Twins in Subdivision Drawings of Hypergraphs

René van Bevern[1,2], Iyad Kanj[3], Christian Komusiewicz[4], Rolf Niedermeier[5], and Manuel Sorge[5(✉)]

[1] Novosibirsk State University, Novosibirsk, Russian Federation
rvb@nsu.ru
[2] Sobolev Institute of Mathematics, Siberian Branch of the Russian Academy of Sciences, Novosibirsk, Russian Federation
[3] DePaul University, Chicago, USA
ikanj@cs.depaul.edu
[4] Friedrich-Schiller-Universität Jena, Jena, Germany
christian.komusiewicz@uni-jena.de
[5] TU Berlin, Berlin, Germany
{rolf.niedermeier,manuel.sorge}@tu-berlin.de

Abstract. Visualizing hypergraphs, systems of subsets of some universe, has continuously attracted research interest in the last decades. We study a natural kind of hypergraph visualization called *subdivision drawings*. Dinkla et al. [Comput. Graph. Forum '12] claimed that only few hypergraphs have a subdivision drawing. However, this statement seems to be based on the assumption (also used in previous work) that the input hypergraph does not contain *twins*, pairs of vertices which are in precisely the same hyperedges (subsets of the universe). We show that such vertices may be necessary for a hypergraph to admit a subdivision drawing. As a counterpart, we show that the number of such "necessary twins" is upper-bounded by a function of the number m of hyperedges and a further parameter r of the desired drawing related to its number of layers. This leads to a linear-time algorithm for determining such subdivision drawings if m and r are constant; in other words, the problem is linear-time fixed-parameter tractable with respect to the parameters m and r.

1 Introduction

Hypergraph drawings are useful as visual aid in diverse applications [1], among them electronic circuit design [11] and relational databases [2,18]. There are several methods for embedding hypergraphs in the plane. The combinatorial problem studied in this work stems from obtaining *subdivision drawings* [14,15].

Herein, 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* [14]. Figure 1 shows an example for such a drawing.
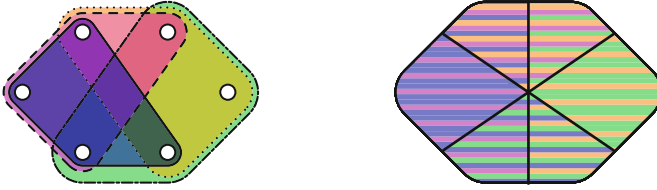

Fig. 1. Two drawings of the same hypergraph. On the left, we see a drawing in the *subset standard* in which the vertices (white circles) are enclosed by curves that correspond to hyperedges. On the right, we see a *subdivision drawing* in which we assign vertices to regions (enclosed by black lines) 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. (Color figure online)

Subdivision drawings are a natural extension of planarity for ordinary graphs: A graph is planar if and only if it has a subdivision drawing when viewed as a hypergraph. For hypergraphs, having a subdivision drawing is a rather general concept of planar embeddings, as, for example, each Zykov planar hypergraph (meaning that the incidence graph is planar) and each hypergraph with a well-formed Euler diagram (see Flower et al. [12]) has a subdivision drawing. Still, Dinkla et al. [10] claimed that "most hypergraphs do not have [subdivision drawings]". However, this claim might have been based on the fact that several works on subdivision drawings assumed that the input hypergraph is twinless, that is, there are no two vertices contained in precisely the same hyperedges (see Mäkinen [18, p. 179], Buchin et al. [6, p. 535], and Kaufmann et al. [15, p. 399]). Twins do not seem useful at first glance: whatever role one vertex can play to obtain a subdivision drawing, its twin can also fulfill. One of our contributions is disproving the general validity of this assumption in Sect. 3. More specifically, we give a hypergraph with two twins that has a subdivision drawing but, removing one twin, it ceases to have one. Thus, twins may indeed be helpful to find a solution.

More generally, we can construct hypergraphs with ℓ twins that allow for subdivision drawings 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 hypergraph with m hyperedges, we can forget all but $\psi(m)$ twins while maintaining the property of having a subdivision drawing. Using well-quasi orderings, one can relatively easily prove the *existence* of such a function ψ, 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 r measuring the number of "layers" in the drawing. A small number r of layers, however, is a relevant special case [4,6].

We study subdivision drawings from a combinatorial point of view, exploiting the fact that it is equivalent for a hypergraph to have a subdivision drawing and to have a support that is planar [14]. Herein, 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 $G[F]$. The outerplanarity number r of the support roughly translates to the number of layers in a corresponding drawing:[1] An r-outerplanar graph admits a planar embedding (without edge crossings) which has the property that, after removing r times all vertices on the outer face, we obtain an empty graph. Similar restrictions were studied before [4,6]. Formally, we study the following problem.

Problem (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?

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.*

In contrast to Theorem 1, r-Outerplanar Support remains NP-complete for $r = \infty$ [14] and even for every fixed $r > 1$ [6] (see below). The constants in the running time of the algorithm in Theorem 1 have a large dependence on m and r. However, 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. Thus, it makes sense to design algorithms particularly for hypergraphs with a small number of hyperedges, as done by Verroust and Viaud [21]. 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 [6].

Related Work. For specifics on the relations of some different planar embeddings for hypergraphs, see Kaufmann et al. [15], Brandes et al. [4].

As mentioned before, Johnson and Pollak [14] showed that finding a *planar* support is NP-complete. Buchin et al. [6] 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. use. Given a formula ϕ

---

[1] We refer to Kaufmann et al. [15] for a method to obtain a subdivision drawing from a planar support.

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 an r-outerplanar support, $r \geq 3$, exists also decides the satisfiability of the corresponding formula.

Towards determining the computational complexity of finding an *outerplanar* hypergraph support, Buchin et al. [4] 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 either a singleton or again a hyperedge in the hypergraph. Getting even more special, a tree support can be found in linear time [2,19]. Buchin et al. [6] gave a polynomial-time algorithm that can deal with an additional upper bound on the vertex degrees in the tree support. Klemz et al. [16] 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 [16,17].

In a wider scope, motivated by drawing metro maps and metro map-like diagrams, Brandes et al. [5] studied the problem of finding *path-based* planar hypergraph supports, that is, planar supports that fulfill the additional constraint that the subgraph induced by each hyperedge contains a Hamiltonian path, giving NP-hardness and tractability results. Finding path-based tree supports is also known as the GRAPH REALIZATION problem, for which several polynomial-time algorithms were already known [3].

Chen et al. [7] 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. The proof is quite different, however.

Organization. In Sect. 2 we provide some technical preliminaries used throughout the work. In Sect. 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 (a proof is deferred to a full version). In Sect. 4 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 promised in Theorem 1. We conclude and give some directions for future research in Sect. 5.

2   Preliminaries

General Notation. 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 tuple $(V, \mathcal{E})$ 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 $v \in V(\mathcal{H})$, we 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 $\forall u, v \in V : (u, v) \in \tau \iff \mathcal{E}(u) = \mathcal{E}(v)$ is an equivalence relation. The equivalence classes $[u]_\tau$, $u \in V$, 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 *subhypergraph 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 [9]. In particular, a *bridge* of a graph is an edge whose removal increases its number of connected components. Analogously, a *cut-vertex* is a vertex whose removal increases its number of connected components. Some special notation including the gluing of graphs is given below. We use the usual notation for planar and plane graphs. An *r-outerplanar* graph admits a planar embedding which has the property that, after r times of removing all vertices on the outer face, we obtain an empty graph. The *ith layer* L_i of a plane graph is defined as the set of vertices on the outer face, after having $i - 1$ times removed all vertices on the outer face.

Boundaried Graphs and Gluing. For a nonnegative integer $b \in \mathbb{N}$, a *b-boundaried graph* is a tuple (G, B, β) where G is a graph, $B \subseteq V(G)$ such that $|B| = b$, and $\beta \colon B \to \{1, \dots, b\}$ 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, we also denote by *β-boundaried* graph G that b-boundaried graph G whose boundary is the domain of β and whose boundary labeling is β.

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 $v \in B_1$ with $\beta_2^{-1}(\beta_1(v)) \in B_2$. We omit the index b in $\circ_b$ if it is clear from the context.

3 Beware of Removing Twins

In Fig. 2, we provide a concrete example that shows that twins can be necessary
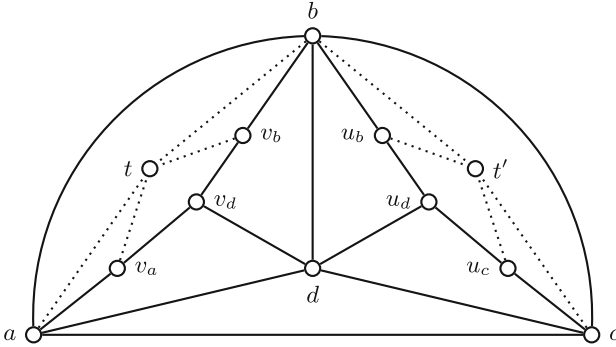to obtain a 2-outer-planar support:

Fig. 2. A hypergraph $\mathcal{H}$ and its support, showing that twins can be essential for obtain-
ing a 2-outer-planar support. The set of hyperedges consists of size-two hyperedges
that are drawn as solid lines between the corresponding vertices and, additionally,
$\{a, v_a, t, t', c\}$, $\{a, v_b, t, t', c\}$, $\{b, v_a, t, t', c\}$, $\{b, v_b, t, t', c\}$, $\{b, u_b, t, t', a\}$, $\{b, u_c, t, t', a\}$,
$\{c, u_b, t, t', a\}$, and $\{c, u_c, t, t', a\}$. Note that the vertices t and t' are twins. The hyper-
graph $\mathcal{H}$ 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.

The vertex set of the hypergraph $\mathcal{H}$ shown in Fig. 2 is $V := \{a, b, c, d, v_a, v_b,$
$v_d, u_b, u_c, u_d, t, t'\}$. We choose the hyperedges in such a way that t and t' are twins
and $\mathcal{H}$ has a planar (more precisely, 2-outerplanar) support but $\mathcal{H} - t$ does not.
First, we add to the set of hyperedges $\mathcal{E}$ of $\mathcal{H}$ the size-two hyperedges represented
by solid lines between the corresponding vertices in Fig. 2. The corresponding
"solid" hyperedges incident with (and only with) a, b, c, d form a K_4 and have the
purpose of essentially fixing the embedding of each support G: Since the complete
graph on four vertices, K_4, is 3-connected, it has only one planar embedding up to
the choice of the outer face [20, p. 747]. The remaining solid hyperedges (incident
with v_a, v_b, v_d and u_a, u_c, u_d) have the purpose of anchoring the u- and v-vertices
within two different faces of the embedding of the K_4: These hyperedges form
two connected components that are adjacent to a, b, d and b, c, d, respectively.
Hence, these connected components reside in those (unique) faces of the K_4 that
are incident with a, b, d and b, c, d, respectively.

With the following additional hyperedges, our goal is to enforce that t and
t' are used as conduits to connect the v-vertices to c via both a and b, and to
connect the u-vertices to a via both b and c. As we explain below, this is achieved
by the following hyperedges:

$$\{a, v_a, t, t', c\}, \quad \{a, v_b, t, t', c\}, \quad \{b, v_a, t, t', c\}, \quad \{b, v_b, t, t', c\},$$
$$\{b, u_b, t, t', a\}, \quad \{b, u_c, t, t', a\}, \quad \{c, u_b, t, t', a\}, \quad \{c, u_c, t, t', a\}.$$

Clearly, t and t' are twins. As can easily be verified, adding t and t' and the dotted edges in Fig. 2 to the graph induced by the solid edges gives a planar support for $\mathcal{H}$.

We now show that t and t' have to reside in different faces for each planar support G for $\mathcal{H}$. First, observe that, in G, either v_a is not adjacent to b or v_b is not adjacent to a. Moreover, neither of v_a and v_b is adjacent to c. Thus, to connect the subgraphs induced by the hyperedges that contain v_a or v_b, either vertex t or its twin t' must be adjacent to one of the two vertices in G. For the same reason, one of t and t' must be adjacent to one of u_b and u_c. Since there is no face in G that is simultaneously incident with one of v_a or v_b and one of u_b or u_c, the twins t and t' thus have to be in different faces. This implies that it is impossible to obtain a planar support if t or t' is missing. Consequently, removing one vertex of a twin class can transform a yes-instance of r-OUTERPLANAR SUPPORT into a no-instance.

The example above is not a pathology of having only one pair of twins, in a full version, we extend it so that an arbitrarily large set of twins is required for the existence of a planar support.

4   Relevant Twins for r-Outerplanar Supports

In this section, we show that there is an explicit function ψ such that, out of each twin class of a given 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 1. 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 1 is correct, Theorem 1 follows.

Proof (Theorem 1). Rule 1 can be applied exhaustively in linear time because the twin classes can be computed in linear time [13]. 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 the terms of parameterized algorithmics, exhaustive application of Rule 1 can be seen as a problem kernel (see [8], for example).

The correctness proof for Rule 1 consists of two parts. First, in Theorem 2, 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 support for our input hypergraph. We then show that, in a long sequence of nested separators for this support as above, 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 is large enough, that is, larger than some function of m and r, then there is always at least one non-crucial vertex which can be removed.

We now formalize our approach. Theorem 2 will guarantee the existence of a long sequence of gluable separators. To formally state it, we need the following notation.

Definition 1. *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,*

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

We call $M(A, B)$ the middle set *of A, B. For an edge set $A \subseteq E(G)$, denote by $G\langle A \rangle := (\bigcup_{e \in A} e, A)$ the subgraph induced by A.*

Recall from Sect. 2 the definitions of graph gluing, boundary, and boundary labeling.

Theorem 2 ($\star$[2]). *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 $\beta_i \colon M(A_i, B_i) \to \{1, \ldots, |M(A_i, B_i)|\}$ 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, where $G\langle A_i \rangle$ is understood to be β_i-boundaried and $G\langle B_j \rangle$ is understood to be β_j-boundaried.

To gain some intuition for Theorem 2 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 sides 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 inbetween 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 because 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

---

[2] Results labeled by $\star$ are deferred to a full version of the paper.

such a structure induces many nested separators from which we can glean 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.

Definition 2 (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 2, 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 whose vertices in between we can remove and reattach as non-crucial vertices, that is, vertices not 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.

Definition 3 (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 tuple where (A, B) is an edge bipartition of G, and $\beta \colon M(A, B) \to \{1, \ldots, |M(A, B)|\}$. Let $\ell := |M(A, B)|$. The* signature *of (A, B, β) is a triple $(\mathcal{T}, \phi, \mathcal{C})$, where*

- *$\mathcal{T} := \{[u]_\tau \mid u \in \bigcup A\}$ is the set of twin classes 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 $\beta^{-1}(i), \beta^{-1}(j) \in F$ and $\beta^{-1}(i)$ is connected to $\beta^{-1}(j)$ in $G\langle B \rangle[F \cap \bigcup B]$. Herein, $G\langle B \rangle[F \cap \bigcup B]$ is the subgraph of $G\langle B \rangle$ induced by $F \cap \bigcup B$.*

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

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

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

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 2 to G. For this we need that G is connected and does not contain any bridges. Indeed, if G is not connected, then add edges between its connected components in a tree-like fashion. 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 2 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 plane graph in G and, since $|M(A_i, B_i)| \leq 2r$, induces at most $\max\{1, 3|M(A_i, B_i)| - 6\} \leq \max\{1, 6r - 6\}$ edges. 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 $G_{ij} := G\langle A_i \rangle \circ G\langle B_j \rangle$, wherein $G\langle A_i \rangle$ is β_i-boundaried and $G\langle B_j \rangle$ is β_j-boundaried. Let $W' := V(G_{ij})$, where we assume that $W' \cap M(A_j, B_j) \subseteq M(A_i, B_i)$ for the sake of a simpler notation. Note that $W \setminus W' \neq \emptyset$ 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 by some vertex in W' in $\mathcal{H}$ and that G_{ij} is a support for $\mathcal{H}|_{W'}$. Since G_{ij} is r-outerplanar by Theorem 2, Statement (iii), this contradicts the choice of G according to the minimum number of vertices, thus proving the lemma.

To prove that each vertex $V \setminus W'$ is covered by some vertex in W', we show that $\{[u]_\tau \mid u \in V\} = \{[u]_\tau \mid u \in W'\}$. 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 $M(A_i, B_j) = S_i = S_j$ of (A_i, B_j) in 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 each vertex in F is 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. $\square$

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 $F \cap W$. $\square$

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 1 is correct.

Proof (Correctness of Rule 1). Consider an instance $\mathcal{H} = (V, \mathcal{E})$ of r-OUTERPLANAR SUPPORT to which Rule 1 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, and v can be safely removed from $\mathcal{H}$.   □

5   Concluding Remarks

The main contribution of this work is to show that twins may be 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 outer-planarity number r of a support. As a result, we can safely remove non-crucial 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, however, 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 reduce the size of the input instance. We also need an efficient algorithm to construct an r-outerplanar support for such an instance. As a first step, it would be interesting to improve on the $n^{O(n)}$-time brute-force algorithm

that simply enumerates all n-vertex planar graphs and tests whether one of them is an r-outerplanar support.[3]

Second, it is interesting to gear the parameters under consideration more towards practice. In Sect. 4 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, this will lead to an illegible drawing of the hypergraph (and still, in absence of better upper bounds, we cannot reduce our input). It seems thus worthwhile to contemplate parameters that capture legibility of the hypergraph drawing by restricting further the number of possible signatures.

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 2) 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 Sect. 4.

References

1. Alsallakh, B., Micallef, L., Aigner, W., Hauser, H., Miksch, S., Rodgers, P.J.: The state-of-the-art of set visualization. Comput. Graph. Forum **35**(1), 234–260 (2016)
2. Beeri, C., Fagin, R., Maier, D., Yannakakis, M.: On the desirability of acyclic database schemes. J. ACM **30**(3), 479–513 (1983)
3. Bixby, R.E., Wagner, D.K.: An almost linear-time algorithm for graph realization. Math. Oper. Res. **13**(1), 99–123 (1988)
4. Brandes, U., Cornelsen, S., Pampel, B., Sallaberry, A.: Blocks of hypergraphs—applied to hypergraphs and outerplanarity. In: Iliopoulos, C.S., Smyth, W.F. (eds.) IWOCA 2010. LNCS, vol. 6460, pp. 201–211. Springer, Heidelberg (2011). doi:10.1007/978-3-642-19222-7_21
5. Brandes, U., Cornelsen, S., Pampel, B., Sallaberry, A.: Path-based supports for hypergraphs. J. Discrete Algorithms **14**, 248–261 (2012)
6. Buchin, K., van Kreveld, M.J., Meijer, H., Speckmann, B., Verbeek, K.: On planar supports for hypergraphs. J. Graph Algorithms Appl. **15**(4), 533–549 (2011)
7. Chen, J., Komusiewicz, C., Niedermeier, R., Sorge, M., Suchý, O., Weller, M.: Polynomial-time data reduction for the subset interconnection design problem. SIAM J. Discrete Math. **29**(1), 1–25 (2015)
8. Cygan, M., Fomin, F.V., Kowalik, L., Lokshtanov, D., Marx, D., Pilipczuk, M., Pilipczuk, M., Saurabh, S.: Parameterized Algorithms. Springer, Berlin (2015)
9. Diestel, R.: Graph Theory. Graduate Texts in Mathematics, vol. 173, 4th edn. Springer, Berlin (2010)

---

[3] Each planar graph has an ordering of the vertices such that each vertex has at most five neighbors later in the ordering. To achieve $n^{O(n)}$ enumeration time we simply guess such an ordering and then for each vertex its at most five later neighbors.

10. Dinkla, K., van Kreveld, M.J., Speckmann, B., Westenberg, M.A.: Kelp diagrams: point set membership visualization. Comput. Graph. Forum **31**(3), 875–884 (2012)
11. Eschbach, T., Günther, W., Becker, B.: Orthogonal hypergraph drawing for improved visibility. J. Graph Algorithms Appl. **10**(2), 141–157 (2006)
12. Flower, J., Fish, A., Howse, J.: Euler diagram generation. J. Vis. Lang. Comput. **19**(6), 675–694 (2008)
13. Habib, M., Paul, C., Viennot, L.: Partition refinement techniques: an interesting algorithmic tool kit. Int. J. Found. Comput. Sci. **10**(2), 147–170 (1999)
14. Johnson, D.S., Pollak, H.O.: Hypergraph planarity and the complexity of drawing Venn diagrams. J. Graph Theory **11**(3), 309–325 (1987)
15. Kaufmann, M., Kreveld, M., Speckmann, B.: Subdivision drawings of hypergraphs. In: Tollis, I.G., Patrignani, M. (eds.) GD 2008. LNCS, vol. 5417, pp. 396–407. Springer, Heidelberg (2009). doi:10.1007/978-3-642-00219-9_39
16. Klemz, B., Mchedlidze, T., Nöllenburg, M.: Minimum tree supports for hypergraphs and low-concurrency Euler diagrams. In: Ravi, R., Gørtz, I.L. (eds.) SWAT 2014. LNCS, vol. 8503, pp. 265–276. Springer, Heidelberg (2014). doi:10.1007/978-3-319-08404-6_23
17. Korach, E., Stern, M.: The clustering matroid and the optimal clustering tree. Math. Program. **98**(1–3), 385–414 (2003)
18. Mäkinen, E.: How to draw a hypergraph. Int. J. Comput. Math. **34**, 178–185 (1990)
19. Tarjan, R., Yannakakis, M.: Simple linear-time algorithms to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs. SIAM J. Comput. **13**(3), 566–579 (1984)
20. Tutte, W.T.: How to draw a graph. Proc. London Math. Soc. **13**(3), 743–768 (1963)
21. Verroust, A., Viaud, M.-L.: Ensuring the drawability of extended Euler diagrams for up to 8 sets. In: Blackwell, A.F., Marriott, K., Shimojima, A. (eds.) Diagrams 2004. LNCS (LNAI), vol. 2980, pp. 128–141. Springer, Heidelberg (2004). doi:10.1007/978-3-540-25931-2_13