node in the path from s to r in $\mathcal{T}$ is empty, hence G admits no upward plane embedding in which e is incident to the outer face. Conversely, if a non-empty feasible set is computed for the child r of the root r_e, then it is concluded that G admits an upward plane embedding in which e is incident to the outer face. The described algorithm can also be suitably modified so that, in case G admits an upward plane embedding in which e is incident to the outer face, it constructs such an embedding. In order to do that, whenever we add a pair (μ, ν) to a feasible set $\mathcal{F}_s$ during the bottom-up visit of $\mathcal{T}$, we also associate to the pair the $\mathcal{O}(k_s)$ values $\rho_0, \ldots, \rho_k, \mu_1, \ldots, \mu_k, \nu_1, \ldots, \nu_k$ and the in-out assignment $\mathcal{A}$ that satisfy Properties $\mathcal{P}_1$–$\mathcal{P}_7$ of Lemma 4. This information can be then exploited in a top-down traversal of $\mathcal{T}$ in order to compute a (μ, ν)-embedding of G. This concludes the proof of Theorem 2.

5. Conclusions

In this paper, we have established a connection between the problem of testing the upward planarity of an n-vertex biconnected outerplanar DAG and a famous complexity theory problem, called Partition, which asks whether a set of positive integers admits a partition into two sets so that the sum of the elements in each set is equal to $n/2$. Our main result is that, if the upward planarity of an n-vertex biconnected outerplanar DAG can be tested in $f(n)$ time, then Partition can also be solved in $\mathcal{O}(f(n))$ time. The time complexity of Partition, when expressed solely as a function of n, is known to be between $\Omega(n)$ and $\mathcal{O}(n \log^2 n)$. Any super-linear lower bound for the Partition problem would hence imply a super-linear lower bound for testing the upward planarity of an n-vertex biconnected outerplanar DAG; this would be, to the best of our knowledge, the first super-linear lower bound for a graph drawing problem not involving real numbers. Conversely, a linear-time algorithm for testing the upward planarity of an n-vertex biconnected outerplanar DAG would imply a linear-time algorithm for Partition, which would be a result improving upon the outcomes of decades of research.

We have also shown a result in the other direction. Namely, suppose that Partition can be solved in $f(n)$ time. Let G be an n-vertex biconnected outerplanar DAG and e be a prescribed edge incident to the outer face of the outerplane embedding of G. Then it can be tested in $\mathcal{O}(f(n))$ time whether G with admits an upward planar drawing in which e is incident to the outer face. A result of Koiliaris and Xu [25] implies that such a function $f(n)$ exists with $f(n) \in \mathcal{O}(n \log^2 n)$. On the one hand, the constraint that e is incident to the outer face of the outerplane embedding $\mathcal{O}$ of G is not a loss of generality per se, as we can prove that if a biconnected outerplanar DAG admits an upward plane embedding, it also admits an upward plane embedding $\mathcal{E}$ in which an edge incident to the outer face of $\mathcal{O}$ is incident to the outer face of $\mathcal{E}$. On the other hand, it would be nice to remove entirely the constraint of having a prescribed edge incident to the outer face and, possibly, to remove the assumption of biconnectivity. It is often the case that an algorithm to test some planarity variant (e.g., upward planarity or rectilinear planarity) with the constraint of having a prescribed edge that is required to be incident to the outer face can be used as a subroutine by an algorithm to solve the same problem, within the same time bound and without the constraint on the prescribed edge. A general methodology for doing that has been recently introduced [14] and already used successfully for problems related to ours [9,15]. However, the techniques introduced in [14] do not seem to apply straight-forwardly to extend our algorithm so to deal with general biconnected and connected outerplanar DAGs.

We conclude by remarking that our paper shows a substantial difference between rectilinear planarity testing and upward planarity testing, two problems that are often tied together. Namely, while it can be tested in linear time whether an outerplanar graph admits a rectilinear planar drawing [15], the results of Section 3 provide evidence that finding a linear-time algorithm for testing the upward planarity of an outerplanar DAG might be an elusive goal. Interestingly, both the linear-time algorithm for testing rectilinear planarity [15] and the algorithm we proposed in Section 4 exploit a bottom-up approach on the extended dual tree, in which the computation at each node of the tree mainly consists of deciding whether each graph component that is attached to the cycle corresponding to the node should be embedded inside or outside such a cycle. However, this choice has a different nature in the two problems. Indeed, in the rectilinear planarity testing problem, choices performed for distinct components are independent from one another. Conversely, in upward planarity testing, the constraints imposed by Theorem 4 bound together some components in the inside/outside decision. This bond between components generates integers: Each integer corresponds to the number of times the cycle has to "roll up" in order to accommodate components that are bound together and that all lie inside (or all lie outside) the cycle. Theorem 4 requires to balance the sum of the integers corresponding to components lying inside the cycle with the sum of the integers corresponding to components lying outside the cycle. This is where the connection to Partition arises, which makes upward planarity seemingly harder to test than rectilinear planarity.

CRediT authorship contribution statement

Fabrizio Frati: Conceptualization, Formal analysis, Funding acquisition, Investigation, Methodology, Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

References

[1] P. Bertolazzi, G. Di Battista, G. Liotta, C. Mannino, Upward drawings of triconnected digraphs, Algorithmica 12 (6) (1994) 476–497.
[2] P. Bertolazzi, G. Di Battista, C. Mannino, R. Tamassia, Optimal upward planarity testing of single-source digraphs, SIAM J. Comput. 27 (1) (1998) 132–169.
[3] T. Bläsius, S. Lehmann, I. Rutter, Orthogonal graph drawing with inflexible edges, Comput. Geom. 55 (2016) 26–40.
[4] G. Borradaile, P.N. Klein, S. Mozes, Y. Nussbaum, C. Wulff-Nilsen, Multiple-source multiple-sink maximum flow in directed planar graphs in near-linear time, SIAM J. Comput. 46 (4) (2017) 1280–1303.
[5] K. Bringmann, A near-linear pseudopolynomial time algorithm for subset sum, in: P.N. Klein (Ed.), Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16–19, SIAM, 2017, pp. 1073–1084.
[6] G. Brückner, M. Himmel, I. Rutter, An SPQR-tree-like embedding representation for upward planarity, in: D. Archambault, C.D. Tóth (Eds.), 27th International Symposium on Graph Drawing and Network Visualization, GD 2019, in: Lecture Notes in Computer Science, vol. 11904, Springer, 2019, pp. 517–531.
[7] H.Y. Chan, A parameterized algorithm for upward planarity testing, in: S. Albers, T. Radzik (Eds.), 12th Annual European Symposium on Algorithms, ESA 2004, in: Lecture Notes in Computer Science, vol. 3221, Springer, 2004, pp. 157–168.
[8] S. Chaplick, E. Di Giacomo, F. Frati, R. Ganian, C.N. Raftopoulou, K. Simonov, Parameterized algorithms for upward planarity, in: X. Goaoc, M. Kerber (Eds.), 38th International Symposium on Computational Geometry (SoCG 2022), in: LIPIcs, vol. 224, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, 26.
[9] S. Chaplick, E. Di Giacomo, F. Frati, R. Ganian, C.N. Raftopoulou, K. Simonov, Testing upward planarity of partial 2-trees, in: P. Angelini, R. von Hanxleden (Eds.), 30th International Symposium on Graph Drawing and Network Visualization (GD 2022), in: LNCS, vol. 13764, Springer, 2022, pp. 175–187.
[10] N. Chiba, T. Nishizeki, S. Abe, T. Ozawa, A linear algorithm for embedding planar graphs using PQ-trees, J. Comput. Syst. Sci. 30 (1) (1985) 54–76.
[11] G. Di Battista, G. Liotta, F. Vargiu, Spirality and optimal orthogonal drawings, SIAM J. Comput. 27 (6) (1998) 1764–1811.
[12] E. Di Giacomo, G. Liotta, F. Montecchiani, Orthogonal planarity testing of bounded treewidth graphs, J. Comput. Syst. Sci. 125 (2022) 129–148.
[13] W. Didimo, F. Giordano, G. Liotta, Upward spirality and upward planarity testing, SIAM J. Discrete Math. 23 (4) (2009) 1842–1899.
[14] W. Didimo, G. Liotta, G. Ortali, M. Patrignani, Optimal orthogonal drawings of planar 3-graphs in linear time, in: S. Chawla (Ed.), SODA '20, SIAM, 2020, pp. 806–825.
[15] F. Frati, Planar rectilinear drawings of outerplanar graphs in linear time, Comput. Geom. 103 (2022) 101854.
[16] M.R. Garey, D.S. Johnson, Computers and Intractability: A Guide to the Theory of NP-Completeness, W. H. Freeman, 1979.
[17] A. Garg, R. Tamassia, On the computational complexity of upward and rectilinear planarity testing, SIAM J. Comput. 31 (2) (2001) 601–625.
[18] P. Healy, K. Lynch, Two fixed-parameter tractable algorithms for testing upward planarity, Int. J. Found. Comput. Sci. 17 (5) (2006) 1095–1114.
[19] J.E. Hopcroft, R.E. Tarjan, Efficient planarity testing, J. ACM 21 (4) (1974) 549–568.
[20] E. Horowitz, S. Sahni, Computing partitions with applications to the knapsack problem, J. ACM 21 (2) (1974) 277–292.
[21] M.D. Hutton, A. Lubiw, Upward planar drawing of single source acyclic digraphs, in: A. Aggarwal (Ed.), 2nd Annual ACM/SIGACT-SIAM Symposium on Discrete Algorithms, SODA 1991, ACM/SIAM, 1991, pp. 203–211.
[22] M.D. Hutton, A. Lubiw, Upward planar drawing of single-source acyclic digraphs, SIAM J. Comput. 25 (2) (1996) 291–311.
[23] N. Karmarkar, R. Karp, The differencing method of set partitioning, Tech. rep., Computer Science Division (EECS), University of California, 1982.
[24] H. Kellerer, U. Pferschy, D. Pisinger, Knapsack Problems, Springer, 2004.
[25] K. Koiliaris, C. Xu, Faster pseudopolynomial time algorithms for subset sum, ACM Trans. Algorithms 15 (3) (2019) 40.
[26] R.E. Korf, A complete anytime algorithm for number partitioning, Artif. Intell. 106 (2) (1998) 181–203.
[27] R.E. Korf, A hybrid recursive multi-way number partitioning algorithm, in: T. Walsh (Ed.), 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011), IJCAI/AAAI, 2011, pp. 591–596.
[28] R.E. Korf, E.L. Schreiber, Optimally scheduling small numbers of identical parallel machines, in: D. Borrajo, S. Kambhampati, A. Oddi, S. Fratini (Eds.), 23rd International Conference on Automated Planning and Scheduling, (ICAPS 2013), AAAI, 2013.
[29] S.L. Mitchell, Linear algorithms to recognize outerplanar and maximal outerplanar graphs, Inf. Process. Lett. 9 (5) (1979) 229–232.
[30] A. Papakostas, Upward planarity testing of outerplanar dags, in: R. Tamassia, I.G. Tollis (Eds.), DIMACS International Workshop on Graph Drawing, GD '94, in: Lecture Notes in Computer Science, vol. 894, Springer, 1994, pp. 298–306.
[31] D. Pisinger, Linear time algorithms for knapsack problems with bounded weights, J. Algorithms 33 (1) (1999) 1–14.
[32] E.L. Schreiber, Optimal Multi-Way Number Partitioning, Ph.D. thesis, University of California, Los Angeles, USA, 2014.
[33] R. Schroeppel, A. Shamir, A $T = O(2^{n/2})$, $S = O(2^{n/4})$ algorithm for certain NP-complete problems, SIAM J. Comput. 10 (3) (1981) 456–464.
[34] R. Tamassia, On embedding a graph in the grid with the minimum number of bends, SIAM J. Comput. 16 (3) (1987) 421–444.
[35] M. Wiegers, Recognizing outerplanar graphs in linear time, in: G. Tinhofer, G. Schmidt (Eds.), WG '86, in: LNCS, vol. 246, Springer, 1987, pp. 165–176.