Journal of Graph Algorithms and Applications
http://jgaa.info/ vol. 28, no. 1, pp. 51–0 (2024)
DOI: 10.7155/jgaa.v28i1.2927

The role of twins in computing planar supports of hypergraphs

René van Bevern [1] *Iyad Kanj* [2] *Christian Komusiewicz* [3]
Rolf Niedermeier [4] *Manuel Sorge* [5]

[1] Huawei Technologies Co., Ltd.
[2] DePaul University, Chicago, USA
[3] Friedrich Schiller University Jena, Jena, Germany
[4] Algorithmics and Computational Complexity, Faculty IV, TU Berlin, Berlin, Germany
[5] Institute of Logic and Computation, TU Wien, Vienna, Austria

Submitted: July 2022	Accepted: February 2024	Published: May 2024
Article type: Regular paper		Communicated by: C. Tóth

Abstract. A *support* or *realization* of a hypergraph $\mathcal{H}$ is a graph G on the same vertex set as $\mathcal{H}$ such that for each hyperedge of $\mathcal{H}$ it holds that its vertices induce a connected subgraph of G. The NP-hard problem of finding a *planar* support has applications in hypergraph drawing and network design. Previous algorithms for the problem assume that *twins*—pairs of vertices that are in precisely the same hyperedges—can safely be removed from the input hypergraph. We prove that this assumption is generally wrong, yet that the number of twins necessary for a hypergraph to have a planar support only depends on its number of hyperedges. We give an explicit upper bound on the number of twins necessary for a hypergraph with m hyperedges to have an r-outerplanar support, which depends only on r and m. Since all additional twins can be safely removed, we obtain a linear-time algorithm for computing r-outerplanar supports for hypergraphs with m hyperedges if m and r are constant; in other words, the problem is fixed-parameter linear-time solvable with respect to the parameters m and r.

Keywords: Subdivision drawings, NP-hard problem, r-outerplanar graphs, sphere-cut branch decomposition

An extended abstract of this work appeared at GD'16 [38]. This version provides full proof details, more illustrative figures, and a simpler example for the necessity of twins with, additionally, smallest-possible hyperedge size. Results reported in this article are unrelated to René van Bevern's work at Huawei.

E-mail addresses: rene.van.bevern@huawei.com (René van Bevern) ikanj@cs.depaul.edu (Iyad Kanj) c.komusiewicz@uni-jena.de (Christian Komusiewicz) rolf.niedermeier@tu-berlin.de (Rolf Niedermeier) manuel.sorge@ac.tuwien.ac.at (Manuel Sorge)

Figure 1: Two drawings of the same hypergraph. On the left, we see a drawing in the *subset standard* in which the vertices (white circles) are enclosed by curves that correspond to hyperedges. On the right, we see a *subdivision drawing* in which we assign vertices to regions (enclosed by black lines) and we color these regions with colors that one-to-one correspond to the hyperedges; for each hyperedge, the union of the regions of the vertices in that hyperedge is connected.

1 Introduction

Hypergraph drawings are useful as visual aid in diverse applications [1], among them electronic circuit design [15, 16] and relational databases [33, 4]. This led to several generalizations of the concept of planarity from graphs to hypergraphs. The earliest among them is the attempt of Zykov [41], who defined a hypergraph to be planar if its incidence graph is.[1] This is equivalent to the requirement that one can draw the hyperedges as closed regions in such a way that each intersection of hyperedges contains exactly one vertex [41]. Voloshina and Feinberg [40] introduced *planar realizations* (for an English reproduction of the results refer to the book of Feinberg, Levin, and Rabinovich [16]) which nowadays are better known as *planar supports* [24, 25]: a *support* for a hypergraph $\mathcal{H} = (V, \mathcal{E})$ is a graph G on the same vertex set as $\mathcal{H}$ such that each hyperedge $F \in \mathcal{E}$ induces a connected subgraph $G[F]$. This is a generalization of planarity: an ordinary graph is planar if and only if it has a planar support when viewed as a hypergraph.

We study the NP-complete [3, 24] problem of recognizing hypergraphs that allow for a planar support. These are exactly the hypergraphs allowing for a *subdivision drawing* [24, 25]: given a hypergraph $\mathcal{H}$, we divide the plane into closed regions that one-to-one correspond to the vertices of $\mathcal{H}$ in such a way that, for each hyperedge F, the union of the regions corresponding to the vertices in F is connected. Subdivision drawings have also been called *vertex-based Venn diagrams* [24]. Figure 1 shows an example for such a drawing.

Having a planar support, or equivalently, a subdivision drawing, is a rather general concept of planar embeddings: for example, each hypergraph that has a planar incidence graph or a well-formed Euler diagram [18] has a planar support. Still, in the same way that most ordinary graphs are not planar, most hypergraphs do not have planar supports [16]. Actually *finding* them might be even more complicated by the fact that several works on planar supports assume that the input hypergraph is twinless, that is, there are no two vertices contained in precisely the same hyperedges (see Mäkinen [33, p. 179], Buchin et al. [9, p. 535], and Kaufmann, van Kreveld, and Speckmann [25, p. 399]). Twins do not seem useful at first glance: whatever role one vertex can play to obtain a planar support, its twin can also fulfill. One of our contributions is disproving the general validity of this assumption in Section 3. More specifically, we give a hypergraph with two twins that has a planar support but after removing one twin it ceases to have one. Thus, twins may be crucial to allow for a planar support.

More generally, we can construct hypergraphs with ℓ twins that allow for a planar support but cease doing so when removing one of the twins. However, the number of hyperedges in the construction grows with ℓ. It is thus natural to ask whether there is a function $\psi\colon \mathbb{N} \to \mathbb{N}$ such that, in each

[1]The *incidence graph* of a hypergraph $\mathcal{H} = (V, \mathcal{E})$ has as vertex set $V \cup \mathcal{E}$ and as edgeset all unordered pairs $\{v, F\}$ of a hypergraph vertex $v \in V$ and a hyperedge $F \in \mathcal{E}$ such that $v \in F$.

hypergraph with m hyperedges, we can forget all but $\psi(m)$ twins while maintaining the property of having a planar support. Using well-quasi orderings, one can prove the *existence* of such a function ψ (see Section 4), yet finding a closed form for ψ turned out to be surprisingly difficult: so far we could only compute a concrete upper bound when considering a second parameter, the outerplanarity number r of the desired planar support. A graph is r-*outerplanar* if it admits a planar embedding (without edge crossings) which has the property that, after r times of removing all vertices on the outer face, we obtain an empty graph. The outerplanarity number r of the support roughly translates to the number of layers in a corresponding drawing which can be seen by examining the construction of a subdivision drawing from a support given by Kaufmann, Kreveld, and Speckmann [25, p. 401] for example. Formally, we study the following problem (special cases of which were also considered previously [7, 9]):

> r-Outerplanar Support
> *Input:* A connected hypergraph $\mathcal{H}$ with n vertices and m hyperedges, and $r \in \mathbb{N}$.
> *Question:* Does $\mathcal{H}$ admit an r-outerplanar support?

Herein, a hypergraph is connected if for every pair of vertices u, v there is an alternating sequence of vertices and hyperedges that begins with u and ends with v such that successive elements are incident with each other. This assumption helps avoiding edge cases; our results easily extend to the non-connected case.

Our main result is a concrete upper bound on the number $\psi(m, r)$ of twins that might be necessary to obtain an r-outerplanar support. Since superfluous twins can then be removed in linear time, this gives the following algorithmic result.

Theorem 1 *There is an algorithm solving* r-Outerplanar Support *which, for constant r and m, has linear running time.*

To put Theorem 1 into perspective, r-Outerplanar Support remains NP-complete for $r = \infty$ [3, 24] and even for every fixed $r > 1$ [9] (see below). The constants in the running time of the algorithm in Theorem 1 have a large dependence on m and r. However, as also the discussion about twins above shows, the number of hyperedges is a natural parameter whose influence on the complexity is interesting to know. Furthermore, it is conceivable that the parameters m and r are small in practical instances: for a large number m of hyperedges, it is plausible that we obtain only hardly legible drawings unless the hyperedges adhere to some special structure, whereas every hypergraph with at most eight hyperedges has a planar support [16, 34]. Thus, it makes sense to design algorithms particularly for hypergraphs with a small number of hyperedges, as done by Verroust and Viaud and Hurtado et al. [39, 23]. Moreover, a small outerplanarity number r leads to few layers in the drawing which may lead to aesthetically pleasing drawings, similarly to path- or cycle-supports [9].

Related work. For specifics on the relations of different planar hypergraph embeddings, see Brandes et al. [7], Feinberg, Levin, and Rabinovich [16], and Kaufmann, Kreveld, and Speckmann [25].

Azarenok and Sarvanov and Johnson and Pollak [3, 24] showed that finding a planar support is NP-complete. Buchin et al. [9] proved that r-Outerplanar Support is NP-complete for $r = 2, 3$. From their proof it follows that r-Outerplanar Support is also NP-complete for every $r > 3$. This is due to a property of the reduction that Buchin et al. [9] use: Given a formula ϕ in 3CNF, they construct a hypergraph $\mathcal{H}$ that has a *planar* support if and only if ϕ is satisfiable. Due to the way in which $\mathcal{H}$ is constructed, if there is any planar support, then it is 3-outerplanar. Thus, deciding whether there is an r-outerplanar support for $r \geq 3$ also decides the satisfiability of the corresponding formula.

Towards determining the computational complexity of finding an outerplanar hypergraph support, Brandes et al. [7] gave a polynomial-time algorithm for cactus supports (graphs in which each edge is contained in at most one cycle). They also showed that finding an outerplanar support (or planar support) can be done in polynomial time if, in the input hypergraph, each intersection or difference of two hyperedges is empty, a singleton, or again a hyperedge in the hypergraph. A tree support can even be found in linear time [4, 37], a tree support of minimum diameter can be found in polynomial time [31], and one can deal with an additional upper bound on the vertex degrees in the tree support in polynomial time [9]. Klemz, Mchedlidze, and Nöllenburg [26] studied so-called area-proportional Euler diagrams, for which the corresponding computational problem reduces to finding a minimum-weight tree support. Such supports can also be found in polynomial time [26, 27]. Furthermore, if there are only two hyperedges and the positions of the vertices in the embedding are specified, then checking for a planar support can be done in polynomial time [23]

In a wider scope, motivated by drawing metro maps and metro map-like diagrams, Brandes et al. [8] studied the problem of finding *path-based* planar hypergraph supports (these are planar supports that fulfill the additional constraint that the subgraph induced by each hyperedge contains a Hamiltonian path) giving NP-hardness and tractability results. Path-based tree supports can also be found in polynomial time [36].

A concept related to subdivision drawings is (overlapping) clustered planarity [12, 2]. Very roughly, a graph G together with a hypergraph $\mathcal{H}$ on the same vertex set is *overlapping clustered planar* if G and $\mathcal{H}$ admit a joint embedding in the plane which is edge-crossing-free for G, a subdivision drawing for $\mathcal{H}$, and no edge of G crosses twice the boundary of a region corresponding to a hyperedge in $\mathcal{H}$. Overlapping clustered planarity is a generalization of clustered planarity; In the latter, one assumes that each pair of hyperedges in the hypergraph $\mathcal{H}$ is either disjoint or one hyperedge is a subset of the other [17]. Clustered planarity has attracted a lot of research interest, see Da Lozzo et al. [29] for a recent overview of the literature. Only recently, a polynomial-time algorithm for testing clustered planarity has been found [19].

Voloshina and Feinberg [40] suggested a data reduction rule for finding planar supports that removes more than just twins: it keeps only one vertex out of each nonempty inclusion-minimal intersection of any number of hyperedges. They prove that the input hypergraph has a planar support if the reduced hypergraph has, but that the reverse direction does generally not hold (Feinberg, Levin, and Rabinovich [16] show an example).

Chen et al. [10] showed that for obtaining minimum-edge supports (not necessarily planar), twins show a similar behavior as for r-outerplanar supports: Removing a twin can increase the minimum number of edges needed for a support and finding a minimum-edge support is linear-time solvable for a constant number of hyperedges via removing superfluous twins.

Organization. In Section 2 we provide some technical preliminaries used throughout the work. In Section 3 we give an example that shows that twins can be crucial for a hypergraph to have a planar support. As mentioned, for each $m \in \mathbb{N}$, there is a number $\psi(m)$ such that in each hypergraph with a planar support we can safely forget all but $\psi(m)$ twins (see Section 4). In Section 5 we give a concrete upper bound for $\psi(m)$ in the case of r-outerplanar supports and derive the linear-time algorithm for r-Outerplanar Support claimed in Theorem 1. We base the proof on a construction of a special sequence of nested separators in r-outerplanar graphs which is given in Section 6. We conclude and give some directions for future research in Section 7.

2 Preliminaries

By $A \uplus B$ we denote the union of two disjoint sets A and B. For a family of sets $\mathcal{F}$, we write $\bigcup \mathcal{F}$ in place of $\bigcup_{S \in \mathcal{F}} S$. For equivalence relations ρ over some set S and $v \in S$, we use $[v]_\rho$ to denote the equivalence class of v in ρ.

Hypergraphs. A *hypergraph* $\mathcal{H}$ is a pair $(V, \mathcal{E})$ consisting of a *vertex set* V, also denoted $V(\mathcal{H})$, and a *hyperedge set* $\mathcal{E}$, also denoted $\mathcal{E}(\mathcal{H})$. The hyperedge set $\mathcal{E}$ is a family of subsets of V, that is, $F \subseteq V$ for every hyperedge $F \in \mathcal{E}$. Where it is not ambiguous, we use $n := |V|$ and $m := |\mathcal{E}|$. When specifying running times, we use $|\mathcal{H}|$ to denote $|V(\mathcal{H})| + \sum_{F \in \mathcal{E}(\mathcal{H})} |F|$. The *size* $|F$— of a hyperedge F is the number of vertices in it. Unless stated otherwise, we assume that hypergraphs do not contain hyperedges of size at most one or multiple copies of the same hyperedge. (These do not play any role for the problem under consideration, and removing them can be done easily and efficiently.)

A vertex $v \in V$ and a hyperedge $F \in \mathcal{E}$ are *incident* with one another if $v \in F$. For a vertex $v \in V(\mathcal{H})$, let $\mathcal{E}_{\mathcal{H}}(v) := \{F \in \mathcal{H} \mid v \in F\}$. If it is not ambiguous, then we omit the subscript $\mathcal{H}$ from $\mathcal{E}_{\mathcal{H}}$. A vertex v *covers* a vertex u if $\mathcal{E}(u) \subseteq \mathcal{E}(v)$. Two vertices $u, v \in V$ are *twins* if $\mathcal{E}(v) = \mathcal{E}(u)$. Clearly, the relation τ on V defined by $\forall u, v \in V \colon (u, v) \in \tau \iff \mathcal{E}(u) = \mathcal{E}(v)$ is an equivalence relation. The equivalence classes $[u]_\tau$ for $u \in V$ are called *twin classes*.

Removing a vertex subset $S \subseteq V(\mathcal{H})$ from a hypergraph $\mathcal{H} = (V, \mathcal{E})$ results in the hypergraph $\mathcal{H} - S := (V \setminus S, \mathcal{E}')$, where $\mathcal{E}'$ is obtained from $\{F \setminus S \mid F \in \mathcal{E}\}$ by removing empty and singleton sets. For brevity, we also write $\mathcal{H} - v$ instead of $\mathcal{H} - \{v\}$. The hypergraph $\mathcal{H}$ *shrunken to* $V' \subseteq V$ is the hypergraph $\mathcal{H}|_{V'} := \mathcal{H} - (V \setminus V')$.

Graphs. Our notation related to graphs is basically standard and heavily borrows from Diestel's book [13]. In particular, a *bridge* of a graph is an edge whose removal increases the graph's number of connected components. Analogously, a *cut-vertex* is a vertex whose removal increases the graph's number of connected components. Some special notation including the *gluing* of graphs is given below.

Boundaried graphs and gluing. For a nonnegative integer $b \in \mathbb{N}$, a *b-boundaried graph* is a triple (G, B, β), where G is a graph, $B \subseteq V(G)$ such that $|B| = b$, and $\beta \colon B \to \{1, \ldots, b\}$ is a bijection. Vertex subset B is called the *boundary* and β the *boundary labeling*. For ease of notation we also refer to (G, B, β) as the *b-boundaried* graph G with boundary B and boundary labeling β. For brevity, a b-boundaried graph G whose boundary is the domain of β and whose boundary labeling is β is also called *β-boundaried*.

For a nonnegative integer b, the *gluing* operation $\circ_b$ maps two b-boundaried graphs to an ordinary graph as follows: Given two b-boundaried graphs G_1, G_2 with corresponding boundaries B_1, B_2 and boundary labelings β_1, β_2, to obtain the graph $G_1 \circ_b G_2$ take the disjoint union of G_1 and G_2, and identify each pair $v \in B_1$ and $w \in B_2$ of vertices such that $\beta_1(v) = \beta_2(w)$. We omit the index b in $\circ_b$ if it is clear from the context.

Topology. A *topological space* is a pair $\mathfrak{X} = (X, \mathcal{F})$ of a set X, called *universe*, and a collection $\mathcal{F}$ of subsets of X, called *topology*, that satisfy the following properties:

- The empty set $\emptyset$ and X are in $\mathcal{F}$.
- The union of the sets of any subcollection of $\mathcal{F}$ is in $\mathcal{F}$.
- The intersection of the sets of any finite subcollection of $\mathcal{F}$ is in $\mathcal{F}$.

Each set in $\mathcal{F}$ is called *open*. A *closed set* is the complement of an open set. (The empty set and X are both open and closed.)

We consider the topological space $\mathfrak{R}^\ell = (\mathbb{R}^\ell, \mathcal{F})$, where $\mathcal{F}$ is the standard topology of $\mathbb{R}^\ell$, that is, $\mathcal{F}$ is the closure under union and finite intersection of the set containing the open ball $\{\vec{x} \in \mathbb{R}^\ell \mid \|\vec{x} - \vec{y}\| < d\}$ for each $d \in \mathbb{R}$, $\vec{y} \in \mathbb{R}^\ell$, where $\|\cdot\|$ is the Euclidean norm.

A *topological subspace* $\mathfrak{Y} \subseteq \mathfrak{X}$ of a topological space $\mathfrak{X}$ is a topological space whose universe is a subset of the universe of $\mathfrak{X}$. We always assume topological subspaces to carry the *subspace topology*, that is, the open sets of $\mathfrak{Y}$ are the intersections of the open sets of $\mathfrak{X}$ with the universe of $\mathfrak{Y}$. We also say that $\mathfrak{Y}$ is the topological subspace *induced* by the universe of $\mathfrak{Y}$.

Important topological subspaces of $\mathfrak{R}^\ell$ are, with a slight abuse of notation,

- the *plane* $\mathfrak{R}^2$,
- the *sphere*, whose universe is $\{(x, y, z) \in \mathbb{R}^3 \mid x^2 + y^2 + z^2 = 1\}$,
- the *closed disk*, whose universe is $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leq 1\}$,
- the *open disk*, whose universe is $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 < 1\}$, and
- the *circle*, whose universe is $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 = 1\}$.

A *homeomorphism* ϕ between two topological spaces is a bijection ϕ between the two corresponding universes such that both ϕ and ϕ^{-1} are continuous. We often refer to a *subspace* $\mathfrak{X}$ *in a topological space* $\mathfrak{Y}$ (for example, a circle on a sphere), by which we mean a topological subspace of $\mathfrak{Y}$ which is homeomorphic to $\mathfrak{X}$.

An *arc* is a topological space that is homeomorphic to the closed interval $[0, 1] \subseteq \mathfrak{R}^1$. The images of 0 and 1 under a corresponding homeomorphism are the *endpoints* of the arc, which *links* them and runs *between* them. Let $\mathfrak{X} = (X, \mathcal{F})$ be a topological space. Being linked by an arc in $\mathfrak{X}$ defines an equivalence relation on X. The topological subspaces induced by the equivalence classes of this relation are called *regions*. We say that a closed set C in a topological space $\mathfrak{S}$ *separates* $\mathfrak{S}$ into the regions of the subspace of $\mathfrak{S}$ induced by $S \setminus C$, where S is the universe of $\mathfrak{S}$.

For more on topology, see Munkres [32], for example.

Embeddings of graphs into the plane and sphere. An *embedding* of a graph $G = (V, E)$ into the plane $\mathfrak{R}^2$ (into the sphere $\mathfrak{S}$) is a tuple $(\mathfrak{V}, \mathcal{E})$ and a bijection $\phi\colon V \to \mathfrak{V}$ such that

- $\mathfrak{V} \subseteq \mathfrak{R}^2$ ($\mathfrak{V} \subseteq \mathfrak{S}$),
- $\mathcal{E}$ is a set of arcs in $\mathfrak{R}^2$ (in $\mathfrak{S}$) with endpoints in $\mathfrak{V}$,
- the interior of any arc in $\mathcal{E}$ (that is, the arc without its endpoints) contains no point in $\mathfrak{V}$ and no point of any other arc in $\mathcal{E}$, and
- $u, v \in V$ are adjacent in G if and only if $\phi(u)$ is linked to $\phi(v)$ by an arc in $\mathcal{E}$.

The regions in $\mathfrak{R}^2 \setminus (\bigcup \mathcal{E})$ (in $\mathfrak{S} \setminus (\bigcup \mathcal{E})$) are called *faces*.

A *planar graph* is a graph which has an embedding in the plane or, equivalently, in the sphere. A *minor* of a graph G is a graph obtained from a subgraph of G by contracting edges, that is, replacing the two endpoints of an edge $\{u, v\}$ by a new vertex which is adjacent to all vertices in $N(u) \cup N(v) \setminus \{u, v\}$. It follows from Kuratowski's theorem that a graph is planar if and only if it does not have a K_5 or a $K_{3,3}$ as a minor [13, Section 4.4]. A *plane graph* $G = (V, E)$ is a planar graph together with a fixed embedding in the plane. An $\mathfrak{S}$-*plane graph* G is a planar graph given with a fixed embedding in the sphere. For notational convenience, we refer to the sets V and $\mathfrak{V}$ as well as E and $\mathcal{E}$ interchangeably. Moreover, we sometimes identify G with the set of points $\mathfrak{V} \cup \bigcup \mathcal{E}$.

A *noose in an* $\mathfrak{S}$-*plane graph* G is a circle in $\mathfrak{S}$ whose intersection with G is contained in $V(G)$ and, moreover, for each face of G the interesection of the face with the circle forms at most one consecutive interval on the circle. (The condition on the intersection with faces is sometimes dropped in the literature. Nooses that satisfy the condition on the intersection are sometimes called strict nooses.) Note that every noose separates $\mathfrak{S}$ into two open disks.

Layer decompositions, outerplanar graphs, face paths. The face of unbounded size in the embedding of a plane graph G is called *outer face*. The *layer decomposition* of G with respect to the embedding is a partition of V into layers $L_1 \uplus \cdots \uplus L_r$ and is defined inductively as follows. Layer L_1 is the set of vertices that lie on the outer face of G. For each $i \in \{2, \ldots, r\}$, layer L_i is the set of vertices that lie on the outer face of $G - (\bigcup_{j=1}^{i-1} L_j)$. The graph G is called *r-outerplanar* if it has an embedding with a layer decomposition consisting of at most r layers. The *outerplanarity number* of G is the minimum r such that G is r-outerplanar. If $r = 1$, then G is said to be *outerplanar*. A *face path* is an alternating sequence of faces and vertices such that two consecutive elements are incident with one another. The first and last element of a face path are called its *ends*. Note that the ends of a face path may be two vertices, two faces, or a face and a vertex. The *length* of a face path is the number of faces in the sequence. Note that a vertex v in layer L_i has a face path of length i from v to the outer face. Moreover, a graph is r-outerplanar if and only if each vertex has a face path of length at most r to the outer face.

Branch decompositions. A *branch decomposition of a graph* G is a pair (T, λ), where T is a ternary tree, that is, each internal vertex has degree three, and λ is a bijection between the leaves of T and $E(G)$. Every edge $e \in E(T)$ defines a bipartition of $E(G)$ into A_e, B_e corresponding to the leaves in the connected components of $T - e$. Define the *middle set* $M(e)$ of an edge $e \in E(T)$ to be the set of vertices in G that are incident with both an edge in A_e and B_e. That is,

$$M(e) \coloneqq \{v \in V(G) \mid \exists a \in A_e \exists b \in B_e \colon v \in a \cap b\}.$$

The *width of an edge* $e \in E(T)$ is $|M(e)|$ and the *width of a branch decomposition* (T, λ) is the largest width of an edge in T. The *branchwidth of a graph* G is the smallest width of a branch decomposition of G.

A *sphere-cut branch decomposition of an* $\mathfrak{S}$*-plane graph* G is a branch decomposition (T, λ) of G fulfilling the following additional condition. For each edge $e \in E(T)$, there is a noose $\mathfrak{N}_e$ whose intersection with G is precisely $M(e)$ and, furthermore, the open disks $\mathfrak{D}_1, \mathfrak{D}_2$ into which the noose $\mathfrak{N}_e$ separates $\mathfrak{S}$ satisfy $\mathfrak{D}_1 \cap G = A_e \setminus M(e)$ and $\mathfrak{D}_2 \cap G = B_e \setminus M(e)$. We use the following theorem.

Theorem 2 ([35, 14, 30]) *Let G be a connected, bridgeless, $\mathfrak{S}$-plane graph of branchwidth at most b. There exists a sphere-cut branch decomposition for G of width at most b.*

Dorn et al. [14] first noted that Seymour and Thomas [35] implicitly proved a variant of Theorem 2 in which G is required to have no degree-one vertices rather than no bridges. Marx and Pilipczuk [30] observed a flaw in Dorn et al.'s derivation, showing that bridgelessness is required (and sufficient). The sphere-cut branch decomposition in Theorem 2 can be computed in $\mathrm{O}(|V(G)|^3)$ time (see Gu and Tamaki [20]), but we do not need to explicitly construct it.

Parameterized algorithms. Let Σ be an alphabet. A *parameter* is a mapping $\Sigma^* \to \mathbb{N}$. For a string q, $\kappa(q)$ is the *parameter value*. A *parameterized problem* is a tuple (Q, κ) of a language Q over some alphabet Σ and a parameter κ. We say that an algorithm is a *fixed-parameter algorithm* with respect to a parameter κ if the algorithm has running time $\phi(\kappa(q)) \cdot \mathrm{poly}(|q|)$ where q is the input and ϕ is some computable function.

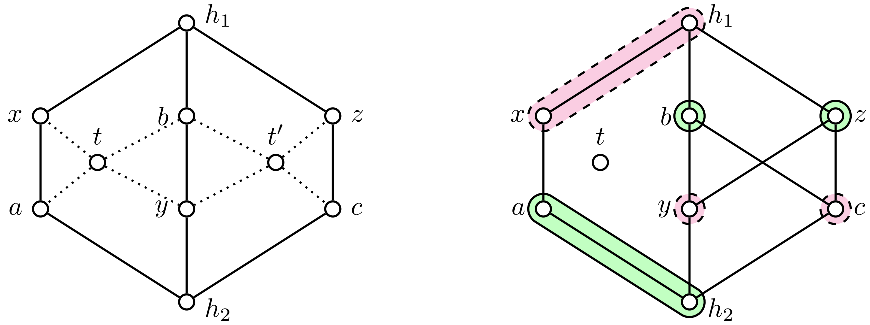

Figure 2: Left: A hypergraph $\mathcal{H}$ and its support, showing that twins can be essential for obtaining a 2-outerplanar support. The set of hyperedges consists of size-two hyperedges that are drawn as solid lines between the corresponding vertices and, additionally, $\{a, b, t, t'\}$, $\{b, c, t, t'\}$, $\{x, y, t, t'\}$, and $\{y, z, t, t'\}$. Note that the vertices t and t' are twins. The hypergraph $\mathcal{H}$ has a (2-outer)planar support whose edges are indicated by the solid and the dotted lines. However, $\mathcal{H} - t'$ does not have a planar support, as witnessed on the right: A $K_{3,3}$-minor of a support of $\mathcal{H} - t'$ obtained in the proof of Theorem 3. One partite vertex set of the minor is encircled with dashed lines and colored in red, the other partite vertex set is encircled solidly and colored in green.

To simplify notation, we omit explit reference to the function κ. For example, if we have a problem of finding a solution of size k, then we write "k" for the solution size parameter, that is, the mapping that takes an instance and extracts the value of k. Moreover, when specifying running times with respect to a parameter κ, we often replace $\kappa(q)$ by the referenced value if the instance q is clear from the context.

A core tool in the development of fixed-parameter algorithms is polynomial-time preprocessing by *data reduction* [21, 6, 28]. The goal is to remove needless information from the input so to reduce its size or to obtain some desirable properties of the input. Such small or well-formed instances can then be exploited by algorithms that produce a solution: small size of the input implies a small search space of the solution algorithm, and similarly, well-formed instances may be easier to solve.

Data reduction is usually presented as a series of *reduction rules*. These are polynomial-time algorithms that take as input an instance of some decision problem and produce another instance of the same problem as output. A reduction rule is *correct* if for each input instance I, the corresponding output instance of the rule is a yes-instance if and only if I is a yes-instance. We call an instance I of a parameterized problem *reduced* with respect to a reduction rule if the reduction rule does not apply to I. That is, carrying out that reduction rule yields an unchanged instance.

The notion of problem kernels captures the idea of reduction rules with effectiveness guarantee. A *kernelization* or *problem kernel* for a parameterized problem (Q, κ) is a parameterized reduction ρ from (Q, κ) to itself such that ρ is computable in polynomial time and there is a function ϕ such that for every $q \in \Sigma^*$ we have $|\rho(q)| \leq \phi(\kappa(q))$. We also call ϕ the *size* of ρ. If ϕ is polynomial, then we also call ρ a *polynomial kernelization* or *polynomial problem kernel*.

3 Beware of removing twins

In Figure 2, we provide a concrete example that shows that twins can be necessary to obtain a 2-outer-planar support. More precisely, the hypergraph on the left in Figure 2 witnesses the following theorem. Recall that, for a hypergraph $\mathcal{H}$ and a vertex v in $\mathcal{H}$, the hypergraph $\mathcal{H} - v$ is obtained from $\mathcal{H}$ by removing v out of each hyperedge in $\mathcal{H}$.

Theorem 3 *There is a hypergraph $\mathcal{H}$ with two twins t and t' such that $\mathcal{H}$ has a 2-outerplanar support, but the hypergraph $\mathcal{H} - t'$ does not have a planar support. Each hyperedge in $\mathcal{H}$ has size at most four.*

The hypergraph improves on the example given in the conference version of this article [38] by having fewer vertices, hyperedges, and smaller maximum hyperedge size. The proof of Theorem 3 is as follows:

Proof: The hypergraph $\mathcal{H}$ has vertex set $\{a, b, c, x, y, z, t, t', h_1, h_2\}$. The hyperedge set consists of the size-two hyperedges shown as solid lines on the left in Figure 2 and, additionally, the four hyperedges $\{a, b, t, t'\}$, $\{b, c, t, t'\}$, $\{x, y, t, t'\}$, and $\{y, z, t, t'\}$.

Consider the graph $\tilde{G}$ shown on the left of Figure 2 that includes both the dashed and solid edges. Observe that each of the size-four hyperedges of $\mathcal{H}$ induces a path in $\tilde{G}$ and thus $\tilde{G}$ is a support of $\mathcal{H}$. Moreover, $\tilde{G}$ is 2-outerplanar.

It remains to prove that $\mathcal{H} - t'$ does not have a planar support. Suppose, for the sake of contradiction, that there is a planar support G. Due to the hyperedges of size two in $\mathcal{H} - t'$, there are two cycles in G, namely C_1, defined as (x, h_1, b, y, h_2, a) and C_2, defined as (z, h_1, b, y, h_2, c). Consider a planar embedding of G. Consider the circle subspaces $\mathfrak{C}_1$ and $\mathfrak{C}_2$ in the plane defined by the embedding of C_1 and C_2, respectively. Circles $\mathfrak{C}_1$ and $\mathfrak{C}_2$ separate the plane into three open regions. These three open regions define three smallest enclosing closed regions:

- The region $\mathfrak{R}_1$, containing the vertices of C_1 and no vertices in $V(C_2) \setminus V(C_1)$,
- region $\mathfrak{R}_2$, containing the vertices of C_2 and no vertices in $V(C_1) \setminus V(C_2)$, and
- region $\mathfrak{R}_3$, containing h_1, h_2, the vertices in the symmetric difference $(V(C_1) \setminus V(C_2)) \cup (V(C_2) \setminus V(C_1))$ of C_1 and C_2, and no vertex in $(V(C_1) \cap V(C_2)) \setminus \{h_1, h_2\}$.

Vertex t is embedded in one of these three regions. Because of symmetry it suffices to obtain a contradiction in the case where t is not embedded in region $\mathfrak{R}_2$, that is, the case where t is not embedded in region $\mathfrak{R}_1$ is analogous. Assume thus that t is not embedded in region $\mathfrak{R}_2$. Then, t is adjacent in G with at most one of $\{b, c\}$ and with at most one of $\{y, z\}$, by planarity of G and the definition of the three regions. Moreover, $G[\{b, c, t\}]$ and $G[\{y, z, t\}]$ are connected, because $\{b, c, t\}$ and $\{y, z, t\}$ are hyperedges of $\mathcal{H} - t'$, and thus G contains the edges $\{b, c\}$ and $\{y, z\}$. Thus, G contains the graph shown on the right in Figure 2 as a subgraph. By contracting $\{x, h_1\}$ and $\{a, h_2\}$ in G we obtain a $K_{3,3}$ as a minor, contradicting the fact that G is planar. Thus, indeed, $\mathcal{H} - t'$ does not have a planar support. $\square$

Theorem 3 shows that removing one vertex of a twin class can transform a yes-instance of r-Outerplanar Support into a no-instance. We next address two features of the construction used in Theorem 3. First, we show that the counterexample requires hyperedges of size at least four. Second, we show that counterexamples exist with twin classes of arbitrary size, rather than just with pairs of twins.

Large twin classes. To show that the example from Theorem 3 is not a pathology of having only one pair of twins, we now extend it so that an arbitrarily large set of twins is required for the existence of a planar support. The basic observation is that in Theorem 3 we indeed have proved that there are two disjoint regions in which the two twins have to reside, for otherwise there is no planar support. By introducing several copies of the hypergraph constructed in Theorem 3 and merging them in an appropriate way, we can ensure that there is an arbitrarily large number ℓ of regions, of which each requires its own private vertex from a twin class of size ℓ.

Theorem 4 *For each integer $\ell \geq 2$, there is a hypergraph $\mathcal{H}$ with a set X of 2ℓ mutual twins such that $\mathcal{H}$ has a 2-outerplanar support, but for each $t \in X$ the hypergraph $\mathcal{H} - t$ does not have a planar support.*

Proof: Fix an integer $\ell \in \mathbb{N}$. To construct the hypergraph $\mathcal{H}$, copy ℓ times the vertex set $V(\mathcal{H})$ from the proof of Theorem 3, and let $V_i := \{a_i, b_i, c_i, x_i, y_i, z_i, h_i^1, h_i^2, t_i, t_i'\}$ denote the vertex set of the ith copy, $i \in \{1, \ldots, \ell\}$. The hyperedges are defined as follows. Within each copy, add the size-two solid hyperedges as shown on the left of Figure 2. Then, add a distinguished new vertex v^*, and add the size-two hyperedges $\{x_i, v^*\}$, $\{a_i, v^*\}$, $\{z_i, v^*\}$, and $\{c_i, v^*\}$. Intuitively, vertex v^* serves to force each of the copies into an embedding with the same outer face and as a conduit to connect subgraphs induced by hyperedges that contain vertices from each of the copies, which we are about to introduce.

Denote by A, B, C, X, Y, Z the sets of copies of the corresponding vertices above, that is, $A := \{a_i \mid i \in \{1, \ldots, \ell\}\}$, $B := \{b_i \mid i \in \{1, \ldots, \ell\}\}$, and so forth. Let $T := \{t_i, t_i' \mid i \in \{1, \ldots, \ell\}\}$ denote the set of designated twins. The final hyperedges in $\mathcal{H}$ are

$$A \cup B \cup T \cup \{v^*\}, \qquad B \cup C \cup T \cup \{v^*\},$$
$$X \cup Y \cup T \cup \{v^*\}, \qquad \text{and } Y \cup Z \cup T \cup \{v^*\}.$$

Note that T forms a twin class in the resulting hypergraph $\mathcal{H}$. Hypergraph $\mathcal{H}$ has a 2-outerplanar support because v^* can be used to join in a star-like fashion into one connected graph the partial supports for each $\mathcal{H}[V_i]$ that are obtained by copying the support for the example on the left of Figure 2.

We claim that, for each $t \in T$, hypergraph $\mathcal{H} - t$ does not have a planar support. Assume, for the sake of a contradiction, that there is a planar support G. Choose a planar embedding of G such that v^* is incident with the outer face. Due to the hyperedges of size two in $\mathcal{H} - t'$, there are 2ℓ cycles in G, namely, for each $i \in [\ell]$, there is the cycle C_i^1, defined as $(x_i, h_i^1, b_i, y_i, h_i^2, a_i)$ and the cycle C_i^2, defined as $(z_i, h_i^1, b_i, y_i, h_i^2, c_i)$. The embeddings of these cycles define $2\ell + 1$ closed regions as follows. The embedding of the cycle C_i^α, for $\alpha \in \{1, 2\}$ and $i \in [\ell]$, separates the plane into two open regions. Of these two regions, pick the one that has empty intersection with the outer face, take the smallest enclosing closed region and denote it by $\mathfrak{R}_{2i-2+\alpha}$. This defines 2ℓ regions $\mathfrak{R}_1, \ldots, \mathfrak{R}_{2\ell}$. Region $\mathfrak{R}_0$ is defined by taking the plane, removing from it all points in the regions $\mathfrak{R}_1, \ldots, \mathfrak{R}_{2\ell}$, and then taking the smallest enclosing closed region.

Since v^* is incident with the outer face, and because of the edges $\{x_i, v^*\}$, $\{a_i, v^*\}$, $\{z_i, v^*\}$, and $\{c_i, v^*\}$, vertex v^* is not contained in any region $\mathfrak{R}_i$ such that $i \in [2\ell]$. It follows that the regions $\mathfrak{R}_i$ are pairwise disjoint except for their boundary cycles. Thus, the regions $\mathfrak{R}_i$, $i \in \{0, \ldots, 2\ell\}$, satisfy the following properties:

- The region $\mathfrak{R}_0$ contains all vertices of each cycle C_i^α, $\alpha \in \{1, 2\}$, $i \in [\ell]$, except for the vertices in $B \cup Y$.
- For each $i \in [\ell]$, region $\mathfrak{R}_{2i-1}$ contains all vertices of C_i^1, no vertex in $V(C_i^2) \setminus V(C_i^1)$, and no vertex of any other cycle C_j^1, C_j^2 for $j \in [2\ell] \setminus \{i\}$.
- For each $i \in [\ell]$, region $\mathfrak{R}_{2i}$ contains all vertices of C_i^2, no vertex in $V(C_i^1) \setminus V(C_i^2)$, and no vertex of any other cycle C_j^1, C_j^2 for $j \in [2\ell] \setminus \{i\}$.

Since there are only $2\ell - 1$ twins from T present in $\mathcal{H} - t$, one of the regions R_i, $i \in \{1, \ldots, 2\ell\}$ does not contain a vertex of T, say R_{i°. By symmetry we may assume that i° is even.

Consider the hyperedge $F \coloneqq B \cup C \cup (T \setminus \{t\}) \cup \{v^\star\}$ and the two vertices b_{i° and c_{i° in F. Observe that the only region $\mathfrak{R}_i$, $i \in \{0, \ldots, 2\ell\}$, containing b_{i° and another vertex of F is $\mathfrak{R}_{i^\circ}$, and the only other vertex of F in $\mathfrak{R}_{i^\circ}$ is c_{i°. Thus, since $G[F]$ is connected, there is an edge $\{b_{i^\circ}, c_{i^\circ}\}$ in G. By an analogous argument for the hyperedge $Y \cup Z \cup (T \setminus \{t\}) \cup \{v^\star\}$ we obtain that $\{y_{i^\circ}, z_{i^\circ}\} \in E(G)$. Thus, $G[V_{i^\circ}]$ contains a subgraph as shown on the right of Figure 2. Thus, G contains a $K_{3,3}$-minor contradicting the fact that G is planar. Hence, $\mathcal{H} - t$ does not have a planar support, as claimed. □

Hyperedge size at most three. We now show that removing twins is correct, as long as they are contained only in hyperedges of size at most three. That is, the following data reduction rule for r-Outerplanar Support is correct.

Rule 5 *Let $\mathcal{H}$ be a hypergraph and u and v be twins in $\mathcal{H}$ such that each hyperedge containing u and v has size at most three. Then, remove u from $\mathcal{H}$.*

Proposition 3.1 *Rule 5 is correct, that is, $\mathcal{H}$ has an r-outerplanar support if and only if $\mathcal{H} - u$ has.*

Proof: If $\mathcal{H} - u$ has an r-outerplanar support G, then $\mathcal{H}$ has an r-outerplanar support G': add u as a degree-one neighbor to v in G to obtain G'. It remains to show that, if $\mathcal{H}$ has an r-outerplanar support G, then $\mathcal{H} - u$ has an r-outerplanar support.

First, assume that $\{u, v\} \in E(G)$. Then $\mathcal{H} - u$ has an r-outerplanar support G': Obtain G' from G by contracting the edge $\{u, v\}$ onto v. Since contracting an edge cannot create new faces in the corresponding embedding, G' is r-outerplanar. Clearly, each hyperedge induces a connected graph in G'.

Now assume that $\{u, v\} \notin E(G)$. Consider an arbitrary hyperedge F containing u and v in $\mathcal{H}$ and observe that F has size three. Let $w \in F \setminus \{u, v\}$. Since u and v are nonadjacent in G yet $G[F]$ is connected, we have $\{u, w\}, \{v, w\} \in E(G)$. Thus, $G - u$ is an r-outerplanar support for $\mathcal{H} - u$, as required. □

In particular, Proposition 3.1 shows that the largest hyperedge size in the example given in Theorem 3 is smallest possible.

4 The existence of an upper bound on the number of important twins

In the proof of Theorem 4, the number of hyperedges increases with the number of necessary twins we seek to enforce. We now show that this is unavoidable. That is, with a fixed number of hyperedges, it is impossible to create arbitrarily large twin classes out of which no twin can be deleted without violating the property of having a planar support.

Theorem 6 *There exists a function $\psi \colon \mathbb{N} \to \mathbb{N}$ with the following property. For each $m \in \mathbb{N}$ and every hypergraph $\mathcal{H}$ that has at most m hyperedges, out of each twin class of $\mathcal{H}$, we can remove all but $\psi(m)$ arbitrary twins such that the resulting hypergraph has a planar support if and only if $\mathcal{H}$ has one.*

The basic observation is that adding twins is not detrimental. If we have a planar support for $\mathcal{H}$, then we can make a new twin adjacent to one of its already present twins, so that the resulting graph remains planar. Reversing this idea, from each hypergraph with a planar support, by deleting twins we can obtain a *minimal* hypergraph $\mathcal{H}'$ which also has a planar support but from which no further twins can be deleted while maintaining the property of having a planar support. Using Dickson's lemma (see below for details) it is not hard to show that there is a function ϕ such that, for each fixed number m of hyperedges, there are only $\phi(m)$ such minimal hypergraphs. Clearly, among these minimal hypergraphs, one has a largest twin class, whose size we can put as the value of $\psi(|\mathcal{E}(\mathcal{H})|)$.

We now formalize the above approach. Denote by $\mathbb{S}$ the set of hypergraphs which have a planar support. As mentioned, $\mathbb{S}$ is *closed under adding twins*, that is, taking an arbitrary hypergraph in $\mathbb{S}$ and adding a twin to it yields another hypergraph in $\mathbb{S}$. We commence with the proof of Theorem 6:

Proof: We first define a quasi-order $\preceq$ on the family of hypergraphs with at most m hyperedges. (A *quasi-order* is reflexive and transitive.) To define $\preceq$, we say that $\mathcal{H} \preceq \mathcal{G}$ if $\mathcal{H}$ can be obtained from $\mathcal{G}$ by iteratively removing a vertex that has a twin. If we allow zero removals so that $\preceq$ is reflexive, it is clear that $\preceq$ is a quasi-order. Moreover, if $\mathcal{H} \in \mathbb{S}$ and $\mathcal{H} \preceq \mathcal{G}$, then $\mathcal{G} \in \mathbb{S}$ since $\mathbb{S}$ is closed under adding twins.

For every $m \in \mathbb{N}$ let $\mathbb{F}_m$ denote the family of hypergraphs in $\mathbb{S}$ that contain at most m hyperedges and are minimal under $\preceq$. Next we show that $\mathbb{F}_m$ is finite. Consider the representation of a hypergraph $\mathcal{H}$ with at most m hyperedges as a 2^m-tuple $t_\mathcal{H} \in \mathbb{N}^{2^m}$, each entry of which represents the size of a distinct twin class. The set of such tuples is quasi-ordered by a natural extension of $\leq$, namely $(a_1, \ldots, a_\ell) \leq (b_1, \ldots, b_\ell)$ if and only if $a_i \leq b_i$ for each $i \in \{1, \ldots, \ell\}$. We now lead the assumption that $\mathbb{F}_m$ is infinite to a contradiction by using Dickson's lemma [11, Lemma A].

If $\mathbb{F}_m$ is infinite, then there is an infinite subset $\mathbb{F}'_m$ of hypergraphs which have the same (nonempty) twin classes. That is, the tuples representing the hypergraphs of $\mathbb{F}'_m$ have the same 0-entries. For hypergraphs $\mathcal{H}, \mathcal{G}$ with the same twin classes, $t_\mathcal{H} \leq t_\mathcal{G}$ implies $\mathcal{H} \preceq \mathcal{G}$. Thus, $\mathbb{F}'_m$ gives an infinite set T of tuples that are pairwise incomparable under $\leq$. Dickson's lemma states that for every set $S \subseteq \mathbb{N}^\ell$ there exists a finite subset $S' \subseteq S$ such that for each $s \in S$ there is an $s' \in S'$ with $s' \leq s$. This is a contradiction to T containing infinitely many incomparable tuples. Hence, $\mathbb{F}_m$ is finite.

Finally, we simply choose the function $\psi(m)$ in Theorem 6 as the largest size of a twin class of a hypergraph in $\mathbb{F}_m$. □

Let us briefly consider the implications of Theorem 6 for our ultimate goal—designing algorithms that compute for any given hypergraph a planar support if there is one and have fixed-parameter running time with respect to the number m of hyperedges in the input hypergraph. Theorem 6 does imply that, for each $m \in \mathbb{N}$, there is such an algorithm which works for all input hypergraph with at most m edges: The algorithm has the minimal yes-instances (in $\mathbb{F}_m$ from the proof) as hard-coded constants and checks whether the given hypergraph $\mathcal{G}$ satisfies $\mathcal{H} \preceq \mathcal{G}$ for some $\mathcal{H} \in \mathbb{F}_m$. However, a priori we do not not know how to compute the set $\mathbb{F}_m$ of minimal yes-instances, making this result nonconstructive and thus not useful in practice. Furthermore, we would like to instead have one algorithm which works for any input hypergraph. Hence, to eventually obtain implementable algorithms that are able to deal with any input, it is important to constructivize Theorem 6. Below, we provide such a constructivization for r-outerplanar supports.

5 Relevant twins for r-outerplanar supports

Towards making the result of Theorem 6 algorithmically exploitable, we now give an explicit upper bound on the function ψ. Concretely, we prove that out of each twin class of a hypergraph $\mathcal{H}$, we can remove all but $\psi(m, r)$ twins such that the resulting hypergraph has an r-outerplanar support if and only if $\mathcal{H}$ has. In other words, we prove that the following data reduction rule is correct.

Rule 7 *Let $\mathcal{H}$ be a hypergraph with m edges. If there is a twin class with more than $\psi(m, r) = 2^{6r \cdot 2^{m \cdot (2r^2+r+1)} \cdot (r+1)^{32r^2+8r}}$ vertices, then remove one vertex out of this class.*

Assuming that Rule 7 is correct, we can show Theorem 1 as follows:

Proof: Rule 7 can be applied exhaustively in linear time because the twin classes can be computed in linear time [22]. After this, each twin class contains at most $\psi(m, r)$ vertices, meaning that, overall, at most $2^m \psi(m, r)$ vertices remain. Testing all possible planar graphs for whether they are a support for the resulting hypergraph thus takes constant time if m and r are constant. Hence, the overall running time is linear in the input size. □

We mention in passing that, in terms of parameterized algorithmics, exhaustive application of Rule 7 yields a so-called problem kernel.

The correctness proof for Rule 7 consists of two parts. First, in Theorem 9, we show that each r-outerplanar graph has a long sequence of nested separators. Here, *nested* means that each separator separates the graph into a *left* side and a *right* side, and each left side contains all previous left sides. Furthermore, the sequence of separators has the additional property that, for any pair of separators S_1, S_2, we can glue the left side of S_1 and the right side of S_2, obtaining another r-outerplanar graph.

In the second part of the proof, we fix an initial r-outerplanar support for our input hypergraph. We then show that, in a long sequence of nested separators for this support, there are two separators such that we can carry out the following procedure. We discard all vertices between the separators, glue their left and right sides, and reattach the vertices which we discarded as degree-one vertices. Furthermore, we can do this in such a way that the resulting graph is an r-outerplanar support. The reattached degree-one vertices hence are not crucial to obtain an r-outerplanar support. We will show that if our input hypergraph contains more than $\psi(m, r)$ vertices, then there is always at least one vertex which can be discarded because it is between two suitable separators.

We now formalize our approach. Theorem 9 will guarantee the existence of a long sequence of gluable separators; it is proven in Section 6. To formally state it, we need the following notation.

Definition 8 (Middle set, subgraph $G\langle A\rangle$ induced by an edge set) *For an edge bipartition $A \uplus B = E(G)$ of a graph G, let $M(A, B)$ be the set of vertices in G which are incident with both an edge in A and in B, that is,*

$$M(A, B) := \{v \in V(G) \mid \exists a \in A \exists b \in B \colon v \in a \cap b\}.$$

We call $M(A, B)$ the middle set *of A, B.*

For an edge set $A \subseteq E(G)$, let $G\langle A\rangle := (\bigcup A, A)$ be the subgraph induced by *A.*

Observe that the definition of a middle set for an edge partition corresponds exactly to the definition of middle sets in branch decompositions. Recall from Section 2 the definitions of graph gluing, boundary, and boundary labeling.

Theorem 9 *For every connected, bridgeless, r-outerplanar graph G with n vertices, there is a sequence $((A_i, B_i, \beta_i))_{i=1}^{s}$ where each pair $A_i, B_i \subseteq E(G)$ is an edge bipartition of G and β_i is a bijection $M(A_i, B_i) \to \{1, \ldots, |M(A_i, B_i)|\}$ such that $s \geq \log(n)/(r+1)^{32r^2+8r}$, and, for every i, j, $1 \leq i < j \leq s$,*

(i) $|M(A_i, B_i)| = |M(A_j, B_j)| \leq 2r$,

(ii) $A_i \subsetneq A_j$, $B_i \supsetneq B_j$, *and*

(iii) $G\langle A_i\rangle \circ G\langle B_j\rangle$ *is r-outerplanar, wherein we think of $G\langle A_i\rangle$ as a β_i-boundaried graph and $G\langle B_j\rangle$ as a β_j-boundaried graph.*

To gain some intuition for Theorem 9 note that each $M(A_i, B_i)$ is a separator, separating its left side $G\langle A_i\rangle$ from its right side $G\langle B_i\rangle$ in G. Statement (ii) ensures that each left side contains all previous left sides, that is, the separators are nested. Statement (iii) ensures that for any two separators in the sequence, we can glue their left and right sides and again obtain an r-outerplanar graph. In this new graph, the vertices between the separators are missing—these will be the vertices which are not crucial to obtain an r-outerplanar support.

The reason why we can prove the lower bound on the length of the sequence is basically that r-outerplanar graphs have a tree-like structure, whence large r-outerplanar graphs have a long "path" in this structure, and a long path in such a structure induces many nested separators. From such a path we can pick the separators that are amenable to Statement (iii).

We next formalize the crucial vertices for obtaining an r-outerplanar support. These are the vertices in a smallest representative support, defined as follows. Recall that a vertex v in a hypergraph *covers* a vertex u if $\mathcal{E}(u) \subseteq \mathcal{E}(v)$, where $\mathcal{E}(u)$ and $\mathcal{E}(v)$ are the hyperedges containing the corresponding vertices.

Definition 10 (Representative support) *Let $\mathcal{H}$ be a hypergraph. A graph G is a* representative support *for $\mathcal{H}$ if $V(G) \subseteq V(\mathcal{H})$, graph G is a support for subhypergraph $\mathcal{H}|_{V(G)}$ shrunken to $V(G)$, and every vertex in $V(\mathcal{H}) \setminus V(G)$ is covered in $\mathcal{H}$ by some vertex in $V(G)$.*

Using the sequence of separators from Theorem 9, we show that the size of a smallest representative r-outerplanar support is upper-bounded by a function of m and r. To this end, we take an initial support, find two separators such that each vertex between them can be removed and reattached as a neighbor of some vertex that covers it and is not between the separators. This implies that the removed vertices need not be contained in a representative support. Intuitively, the two separators have to have the same "status" with respect to the hyperedges that cross them. We formalize this as follows (some further intuition is given after the definition).

Definition 11 (Edge-bipartition signature) *Let $\mathcal{H} = (V, \mathcal{E})$ be a hypergraph and let G be a representative planar support for $\mathcal{H}$. Let (A, B, β) be a triple where (A, B) is an edge bipartition of G and β is a bijection $M(A, B) \to \{1, \ldots, |M(A, B)|\}$. Let $\ell := |M(A, B)|$. The* signature *of (A, B, β) is a triple $(\mathcal{T}, \phi, \mathcal{C})$, where*

- $\mathcal{T} := \{\{F \in \mathcal{E} \mid u \in F\} \mid u \in \bigcup A\}$ *is the family of sets of hyperedges that contain some vertex in $\bigcup A$,*
- $\phi : \{1, \ldots, \ell\} \to \{[u]_\tau \mid u \in V\} \colon j \mapsto [\beta^{-1}(j)]_\tau$ *maps each index of a vertex in $M(A, B)$ to the twin class of that vertex, and*
- $\mathcal{C} := \{(F, \gamma_F) \mid F \in \mathcal{E}\}$, *where γ_F is the relation on $\{1, \ldots, \ell\}$ defined by $(i, j) \in \gamma_F$ whenever both $\beta^{-1}(i), \beta^{-1}(j) \in F$ and $\beta^{-1}(i)$ is connected to $\beta^{-1}(j)$ in $G\langle B\rangle[F \cap \bigcup B]$.*

For easier notation, we will in the following refer to the sets in the family $\mathcal{T}$ in an edge bipartition signature simply as a *twin class*.

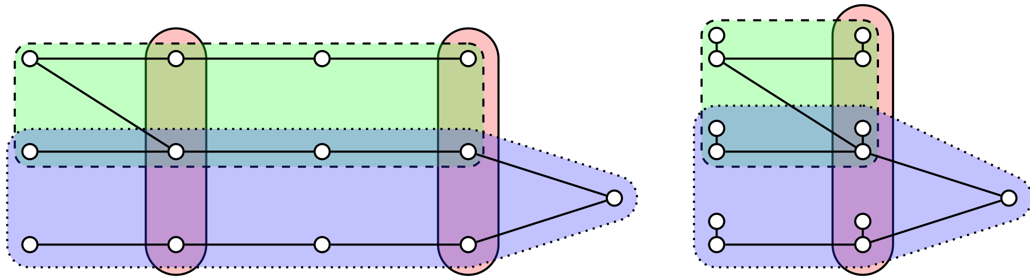

Figure 3: Intuition behind signatures and how we use them. Left: A hypergraph $\mathcal{H}$ containing the white vertices and two hyperedges (dashed green and dotted blue encircled vertices) together with a support (solid lines between vertices) and two separators (solid red encircled vertices) that have the same signature. Note that each twin class on the left (A-side) of the right separator also occurs on the left (A-side) of the left separator and vice versa (this is the first entry of the signature tuple). The vertices in the left separator have the same twin classes as the ones in the right separator (second entry of the signature tuple). Furthermore, for both separators the two vertices of the green dashed hyperedge are not connected on the right (B-side) of both separators; the two vertices of the blue hyperedge are connected on the right (B-side) of both separators (third entry of the signature tuple). Right: The same hypergraph $\mathcal{H}$ together with a support obtained from the support on the left by gluing the left side of the first separator and the right side of the second separator and modifying all the vertices between separators in the support (and those that have been identified during the gluing operation) so that they become degree-one vertices in the support.

Intuitively, the first component encodes the twin classes present on the "A-side" of the edge bipartition. The second component encodes the twin classes of the vertices in the middle set of the edge bipartition. The third component encodes for each hyperedge F, which pairs of vertices on the noose are connected if we restrict F to the B-side of the edge bipartition. See the left side of Figure 3 for an illustration.

Intuitively, the signatures are used as follows (see also Figure 3). If we have two separators $M(A_i, B_i)$, $M(A_j, B_j)$, $i < j$, from Theorem 9, then the left side, $G\langle A_i\rangle$, of $M(A_i, B_i)$ and the right side, $G\langle B_j\rangle$, of $M(A_j, B_j)$ can be glued to obtain an r-outerplanar graph. If the two separators additionally have the same edge-bipartition signature, then the three elements of the signature ensure in the following way that we maintain that the glued graph remains a representative support for the hypergraph $\mathcal{H}$. First, since there are the same twin classes on the two left sides $G\langle A_i\rangle$, $G\langle A_j\rangle$ due to element $\mathcal{T}$, and the left side $G\langle A_j\rangle$ of $M(A_j, B_j)$ contains all vertices between the two separators, all removed vertices between the two separators are still covered in the new graph. Second, elements ϕ and $\mathcal{C}$ ensure that the connectivity relation of each hyperedge in $\bigcup B_i$ is the same as in $\bigcup B_j$. That is, regardless of how the concrete connectivity in the "forgotten part" between the two separators looks like, replacing B_i with B_j preserves the fact that each hyperedge induces a connected subgraph.

We have the following upper bound on the number of different separator states.

Lemma 1 *In a sequence $((A_i, B_i, \beta_i))_{i=1}^{s}$ as in Theorem 9 the number of distinct edge-bipartition signatures is upper-bounded by $2^{m\cdot(2r^2+r+1)}$.*

Proof: Denote the signature of (A_i, B_i, β_i) by $(\mathcal{T}_i, \phi_i, \mathcal{C}_i)$. There are at most $2^m - 1$ twin classes in $\mathcal{T}_i$. Furthermore, for every i, j, $i < j$, we have $A_i \subsetneq A_j$, which implies $\mathcal{T}_i \subseteq \mathcal{T}_j$. Thus, either

$\mathcal{T}_i = \mathcal{T}_{i+1}$ or $\mathcal{T}_{i+1}$ comprises at least one additional twin class. Since the number of twin classes can increase at most $2^m - 1$ times, the number of different $\mathcal{T}_i$ is less than 2^m. Next, there are at most 2^m choices for a twin class for each $\beta^{-1}(i) \in M(A_i, B_i)$, leading to at most $2^{m\ell}$ different possibilities where $\ell = |M(A_i, B_i)|$. For the last part of the signature, $\mathcal{C}_i$, for each γ_e there are $2^{(\ell^2-\ell)/2}$ possibilities, leading to $2^{m(\ell^2-\ell)/2}$ possibilities for $\mathcal{C}_i$. Since the size ℓ of the middle sets in Theorem 9 is at most $2r$, the number of possible signatures is at most $2^m \cdot 2^{2mr} \cdot 2^{m\cdot(2r^2-r)} \;=\; 2^{m\cdot(2r^2+r+1)}$. □

As before, let $\psi(m, r) \coloneqq 2^{6r \cdot 2^{m\cdot(2r^2+r+1)} \cdot (r+1)^{32r^2+8r}}$.

Lemma 2 *If a hypergraph $\mathcal{H} = (V, \mathcal{E})$ has an r-outerplanar support, then it has a representative r-outerplanar support with at most $\psi(m, r)$ vertices.*

Proof: Let $G = (W, E)$ be a representative r-outerplanar support for $\mathcal{H}$ with the minimum number of vertices and fix a corresponding planar embedding. Assume towards a contradiction that $|W| > \psi(m, r)$. We show that there is a representative support for $\mathcal{H}$ with less than $\psi(m, r)$ vertices.

We aim to apply Theorem 9 to G. For this we need that G is connected and does not contain any bridges. Indeed, if G is not connected, then pick a vertex incident with the outer face for each connected component and add edges to make the picked vertices induce a path. This does not affect the outerplanarity number of G (although it adds bridges). If G has a bridge $\{u, v\}$, then proceed as follows. At least one of the ends of the bridge, say v, has degree at least two because $|W| > \psi(m, r) \geq 2$. One neighbor $w \neq u$ of v is incident with the same face as u, because $\{u, v\}$ is a bridge. Add the edge $\{u, w\}$. Thus, edge $\{u, v\}$ ceases to be a bridge. We can embed $\{u, w\}$ in such a way that the face $\mathfrak{F}$ incident with u, v, and w is split into one face that is incident with only $\{u, v, w\}$ and devoid of any other vertex, and one face $\mathfrak{F}'$ that is incident with all the vertices that are incident with $\mathfrak{F}$ including u, v, and w. This implies that each vertex retains its layer L_i, meaning that G remains r-outerplanar. Thus, we may assume that G is connected, bridgeless, and r-outerplanar.

Since G contains more than $\psi(m, r)$ vertices, there is a sequence $\mathcal{S} = ((A_i, B_i, \beta_i))_{i=1}^{s}$ as in Theorem 9 of length at least

$$s \;\geq\; \frac{\log(\psi(m,r))}{(r+1)^{32r^2+8r}} \;=\; \frac{6r \cdot 2^{m\cdot(2r^2+r+1)} \cdot (r+1)^{32r^2+8r}}{(r+1)^{32r^2+8r}} \;=\; 6r \cdot 2^{m\cdot(2r^2+r+1)}.$$

Since there are less than $2^{m\cdot(2r^2+r+1)}$ different signatures in $\mathcal{S}$ (Lemma 1), there are $6r$ elements of $\mathcal{S}$ with the same signature. Note that each middle set $M(A_i, B_i)$ induces a planar graph in G and, since $|M(A_i, B_i)| \leq 2r$, this graph has at most $\max\{1, 3|M(A_i, B_i)| - 6\} \leq \max\{1, 6r - 6\}$ edges. Recall that $A_i \subsetneq A_{i+1}$ for each $i \in [s-1]$, and thus, $|A_j| - |A_i| \geq j - i$ for each $i, j \in [s]$ with $i < j$. Thus, if $i, j \in [s]$ with $i + 6r \leq j$, $M(A_i, B_i)$ and $M(A_j, B_j)$ differ in at least one vertex. Thus, there are two edge bipartitions (A_i, B_i, β_i) and (A_j, B_j, β_j), $i < j$, in $\mathcal{S}$ with the same signature such that the middle sets $M(A_i, B_i)$, $M(A_j, B_j)$ differ in at least one vertex.

Let $G_{ij} \coloneqq G\langle A_i\rangle \circ G\langle B_j\rangle$, wherein $G\langle A_i\rangle$ is β_i-boundaried and $G\langle B_j\rangle$ is β_j-boundaried, and let $W' \coloneqq V(G_{ij})$. Note that $|W'| < |W|$ since the middle sets of the two edge bipartitions differ in at least one vertex and since $A_i \subsetneq A_j$.

We prove that G_{ij} is a representative support for $\mathcal{H}$, that is, each vertex $V \setminus W'$ is covered in $\mathcal{H}$ by some vertex in W' and that G_{ij} is a support for $\mathcal{H}|_{W'}$. Herein, for the sake of simpler notation, when referring to the covering condition of representative supports, we assume that each vertex in G_{ij} that results from identifying two vertices in G is equal to an arbitrary one of the two identified vertices. Since G_{ij} is r-outerplanar by Theorem 9, Statement (iii), the existence of G_{ij} contradicts the choice of G according to the minimum number of vertices, thus proving the lemma.

To prove that each vertex in $V \setminus W'$ is covered by some vertex in W', we show that $\{[u]_\tau \mid u \in V\} = \{[u]_\tau \mid u \in W'\}$. (Hence, every vertex in $V \setminus W'$ has a twin in W'; observe that twins cover each other.) Since $G = (W, E)$ is a representative support, $\{[u]_\tau \mid u \in V\} = \{[u]_\tau \mid u \in W\}$. Furthermore, by the definition of signature, we have $\{[u]_\tau \mid u \in \bigcup A_i\} = \{[u]_\tau \mid u \in \bigcup A_j\}$. Thus, for each vertex $u \in W \setminus W'$, there is a vertex $v \in W'$ with $[u]_\tau = [v]_\tau$, meaning that, indeed, $\{[u]_\tau \mid u \in V\} = \{[u]_\tau \mid u \in W'\}$.

To show that G_{ij} is a representative support it remains to show that it is a support for $\mathcal{H}|_{W'}$, that is, each hyperedge F' of $\mathcal{H}|_{W'}$ induces a connected graph $G_{ij}[F']$. Let F be a hyperedge of $\mathcal{H}$ such that $F \cap W' = F'$. Observe that such a hyperedge F exists and that $G[F \cap W]$ is connected since G is a representative support of $\mathcal{H}$. Denote by S_k the middle set $M(A_k, B_k)$ of (A_k, B_k) in G for $k \in \{i, j\}$ and by S the middle set $M(A_i, B_j)$ of (A_i, B_j) in G_{ij}. Note that the set S is obtained by identifying the vertices of S_i with S_j in the construction of G_{ij}.

To show that $G_{ij}[F']$ is connected, consider first the case that $F \cap (S_i \cup S_j) = \emptyset$. Since each vertex in $V \setminus W'$ is covered by a vertex in W', we have that all vertices in F are contained in either $G\langle A_i\rangle$ or $G\langle B_j\rangle$ along with all edges of $G[F]$. All these edges are also present in G_{ij} whence $G_{ij}[F']$ is connected.

Now consider the case that $F \cap (S_i \cup S_j) \neq \emptyset$. Since S_i and S_j are separators in G, each vertex in $F \setminus (S_i \cup S_j)$ is connected in $G[F]$ to some vertex in S_i or S_j via a path with internal vertices in $F \setminus (S_i \cup S_j)$. We consider the connectivity relation of their corresponding vertices in S. To this end, for a graph H and $T \subseteq V(H)$ use $\gamma(T, H)$ for the equivalence relation on T of connectivity in H. That is, for $u, v \in T$ we have $(u, v) \in \gamma(T, H)$ if u and v are connected in H. Using this terminology, since both S_i and S_j equal S in G_{ij}, to show that $G_{ij}[F']$ is connected, it is enough to prove that the transitive closure δ of $\gamma(F' \cap S, G_{ij}\langle A_i\rangle) \cup \gamma(F' \cap S, G_{ij}\langle B_j\rangle)$ contains only one equivalence class.

Denote by $\hat{G}$ the graph obtained from G by identifying each $v \in S_i$ with $\beta_j^{-1}(\beta_i(v)) \in S_j$, hence, identifying S_i and S_j, resulting in the set S. Relation $\alpha := \gamma(F \cap S, \hat{G})$ has only one equivalence class and, moreover, it is the transitive closure of $\gamma(F \cap S_i, G\langle A_i\rangle) \cup \gamma(F \cap S, \hat{G}\langle B_i \setminus B_j\rangle) \cup \gamma(F \cap S_j, G\langle B_j\rangle)$, wherein we identify each $v \in S_i$ with $\beta_j^{-1}(\beta_i(v)) \in S_j$ as above and, thus, $S_i = S_j = S$. We have $\gamma(F' \cap S, G_{ij}\langle A_i\rangle) = \gamma(F \cap S_i, G\langle A_i\rangle)$ and $\gamma(F' \cap S, G_{ij}\langle B_j\rangle) = \gamma(F \cap S_j, G\langle B_j\rangle)$. Thus for $\alpha = \delta$ it suffices to prove that $\gamma(F \cap S, \hat{G}\langle B_i \setminus B_j\rangle) \subseteq \gamma(F' \cap S_j, G_{ij}\langle B_j\rangle)$. Indeed, the left-hand side $\gamma(F \cap S, \hat{G}\langle B_i \setminus B_j\rangle)$ is contained in $\gamma(F \cap S_i, G\langle B_i\rangle)$: Let $(\mathcal{T}, \phi, \mathcal{C})$ be the signature of (A_i, B_i, β_i) and (A_j, B_j, β_j) and $(F, \gamma_F) \in \mathcal{C}$. Note that $\gamma(F \cap S_i, G\langle B_i\rangle) = \gamma_F = \gamma(F \cap S_j, G\langle B_j\rangle)$ where we abuse notation and set $u = \beta_i(u)$ for $u \in S_i$ and $v = \beta_j(v)$ for $v \in S_j$. Hence, $\gamma(F \cap S, \hat{G}\langle B_i \setminus B_j\rangle) \subseteq \gamma(F \cap S_j, G\langle B_j\rangle) = \gamma(F' \cap S_j, G\langle B_j\rangle) = \gamma(F' \cap S_j, G_{ij}\langle B_i\rangle)$. Thus, indeed, $\delta = \alpha$, that is, F' is connected. □

We now use the upper bound on the number of vertices in representative supports to get rid of superfluous twins. First, we show that representative supports can be extended to obtain a support.

Lemma 3 *Let $G = (W, E)$ be a representative r-outerplanar support for a hypergraph $\mathcal{H} = (V, \mathcal{E})$. Then $\mathcal{H}$ has an r-outerplanar support in which all vertices of $V \setminus W$ have degree one.*

Proof: Let G' be the graph obtained from G by making each vertex v of $V \setminus W$ a degree-one neighbor of a vertex in W that covers v (such a vertex exists by the definition of representative support). Clearly, the resulting graph is planar. It is also r-outerplanar, which can be seen by adapting an r-outerplanar embedding of G for G': If the neighbor v of a new degree-one vertex u

is in L_1, then place u in the outer face. If $v \in L_i$, $i > 1$, then place u in a face which is incident with v and a vertex in L_{i-1} (such a face exists by the definition of L_i).

It remains to show that G' is a support for $\mathcal{H}$. Consider a hyperedge $F \in \mathcal{E}$. Since G is a representative support for $\mathcal{H}$, we have that $F \cap W$ is nonempty and that $G[F \cap W]$ is connected. In G', each vertex $u \in F \setminus W$ is adjacent to some vertex $v \in W$ that covers u. Hence $v \in F$. Thus, $G'[F]$ is connected as $G'[F \cap W]$ is connected and all vertices in $F \setminus W$ are neighbors of a vertex in $F \cap W$. ☐

We now use Lemma 3 to show that, if there is a twin class that contains more vertices than a small representative support, then we can safely remove one vertex from this twin class.

Lemma 4 *Let $\ell \in \mathbb{N}$, let $\mathcal{H}$ be a hypergraph, and let $v \in V(\mathcal{H})$ be a vertex such that $|[v]_\tau| \geq \ell$. If $\mathcal{H}$ has a representative r-outerplanar support with less than ℓ vertices, then $\mathcal{H} - v$ has an r-outerplanar support.*

Proof: Let $G = (W, E)$ be a representative r-outerplanar support for $\mathcal{H}$ such that $|W| < \ell$. Then at least one vertex of $[v]_\tau$ is not in W and we can assume that this vertex is v without loss of generality. Thus, $\mathcal{H}$ has an r-outerplanar support G' in which v has degree one by Lemma 3. The graph $G' - v$ is an r-outerplanar support for $\mathcal{H} - v$: For each hyperedge F in $\mathcal{H} - v$, we have that $G'[F \setminus \{v\}]$ is connected because v is not a cut-vertex in $G'[F]$ (since it has degree one). ☐

Now we combine the observations above with the fact that there are small r-outerplanar supports to prove that Rule 7 is correct.

Proof: [Correctness proof for Rule 7] Consider an instance $\mathcal{H} = (V, \mathcal{E})$ of r-Outerplanar Support to which Rule 7 is applicable and let $v \in V$ be a vertex to be removed, that is, v is contained in a twin class of size more than $\psi(m, r)$. By Lemma 2, if $\mathcal{H}$ has an r-outerplanar support, then it has a representative r-outerplanar support with at most $\psi(m, r)$ vertices. By Lemma 4, this implies that $\mathcal{H} - v$ has an r-outerplanar support. Moreover, if $\mathcal{H} - v$ has an r-outerplanar support, then this r-outerplanar support is a representative r-outerplanar support for $\mathcal{H}$. By Lemma 3, this implies that $\mathcal{H}$ has an r-outerplanar support. Therefore, $\mathcal{H}$ and $\mathcal{H} - v$ are equivalent instances of r-Outerplanar Support, that is, $\mathcal{H}$ has an r-outerplanar support if and only if $\mathcal{H} - v$ has, and v can be safely removed from $\mathcal{H}$. ☐

6 A sequence of gluable edge bipartitions

In this section, we prove Theorem 9. For convenience, it is restated below. Recall also from Section 5 the intuitive description of the theorem statement, the definition of middle set for an edge bipartition, the subgraph $G\langle A\rangle$ induced by an edge set, and from Section 2 the definitions of graph gluing, boundary, and boundary labeling.

Theorem 9 *For every connected, bridgeless, r-outerplanar graph G with n vertices, there is a sequence $((A_i, B_i, \beta_i))_{i=1}^{s}$ where each pair $A_i, B_i \subseteq E(G)$ is an edge bipartition of G and β_i is a bijection $M(A_i, B_i) \to \{1, \ldots, |M(A_i, B_i)|\}$ such that $s \geq \log(n)/(r+1)^{32r^2+8r}$, and, for every i, j, $1 \leq i < j \leq s$,*

(i) $|M(A_i, B_i)| = |M(A_j, B_j)| \leq 2r$,

(ii) $A_i \subsetneq A_j$, $B_i \supsetneq B_j$, *and*

(iii) $G\langle A_i\rangle \circ G\langle B_j\rangle$ *is r-outerplanar, wherein we think of $G\langle A_i\rangle$ as a β_i-boundaried graph and $G\langle B_j\rangle$ as a β_j-boundaried graph.*

The proof relies crucially on sphere-cut branch decompositions [14, 30]. Recall the corresponding definitions from Section 2.

Outline of the proof of Theorem 9. We first transform the planar embedding of G into an embedding in the sphere. Using the fact that r-outerplanar graphs have branchwidth at most $2r$ [5], we may apply Theorem 2, from which we obtain a sphere-cut branch decomposition for G of width at most $2r$ (recall that the width of a branch decomposition is the size of its largest middle set). The edge bipartitions in Theorem 9 are defined based on the edges in a longest path in the decomposition tree corresponding to the sphere-cut branch decomposition. The longest path in the decomposition tree has length at least $2\log(n)$, and the edges on this path define a sequence of edge bipartitions, a supersequence of the one in Theorem 9. We define a labeling for each bipartition, which is a string containing $(32r^2 + 8r) \cdot \log(r+1) + 1$ bits, that determines the pairs of edge bipartitions that can be glued so as to obtain an r-outerplanar graph. The sequence in Theorem 9 is then obtained from those bipartitions that have the same labeling. The sphere-cut property of the branch decomposition gives one noose in the sphere for each edge bipartition in the sequence, such that it separates the parts in the edge bipartition from one another. The nooses of the sphere-cut branch decomposition will be crucial in the proof of Statement (iii) in Theorem 9, that is, the r-outerplanarity of the glued graphs.

Let us give some more details concerning the r-outerplanarity of the glued graphs. After slightly modifying the nooses if necessary, we can assume that they separate the sphere into *left* disks and *right* disks in such a way that each left disk contains all left disks with smaller indices. Hence, for each pair of nooses, we can cut out a left disk and a right disk, and glue them along their corresponding nooses such that we again get a sphere. Alongside the sphere, we get a graph embedded in it that corresponds to the left and right sides of the separators induced by the nooses. It then remains to make the gluing so that the graph remains r-outerplanar, that is, it results in a graph embedded without edge crossings such that each vertex has a face path of length at most r to the outer face. For this we define a labeling for each edge bipartition and we keep only the largest subsequence of edge bipartitions that have the same labeling.

The labelings roughly work as follows. We want to use the labelings to ensure that the layer of each vertex in the layer decomposition of $G\langle A_i \rangle \circ G\langle B_j \rangle$ does not increase in comparison to G. For this, for each face touched by the nooses that correspond to (A_i, B_i) and (A_j, B_j), we note in the labelings how far it is away from the outer face (or, rather, the face in the sphere corresponding to the outer face in the plane), and we note for each pair of faces touched by the noose how far they are away from each other. Then, if two edge bipartitions have the same labeling, each vertex in the glued graph will be at most as far away from the faces touched by the noose and, hence, at most as far away from the outer face.

As we will see below, each edge-bipartition labeling can be encoded in $(32r^2 + 8r) \cdot \log(r+1) + 1$ bits. Thus, out of the $2\log(n)$ edge bipartitions that we obtain from the longest path in the decomposition tree, there are at least $\log(n)/(r+1)^{32r^2+8r}$ edge bipartitions with the same labeling.

The rest of this section is dedicated to the formal proof of Theorem 9. In the following, fix an arbitrary r-outerplanar embedding of G.

An initial sequence $\mathcal{T}$ of edge bipartitions. Consider the canonical embedding of G into a sphere $\mathfrak{S}$ that we obtain by taking a circle that encloses but does not intersect G and identifying all points in the unbounded region of the plane that is separated off by this circle. Since G is r-outerplanar, it has branchwidth at most $2r$ [5, Lemma 3]. By Theorem 2, there is a sphere-cut

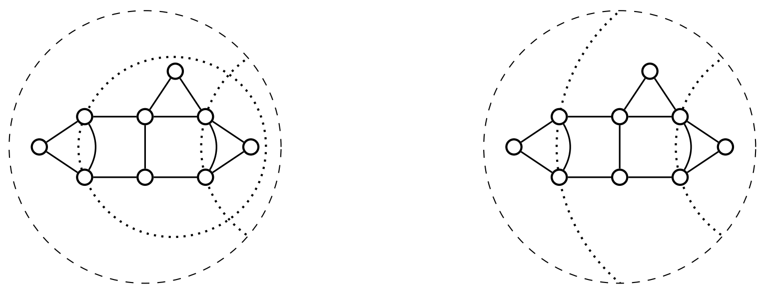

Figure 4: A graph embedded in the sphere and two crossing nooses (dotted, left) and two noncrossing nooses (dotted, right). We projected the sphere into the plane by replacing a point in the sphere with a circle (dashed) and drawing all remaining points inside this circle. Both pairs of nooses represent the same edge bipartitions. Note that the two nooses on the right share a point on the sphere.

branch decomposition (T, λ) for G of width at most $2r$. We define the sequence in Theorem 9 based on (T, λ).

Consider a longest path P in T. Note that the endpoints of P are leaves of T. Denote by e_1 that edge of G that is the preimage of the first vertex of P under the mapping λ of leaves of T to edges of G. Since each edge in T induces a bipartition of the edges in G, so does each edge on P. Define the sequence $\mathcal{T} := ((C_i, D_i))_{i=1}^{t}$, where (C_i, D_i) is the bipartition of $E(G)$ induced by the ith edge on P such that $e_1 \in C_i$. We have $C_i \subsetneq C_{i+1}$ and $D_i \supsetneq D_{i+1}$ because T is a ternary tree and λ is a bijection. We later need a lower bound on the length of $\mathcal{T}$. For this, observe that P contains at least $2\log(n)$ edges, because G contains at least n edges (there are no vertices of degree one) and T is a ternary tree. Thus, we have:

Lemma 5 *Sequence $\mathcal{T}$ is of length at least* $2\log(n)$.

The sequence in Theorem 9 is defined based on a subsequence of $\mathcal{T}$.

Obtaining a sequence of noncrossing nooses. As mentioned, we aim for the nooses associated with the edge bipartitions in $\mathcal{T}$ to be nested, that is, they separate the sphere of the embedding into left and right disks, and each left disk of some noose contains all left disks of nooses with smaller index. For this, we ensure that the nooses do not cross each other: Let $\mathfrak{N}_i$ and $\mathfrak{N}_j$ be two nooses of the sphere-cut branch decomposition associated with (C_i, D_i) and (C_j, D_j), respectively. Denote by $\mathfrak{C}_i, \mathfrak{D}_i$ the open disks in which $\mathfrak{N}_i$ separates $\mathfrak{S}$ such that $C_i \subseteq \mathfrak{C}_i$ and $D_i \subseteq \mathfrak{D}_i$ and analogously for j. The nooses $\mathfrak{N}_i$ and $\mathfrak{N}_j$ are *noncrossing* if it holds that $\mathfrak{C}_i \subsetneq \mathfrak{C}_j$ and $\mathfrak{D}_i \supsetneq \mathfrak{D}_j$. Otherwise the nooses *cross* each other.

To define our desired subsequence of $\mathcal{T}$, we choose one noose $\mathfrak{N}_i$ for each $(C_i, D_i) \in \mathcal{T}$ such that the resulting sequence of nooses is noncrossing. To see that we can choose the nooses in this way, first choose them arbitrarily and then consider two crossing nooses $\mathfrak{N}_i, \mathfrak{N}_j$, $i < j$, that is, $\mathfrak{C}_i \cap \mathfrak{D}_j \neq \emptyset$. We define a noose $\tilde{\mathfrak{N}}_i$ which we obtain from $\mathfrak{N}_i$ by replacing each maximal subsegment contained in $\mathfrak{D}_j$ by the corresponding subsegment of $\mathfrak{N}_j$ which is contained in $\mathfrak{C}_i$. There is no edge of G contained in $\mathfrak{C}_i \cap \mathfrak{D}_j$ because such an edge then would also be in $C_i \cap D_j \subseteq C_i \cap D_i$, a contradiction to the fact that C_i, D_i is a bipartition of $E(G)$. Hence, noose $\tilde{\mathfrak{N}}_i$ separates $\mathfrak{S}$ into two open disks $\tilde{\mathfrak{C}}_i, \tilde{\mathfrak{D}}_i$ such that $C_i = \tilde{\mathfrak{C}}_i \cap E(G)$ and $D_i = \tilde{\mathfrak{C}}_i \cap E(G)$. Thus, $\tilde{\mathfrak{N}}_i$ fulfills the conditions for the nooses in sphere-cut branch decompositions and we may choose $\tilde{\mathfrak{N}}_i$ for (C_i, D_i) instead of $\mathfrak{N}_i$.

Clearly, $\tilde{\mathfrak{N}}_i$ and $\mathfrak{N}_j$ are noncrossing. Moreover, any noose $\mathfrak{N}_k$, $k > i$, that crosses $\tilde{\mathfrak{N}}_i$ also crosses $\mathfrak{N}_i$ because $\tilde{\mathfrak{C}}_i \subseteq \mathfrak{C}_i$. Thus, by replacing $\mathfrak{N}_i$ with $\tilde{\mathfrak{N}}_i$, the number of pairs of crossing nooses with indices at least i is strictly decreased. This means that after a finite number of such replacements we reach a sequence of pairwise noncrossing nooses.

Labelings that allow gluing. Based on the sequence $\mathcal{T}$ of edge bipartitions of G and the nooses we have fixed above for each edge bipartition, we now define a tuple, the labeling, for each edge bipartition that can be encoded using $(32r^2 + 8r) \cdot \log(r+1) + 1$ bits and that has the property that, if two edge bipartitions have the same labeling, then the corresponding graphs can be glued in a way that results in an r-outerplanar graph, as stated in Theorem 9. As mentioned in the outline of the proof, the labeling essentially specifies for each vertex and each face on the corresponding noose, how far away it is from the outer face. Note that the corresponding face paths may cross the noose several times and hence we also need to encode the distance, in face-path length, between vertices and faces on the noose. These two components essentially make up the labeling. Furthermore, we need to take care to make sure that the labeling is being encoded using a number of bits that only depends on m and r.

We need some notation and definitions. Denote by $\mathfrak{F}$ the face in the sphere embedding of G that corresponds to the outer face of G in the planar embedding. Pick a point $y \in \mathfrak{F}$ in such a way that y is not equal to any vertex and not contained in any edge or noose $\mathfrak{N}_i$. For every noose $\mathfrak{N}_i$ we define a bijection $\beta_i \colon \{1, 2, \ldots, |M(C_i, D_i)|\} \to M(C_i, D_i)$ corresponding to the order in which the vertices in $M(C_i, D_i)$ appear in a clockwise (in the plane embedding) traversal of $\mathfrak{N}_i$ that starts in an arbitrary point. We furthermore define a mapping γ_i from $\{1, 2, \ldots, |M(C_i, D_i)|\}$ to the set of faces touched by $\mathfrak{N}_i$ according to their occurrence in the traversal of $\mathfrak{N}_i$ above. More precisely, if face $\mathfrak{G}$ occurs in the traversal of $\mathfrak{N}_i$ between vertex $\beta_i(\ell)$ and $\beta_i(\ell+1)$ (wherein we put $\beta_i(|M(C_i, D_i)|+1)$ equal to $\beta_i(1)$), then $\gamma_i(\ell) = \mathfrak{G}$. Finally, say that a face path P is *contained* in a closed disk $\mathfrak{E}$ if each vertex in P is contained in $\mathfrak{E}$.

The *labeling* of (C_i, D_i) is a tuple $(b, \Upsilon_1, \Upsilon_2)$ defined as follows.

- $b = 1$ if $y \in \mathfrak{C}_i$ and $b = 0$ otherwise (this encodes which of the disks $\mathfrak{C}_i$ or $\mathfrak{D}_i$ contains the "left" side of the graph).
- Υ_1 encodes the distance of the vertices and faces touched by the noose $\mathfrak{N}_i$ from the outer face $\mathfrak{F}$. Formally, it is defined as follows. Let $\beta, \gamma, \mathfrak{C}, \mathfrak{D}$ be symbols. Function Υ_1 maps each triple $(k, \xi, \mathfrak{X})$ such that $k \in \{1, 2, \ldots, |M(C_i, D_i)|\}$, $\xi \in \{\beta, \gamma\}$, and $\mathfrak{X} \in \{\mathfrak{C}, \mathfrak{D}\}$ to the length of a shortest face path that is contained in $\mathfrak{X}_i \cup \mathfrak{N}_i$ and that runs from $\xi_i(k)$ to $\mathfrak{F}$. (Herein, ξ_i refers to β_i if $\xi = \beta$ for example and analogously for $\mathfrak{X}$.) If such a path does not exist, or its length is larger than r, then put ∞ instead of the length.
  For instance, $\Upsilon_1(p, \gamma, \mathfrak{C})$ is the length of a shortest face path contained in $\mathfrak{C}_i$ from the pth face touched by $\mathfrak{N}_i$ to $\mathfrak{F}$ if such a path exists and is of length at most r.
- Υ_2 encodes the distance of the vertices and faces touched by the noose $\mathfrak{N}_i$ from each other. Formally, Υ_2 is the function that maps each quintuple $(k_1, k_2, \xi, \psi, \mathfrak{X})$ such that $k_1, k_2 \in \{1, 2, \ldots, |M(C_i, D_i)|\}$, $\xi, \psi \in \{\beta, \gamma\}$, and $\mathfrak{X} \in \{\mathfrak{C}, \mathfrak{D}\}$ to the length of a shortest face path that is contained in $\mathfrak{X}_i \cup \mathfrak{N}_i$ from $\xi_i(k_1)$ to $\psi_i(k_2)$. Again, if the path above does not exist, or its length is larger than r, then put ∞ instead of the length.

The desired edge bipartition sequence and its properties. Take

$$\mathcal{S} := ((C_i, D_i, \beta_i))_{i=1}^{s}$$

where, in a slight abuse of notation, $((C_i, D_i))_{i=1}^{s}$ is the longest subsequence of $\mathcal{T}$ in which all edge bipartitions (C_i, D_i) have the same labeling. Two edge bipartitions (defined via nooses) which have the same labeling are shown in Figure 4 and in Figure 5. We claim that $\mathcal{S}$ fulfills the conditions of Theorem 9. First, we consider the length.

Lemma 6 *The length s of $\mathcal{S}$ is at least* $\log(n)/(r+1)^{32r^2+8r}$.

Proof: By Lemma 5, $\mathcal{T}$ contains at least $2\log(n)$ entries. The longest subsequence of $\mathcal{T}$ with pairwise equal labelings has length at least $2\log(n)$ divided by the number of different labelings $(b, \Upsilon_1, \Upsilon_2)$. It is not hard to see that there are at most two possibilities for b, at most $(r+1)^{2r\cdot 2\cdot 2} = (r+1)^{8r}$ possibilities for Υ_1, and at most $(r+1)^{2r\cdot 2r\cdot 2\cdot 2\cdot 2} = (r+1)^{32r^2}$ possibilities for Υ_2, giving an overall upper bound on the number of different labelings of

$$2\cdot(r+1)^{8r}\cdot(r+1)^{32r^2} \;=\; 2\cdot(r+1)^{32r^2+8r}.$$

This shows the lemma. □

For each (C_i, D_i), $(C_j, D_j) \in \mathcal{S}$, $i < j$, we have $C_i \subsetneq C_j$ and $D_i \supsetneq D_j$. Thus to prove Theorem 9 it remains to show the following.

Lemma 7 *The graph $G_{ij} := G\langle C_i\rangle \circ G\langle D_j\rangle$ is r-outerplanar.*

The rest of this section contains the proof of Lemma 7. To this end, we first describe how to obtain an r-outerplanar embedding for a supergraph G' of G_{ij} from G's embedding in the sphere. Graph G' is defined below and is isomorphic to G_{ij} except that it may contain multiple copies of some edges in G_{ij}.

Recall that the nooses $\mathfrak{N}_i$ and $\mathfrak{N}_j$ are noncrossing. Hence the closed disks $\mathfrak{C}_i \cup \mathfrak{N}_i$ and $\mathfrak{D}_j \cup \mathfrak{N}_j$ can intersect only in their nooses $\mathfrak{N}_i$ and $\mathfrak{N}_j$. We now consider dislocating these disks from the sphere, and identifying their boundaries $\mathfrak{N}_i$ and $\mathfrak{N}_j$, thus creating a new sphere. Figure 5 serves as an example: Consider the left picture as a plane embedding of the sphere. Take $\mathfrak{C}_i$ to be the outer region of the outermost noose and $\mathfrak{D}_j$ to be the the inner region of the innermost noose. Then dislocate $\mathfrak{C}_i$ and $\mathfrak{D}_j$ from this sphere and glue them on their nooses. The right picture shows a plane embedding of a sphere that can be obtained in this way.

Using an arbitrary way to glue $\mathfrak{C}_i\cup\mathfrak{N}_i$ and $\mathfrak{D}_j\cup\mathfrak{N}_j$ may yield a graph with a larger outerplanarity number. We hence now specify how exactly we do the glueing operation. Observe that there is a homeomorphism $\phi\colon \mathfrak{C}_i \cup \mathfrak{N}_i \to \mathfrak{C}_j \cup \mathfrak{N}_j$ since both point sets are closed discs. Recall that, since (C_i, D_i) and (C_j, D_j) have the same labeling, we have that the vertex $y \in \mathfrak{F}$ which we fixed above is either both in $\mathfrak{C}_i$ and $\mathfrak{C}_j$ or both in $\mathfrak{D}_i$ and $\mathfrak{D}_j$. Intuitively, the nooses are thus nested inside each other in the plane. Recall furthermore that the vertices in $M(C_i, D_i)$ and $M(C_j, D_j)$ are enumerated by β_i and β_j, respectively, according to clockwise traversals of the corresponding nooses. Hence, we may choose the homeomorphism $\phi\colon \mathfrak{C}_i \cup \mathfrak{N}_i \to \mathfrak{C}_j \cup \mathfrak{N}_j$ such that it has the following two properties. Intuitively, it preserves the order of the vertices on the noose and it maps each vertex in $\mathfrak{N}_i$ to a vertex in $\mathfrak{N}_j$ with the same index.

(i) For the two traversals of the nooses that define β_i and β_j, respectively, we have that the initial points of the traversals are mapped onto each other by ϕ and, if point x comes before point z in the traversal of $\mathfrak{N}_i$ used to define β_i, then $\phi(z)$ comes after $\phi(x)$ in the traversal of $\mathfrak{N}_j$ used to define β_j.

(ii) For each $k \in \{1, 2, \ldots, |M(C_i, D_i)|\}$ we have $\phi(\beta_i(k)) = \beta_j(k)$.

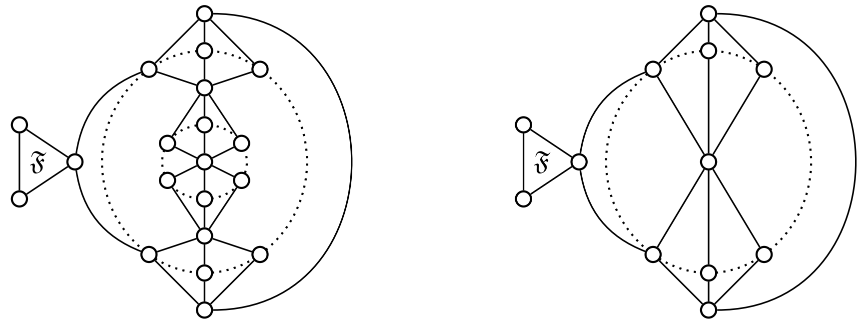

Figure 5: Left: A graph embedded in a subdisk of the sphere which has been projected onto the plane. We show two nooses (dotted) that induce edge bipartitions. The labelings of the two edge bipartitions are the same if we assume that both left sides (the C_i's) of the bipartitions contain the outermost edges in the drawing and if we furthermore assume that the corresponding mappings β_i's are the clockwise orderings of the vertices on the noose starting with the topmost vertex. Right: The graph resulting from gluing along the two nooses.

Denote by G' the $\mathfrak{S}$-plane graph induced by the point set $\phi(G \cap \mathfrak{C}_i) \cup (G \cap \mathfrak{D}_i)$. We claim that from G' we can derive an r-outerplanar embedding of G_{ij}.

We first prove that G_{ij} is an edge-induced subgraph of G' without loss of generality: We may assume that G' and G_{ij} have the same vertex set without loss of generality by Property (ii) of homeomorphism ϕ. Since each edge $e \in C_i$ is contained in $\mathfrak{C}_i$, it is also present in $\phi(\mathfrak{C}_i)$ and thus in G'. Moreover, each edge $e \in D_j$ is trivially contained in $\mathfrak{D}_j$, hence, also in G'. Thus, we may assume that G_{ij} is an edge-induced subgraph of G'. Thus each r-outerplanar embedding of G' induces an r-outerplanar embedding of G_{ij}.

It remains to show that G' is r-outerplanar. Graph G' has a sphere embedding due to the way it was constructed. We now prove that from this embedding we can obtain an r-outerplanar one. This then finishes the proof of Theorem 9. Note that there is a face in the sphere embedding of G' that contains y or $\phi(y)$ due to the flag b in the labelings (i.e., if $b = 1$ then there is a face containing $\phi(y)$, otherwise there is a face containing y). We denote this face by $\mathfrak{F}$. By removing a point contained in the face $\mathfrak{F}$ from the sphere, we obtain a topological space homeomorphic to the plane. Fix a corresponding homeomorphism δ and note that, applying δ to G', we obtain a planar embedding of G' with the outer face $\delta(\mathfrak{F})$. In the following we assume that G' is embedded in this way and, for the sake of simplicity, denote $\delta(\mathfrak{F})$ by $\mathfrak{F}$.

Recall that a graph is r-outerplanar if and only it has an embedding in the plane such that each vertex v has an incident face with a face path of length at most r to the outer face $\mathfrak{F}$. Call such a path *good* with respect to v.

We claim that each vertex in G' has a good path.

Lemma 8 *Each vertex in G' has a good path.*

Proof: It suffices to prove this for vertices in $\mathfrak{C}_i$ whose good paths in G are not contained in $\mathfrak{C}_i$ and vertices in $\mathfrak{D}_j$ whose good paths in G are not contained in $\mathfrak{D}_j$ as the remaining ones are also present in G'. Consider a vertex in $\mathfrak{C}_i$ whose good path P is not contained in $\mathfrak{C}_i$. Observe that each subpath of P that is not contained in $\mathfrak{C}_i$ is contained in $\mathfrak{D}_i \cup \mathfrak{N}_i$. We claim that we can replace every maximal face subpath of P which is contained in $\mathfrak{D}_i \cup \mathfrak{N}_i$ by a face path contained in $\mathfrak{D}_j \cup \mathfrak{N}_j$ in such a way that the resulting sequence P' is a face path in G'. Moreover, P' is at most as long as P.

Consider a maximal face subpath S of P which is contained in $\mathfrak{D}_i \cup \mathfrak{N}_i$. Each end of S is either a vertex in $M(C_i, D_i)$, or a face. If an end of S is a face, then it can either be the outer face $\mathfrak{F}$ or a face $\mathfrak{G} \neq \mathfrak{F}$ which is intersected by $\mathfrak{N}_i$. (Note that not both ends of S can be $\mathfrak{F}$ as P is a shortest path to $\mathfrak{F}$.) We now use the labelings to show that there is a suitable replacement for S. In each step we maintain that (i) each maximal subsegment of P contained in G' is a face path, (ii) the number of maximal segments of P that are not contained in G' strictly decreases, and (iii) the length of P does not increase.

Consider the case where one end of S is $\mathfrak{F}$; we will indeed only treat this case explicitly. The other case is analogous.

Consider the subcase where the other end of S is a vertex v. We aim to use the Υ functions of the labelings to find a replacement for S. To this end, associate with S the tuple $(k, \beta, \mathfrak{D})$, where β and $\mathfrak{D}$ are the corresponding symbols from the labeling of (C_i, D_i) (or equivalently of (C_j, D_j)). The first entry, k, is the index of the vertex v in the traversal of the noose $\mathfrak{N}_i$, that is, $k = \beta_i^{-1}(v)$. Let Υ_1^i be the function Υ_1 from the labeling of (C_i, D_i) and Υ_1^j the function Υ_1 from the labeling of (C_j, D_j). Since the two labelings are the same, $\Upsilon_1^i(k, \beta, \mathfrak{D}) = \Upsilon_2^j(k, \beta, \mathfrak{D}) \leq \ell$, where ℓ is the length of S (observe that $\ell \leq r$). Hence, there is a face path S' in $\mathfrak{D}_j$ with the ends $\mathfrak{F}$ and $\beta_j(k)$ and S' has length at most that of S. By Property (ii) of homeomorphism ϕ we have that $\beta_i(k) = \beta_j(k)$ in G'. Thus we can replace S by S' in P such that properties (i), (ii), and (iii) are maintained.

Now consider the subcase where the end of S that is different from $\mathfrak{F}$ is a face $\mathfrak{G}$. We again aim to use the Υ functions of the labelings to find a replacement for S. To this end, associate with S the tuple $(k, \gamma, \mathfrak{D})$, where γ and $\mathfrak{D}$ are the corresponding symbols from the labeling of (C_i, D_i). The first entry, k, intuitively is the index of an interval on $\mathfrak{N}_i$ between two vertices that corresponds to the face $\mathfrak{G}$. Formally, it is defined as follows. Draw an arc $\mathfrak{A}$ contained in $\mathfrak{G}$ between the two vertices that P visits before and after $\mathfrak{G}$ such that $\mathfrak{A}$ and $\mathfrak{N}_i$ intersect in only one point. Call this intersection point x. Define $k \in \mathbb{N}$ such that, in the traversal of $\mathfrak{N}_i$ that defines β_i, vertex $\beta_i(k)$ comes before x and vertex $\beta_i(k+1)$ comes after x (where we set $k+1 = 1$ if $k = |M(C_i, D_i)|$). Let Υ_1^i be the function Υ_1 from the labeling of (C_i, D_i) and Υ_1^j the function Υ_1 from the labeling of (C_j, D_j). Since the two labelings are the same, $\Upsilon_1^i(k, \gamma, \mathfrak{D}) = \Upsilon_1^j(k, \gamma, \mathfrak{D}) \leq \ell$, where ℓ is the length of S (observe that $\ell \leq r$). Hence, by the definition of the labelings, there is a face path S' in $\mathfrak{D}_j \cup \mathfrak{N}_j$ with the ends $\mathfrak{F}$ and $\gamma_j(k)$ of length at most that of S.

We claim that we can replace S by S', maintaining the three properties. By the definition of S', the length does not increase and thus property (iii) is maintained. Clearly, property (ii) is maintained as well. We claim that also property (i) holds. That is, S' and the segment of P before that is contained in $\mathfrak{C}_i$ form a valid face path. Note that the segment on $\mathfrak{N}_i$ occurring between vertex $\beta_i(k)$ and $\beta_i(k+1)$ and that of $\mathfrak{N}_j$ occurring between vertex $\beta_j(k)$ and $\beta_j(k+1)$ have been identified during gluing $\mathfrak{C}_i \cup \mathfrak{N}_i$ and $\mathfrak{D}_j \cup \mathfrak{N}_j$ due to the two properties of the homeomorphism ϕ. Thus the face of G' occurring on $\mathfrak{N}_i$ between vertex $\beta_i(k)$ and $\beta_i(k+1)$ consists of the union of the faces $\gamma(k)_i$ and $\gamma(k)_j$ of G, intersected with $\mathfrak{C}_i \cup \mathfrak{N}_i$ and $\mathfrak{D}_j \cup \mathfrak{N}_j$, respectively. (Recall that a noose intersects each face at most once.) Hence, property (iii) holds as well.

The proof that we can replace S by a corresponding path S' in P in the case that S does not have $\mathfrak{F}$ as an end is analogous to the above: Note that, in this case both ends are a vertex or face intersected by the noose $\mathfrak{N}_i$. We may then use the labelings in an analogous way to find a replacement, relying instead of the Υ_1 functions on the Υ_2 functions.

Hence, replacing all maximal face subpaths of P that are not contained in $\mathfrak{C}_i$, we obtain a good path in G'. Finally, the case that the good path of a vertex in $\mathfrak{D}_j$ is not contained in $\mathfrak{C}_i$ is symmetric to the above and also omitted. $\square$

Summarizing, each vertex in G' has a good path by Lemma 8, meaning that G' is r-outerplanar. Since G_{ij} is an edge-induced subgraph of G', also G_{ij} is r-outerplanar. This concludes the proof of Lemma 7 and thus of Theorem 9.

7 Concluding remarks

The main contribution of this work is to show that twins are crucial for instances of r-Outerplanar Support but the number of crucial twins is upper-bounded in terms of the number m of hyperedges and the outerplanarity number r of a support. As a result, we can safely remove noncrucial twins. More specifically, in linear time we can transform any instance of r-Outerplanar Support into an equivalent one whose size is upper-bounded by a function of m and r only. In turn, this implies fixed-parameter tractability with respect to $m + r$. It is fair to say that due to the strong exponential growth in m and r this result is mainly of classification nature. Improved bounds (perhaps based on further data reduction rules) are highly desirable for practical applications.

Two further directions for future research are as follows. First, above we only showed how to shrink the size of the input instance. We also need an efficient algorithm to construct an r-outerplanar support for such an instance. A naive algorithm for this task has running time $n^{\mathrm{O}(n)}$: Since every planar graph on n vertices has $\mathrm{O}(n)$ edges, we may enumerate all planar graphs in $n^{\mathrm{O}(n)}$ time by considering all possible endpoints for each edge. For each enumerated graph, we then test whether one of them is an r-outerplanar support. Can we improve over this brute-force algorithm?

Second, it is interesting to gear the parameters under consideration more towards practice. In Section 5 above we attached signatures to each edge bipartition in a sequence of edge bipartitions of a support and we could reduce our input only if there were sufficiently many edge bipartitions with the same signature. This signature contained, among other information, the twin class of each vertex of the separator induced by the edge bipartition. Clearly, if all of these at least 2^{mr} different types of signatures are present, then this will lead to an illegible drawing of the hypergraph (and still, in absence of better upper bounds, we cannot shrink our input). It seems thus worthwhile to contemplate parameters that capture legibility of the hypergraph drawing by restricting further the number of possible signatures. Furthermore, it would be interesting to consider hypergraphs that have bounded VC-dimension or hypergraphs that stem from other restricted settings such as hypergraphs defined via geometric configurations.

Finally, an obvious open question is whether finding a *planar* support is (linear-time) fixed-parameter tractable with respect to the number m of hyperedges only. A promising direction might be to show that there is a planar representative support (as in Definition 10) which has treewidth upper-bounded by a function of m. From this, we would get a sequence of gluable subgraphs similarly to the one we have used here, amenable to the same approach as in Section 5.

Acknowledgements

We thank anonymous referees for carefully reading the manuscript and suggesting improvements to the presentation. We thank an anonymous referee who pointed out the ideas for the proof in Section 4.

Funding. Significant parts of the work were done while all authors were with TU Berlin and while RvB, IK, and MS were supported by the DFG project NI 369/12. CK was supported by DFG project KO 3669/4-1. MS was supported by the People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme (FP7/2007-2013) under REA grant agreement number 631163.11, by the Israel Science Foundation (grant no. 551145/14), and by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme under grant agreement number 714704, and by the Alexander von Humboldt Foundation. Some work of MS was done while with Ben-Gurion University of the Negev and with University of Warsaw.

References

[1] B. Alsallakh, L. Micallef, W. Aigner, H. Hauser, S. Miksch, and P. J. Rodgers. The state-of-the-art of set visualization. *Computer Graphics Forum*, 35(1):234–260, 2016. doi:10.1111/cgf.12722.

[2] J. C. Athenstädt and S. Cornelsen. Planarity of overlapping clusterings including unions of two partitions. *Journal of Graph Algorithms and Applications*, 21(6):1057–1089, 2017. doi:10.7155/jgaa.00450.

[3] A. S. Azarenok and V. I. Sarvanov. O slozhnosti planarnoy realizatsii gipergrafa [on the complexity of the planar realization of hypergraphs]. *Izvestiya Akademii nauk Belorusskoi SSR*, 4:10–12, 1987.

[4] C. Beeri, R. Fagin, D. Maier, and M. Yannakakis. On the desirability of acyclic database schemes. *Journal of the ACM*, 30(3):479–513, 1983. doi:10.1145/2402.322389.

[5] T. Biedl. On triangulating k-outerplanar graphs. *Discrete Applied Mathematics*, 181:275–279, 2015. doi:10.1016/j.dam.2014.10.017.

[6] H. L. Bodlaender. Kernelization: New upper and lower bound techniques. In *Proceedings of the 4th International Workshop on Parameterized and Exact Computation (IWPEC '09)*, volume 5917 of *Lecture Notes in Computer Science*, pages 17–37. Springer, 2009.

[7] U. Brandes, S. Cornelsen, B. Pampel, and A. Sallaberry. Blocks of hypergraphs—Applied to hypergraphs and outerplanarity. In C. S. Iliopoulos and W. F. Smyth, editors, *Proceedings of the 21st International Workshop on Combinatorial Algorithms (IWOCA '10)*, volume 6460 of *LNCS*, pages 201–211. Springer Berlin Heidelberg, 2011. doi:10.1007/978-3-642-19222-7_21.

[8] U. Brandes, S. Cornelsen, B. Pampel, and A. Sallaberry. Path-based supports for hypergraphs. *Journal of Discrete Algorithms*, 14:248–261, 2012. doi:10.1016/j.jda.2011.12.009.

[9] K. Buchin, M. J. van Kreveld, H. Meijer, B. Speckmann, and K. Verbeek. On planar supports for hypergraphs. *Journal of Graph Algorithms and Applications*, 15(4):533–549, 2011. doi:10.7155/jgaa.00237.

[10] J. Chen, C. Komusiewicz, R. Niedermeier, M. Sorge, O. Suchý, and M. Weller. Polynomial-time data reduction for the subset interconnection design problem. *SIAM Journal on Discrete Mathematics*, 29(1):1–25, 2015. doi:10.1137/140955057.

[11] L. E. Dickson. Finiteness of the odd perfect and primitive abundant numbers with n distinct prime factors. *American Journal of Mathematics*, 35(4):413–422, 1913. doi:10.2307/2370405.

[12] W. Didimo, F. Giordano, and G. Liotta. Overlapping cluster planarity. *Journal of Graph Algorithms and Applications*, 12(3):267–291, 2008. doi:10.7155/jgaa.00167.

[13] R. Diestel. *Graph Theory*, volume 173 of *GTM*. Springer, 5th edition, 2016.

[14] F. Dorn, E. Penninkx, H. L. Bodlaender, and F. V. Fomin. Efficient exact algorithms on planar graphs: Exploiting sphere cut decompositions. *Algorithmica*, 58(3):790–810, 2010. doi:10.1007/s00453-009-9296-1.

[15] T. Eschbach, W. Günther, and B. Becker. Orthogonal hypergraph drawing for improved visibility. *Journal of Graph Algorithms and Applications*, 10(2):141–157, 2006. doi:10.7155/jgaa.00122.

[16] V. Feinberg, A. Levin, and E. Rabinovich. *Hypergraph Planarization*, pages 45–86. Springer Netherlands, Dordrecht, 1997. doi:10.1007/978-94-011-5740-7_4.

[17] Q.-W. Feng, R. F. Cohen, and P. Eades. Planarity for clustered graphs. In *Proceedings of the 3rd European Symposium on Algorithms (ESA '95)*, volume 979 of *Lecture Notes in Computer Science*, pages 213–226. Springer, 1995. doi:10.1007/3-540-60313-1_145.

[18] J. Flower, A. Fish, and J. Howse. Euler diagram generation. *Journal of Visual Languages & Computing*, 19(6):675–694, 2008. doi:10.1016/j.jvlc.2008.01.004.

[19] R. Fulek and C. D. Tóth. Atomic embeddability, clustered planarity, and thickenability. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA '20)*, pages 2876–2895. Society for Industrial and Applied Mathematics, 2020. doi:10.1137/1.9781611975994.175.

[20] Q. Gu and H. Tamaki. Optimal branch-decomposition of planar graphs in $o(n^3)$ time. *ACM Transactions on Algorithms*, 4(3), 2008. doi:10.1145/1367064.1367070.

[21] J. Guo and R. Niedermeier. Invitation to data reduction and problem kernelization. *ACM SIGACT News*, 38(1):31–45, 2007.

[22] M. Habib, C. Paul, and L. Viennot. Partition refinement techniques: An interesting algorithmic tool kit. *International Journal of Foundations of Computer Science*, 10(2):147–170, 1999. doi:10.1142/S0129054199000125.

[23] F. Hurtado, M. Korman, M. van Kreveld, M. Löffler, V. Sacristán, A. Shioura, R. I. Silveira, B. Speckmann, and T. Tokuyama. Colored spanning graphs for set visualization. *Computational Geometry*, 68:262–276, 2018. doi:10.1016/j.comgeo.2017.06.006.

[24] D. S. Johnson and H. O. Pollak. Hypergraph planarity and the complexity of drawing Venn diagrams. *Journal of Graph Theory*, 11(3):309–325, 1987. doi:10.1002/jgt.3190110306.

[25] M. Kaufmann, M. J. van Kreveld, and B. Speckmann. Subdivision drawings of hypergraphs. In I. G. Tollis and M. Patrignani, editors, *Proceedings of the 16th International Symposium on Graph Drawing (GD '08)*, volume 5417 of *LNCS*, pages 396–407. Springer Berlin Heidelberg, 2008. doi:10.1007/978-3-642-00219-9_39.

[26] B. Klemz, T. Mchedlidze, and M. Nöllenburg. Minimum tree supports for hypergraphs and low-concurrency euler diagrams. In R. Ravi and I. L. Gørtz, editors, *Proceedings of the 14th Scandinavian Symposium on Algorithm Theory (SWAT '14)*, volume 8503 of *LNCS*, pages 265–276. Springer International Publishing, 2014. doi:10.1007/978-3-319-08404-6_23.

[27] E. Korach and M. Stern. The clustering matroid and the optimal clustering tree. *Mathematical Programming*, 98(1-3):385–414, 2003. doi:10.1007/s10107-003-0410-x.

[28] S. Kratsch. Recent developments in kernelization: A survey. *Bulletin of the EATCS*, 113, 2014.

[29] G. D. Lozzo, D. Eppstein, M. T. Goodrich, and S. Gupta. C-planarity testing of embedded clustered graphs with bounded dual carving-width. In *Proceedings of the 14th International Symposium on Parameterized and Exact Computation (IPEC '19)*, volume 148 of *LIPIcs*, pages 9:1–9:17. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2019. doi:10.4230/LIPIcs.IPEC.2019.9.

[30] D. Marx and M. Pilipczuk. Optimal parameterized algorithms for planar facility location problems using voronoi diagrams. In N. Bansal and I. Finocchi, editors, *Proceedings of the 23rd Annual European Symposium on Algorithms (ESA '15)*, pages 865–877. Springer, 2015. doi:10.1007/978-3-662-48350-3_72.

[31] O. I. Melnikov. Realization of hypergraphs by trees of minimal diameter. *Discrete Mathematics and Applications*, 7(3):279–286, 1997. doi:10.1515/dma.1997.7.3.279.

[32] J. Munkres. *Topology*. Prentice Hall, 2nd edition edition, 2000.

[33] E. Mäkinen. How to draw a hypergraph. *International Journal of Computer Mathematics*, 34:178–185, 1990. doi:10.1080/00207169008803875.

[34] T. E. Piliposyan. O planarno realizuemykh gipergrafakh [on planarly realizable hypergraphs]. *Doklady Akademii nauk Armyanskoi SSR*, 82:102–106, 1986. English abstract in zbMATH 0639.05038. URL: https://arar.sci.am/dlibra/publication/49904/edition/44755.

[35] P. D. Seymour and R. Thomas. Call routing and the ratcatcher. *Combinatorica*, 14(2):217–241, 1994. doi:10.1007/BF01215352.

[36] R. Swaminathan and D. K. Wagner. On the consecutive-retrieval problem. *SIAM Journal on Computing*, 23(2):398–414, 1994. doi:10.1137/S0097539792235487.

[37] R. Tarjan and M. Yannakakis. Simple linear-time algorithms to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs. *SIAM Journal on Computing*, 13(3):566–579, 1984. doi:10.1137/0213035.

[38] R. van Bevern, I. Kanj, C. Komusiewicz, R. Niedermeier, and M. Sorge. Twins in subdivision drawings of hypergraphs. In *Proceedings of the 24th International Symposium on Graph Drawing and Network Visualization (GD '16)*, volume 9801 of *LNCS*, pages 67–80. Springer, 2017. URL: http://arxiv.org/abs/1511.09389v3, doi:10.1007/978-3-319-50106-2_6.

[39] A. Verroust and M. Viaud. Ensuring the drawability of extended euler diagrams for up to 8 sets. In A. F. Blackwell, K. Marriott, and A. Shimojima, editors, *Proceedings of the Third International Conference on Diagrammatic Representation and Inference (Diagrams '04)*, volume 2980 of *LNCS*, pages 128–141. Springer Berlin Heidelberg, 2004. doi:10.1007/978-3-540-25931-2_13.

[40] A. A. Voloshina and V. Z. Feinberg. O planarnosti gipergrafov [on the planarity of hypergraphs]. *Doklady Akademii nauk Belorusskoi SSR*, 28(4):309–311, 1984.

[41] A. A. Zykov. Gipergrafy [hypergraphs]. *Uspekhi matematicheskikh nauk*, 29(6):89–154, 1974. In Russian, English translation in Mathematical Surveys 29.6. doi:10.1070/rm1974v029n06abeh001303.