BOLYAI SOCIETY
MATHEMATICAL STUDIES, 19

Building Bridges
pp. 1–33.

Deformable Polygon Representation and Near-Mincuts

András A. Benczúr* and Michel X. Goemans†

We derive a necessary and sufficient condition for a symmetric family of sets to have a geometric representation involving a convex polygon and some of its diagonals. We show that cuts of value less than 6/5 times the edge-connectivity of a graph admit such a representation, thereby extending the cactus representation of all mincuts.

1. Introduction

Given an undirected graph $G = (V, E)$ possibly with multiple edges (or nonnegative weights on the edges), let $d(S)$ represent the size (or the weight) of the cut $(S\colon \overline{S}) = \big\{(i,j) \in E\colon \big|\{i,j\} \cap S\big| = 1\big\}$ (where $\overline{S} = V \setminus S$). The edge-connectivity λ of G is equal to $\min_{\emptyset \neq S \neq V} d(S)$, and any cut $(S, \overline{S})$ achieving the minimum is called a mincut. A cut $(S, \overline{S})$ for which $d(S) < \alpha\lambda$ for some $\alpha > 1$ is called an α-near-mincut. In 1976, Dinitz, Karzanov and Lomonosov [4] have given a compact representation of all mincuts; this is known as the cactus representation. Informally, the cactus representation is a multigraph H in which every edge is in exactly one cycle[1] and every vertex of G is mapped to a node[2] of H (see Figure 1). This mapping does not need to be bijective, surjective or injective. A node of H can correspond

*Research supported under grant OTKA NK 72845.
†Research supported under NSF grant CCF-0515221 and ONR grant N00014-05-1-0148.

[1]In descriptions of the cactus representation, the cycles of length 2 are sometimes replaced by a single edge (bridge) of weight 2.

[2]To easily distinguish them, we use vertices for the graph G and nodes for the cactus H.

to one, several or even no vertex of G; in the latter case, the node is said to be empty. The set of all mincuts of size 2 in H, i.e. those obtained by removing any two edges of the same cycle, correspond to the set of all mincuts in G. Because of the presence of empty nodes, observe that several mincuts in H can correspond to the same cut in G.

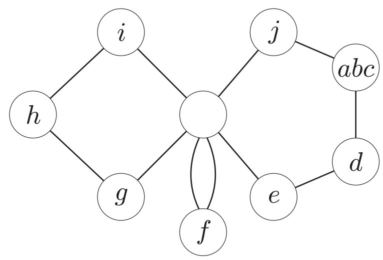

Fig. 1. A cactus of a graph with vertex set $V = \{a, b, c, d, e, f, g, h, i, j\}$.

The cactus is not necessarily unique, and Nagamochi and Kameda [15] describe two canonical cactus representations, one with no cycles of length 3 and the other without precisely 3 cycles meeting at the same empty node. Nagamochi and Kameda also show that these canonical representations (and many others) have at most $n = |V|$ empty nodes. A cactus representation can be constructed efficiently, see [11, 8, 20, 16, 3, 7, 18, 17]. A 2-level cactus representing all cuts of value λ and $\lambda + 1$ in an unweighted graph has been derived by Dinitz and Nutov [5].

In this paper, we consider extensions of the cactus representation to arbitrary symmetric[3] families $\mathcal{F} \subseteq 2^V$, in which every cycle is replaced by a convex polygon P with some of its diagonals drawn and the elements of V are mapped to the cells defined by the diagonals within the polygon, see Figure 2. A cell can have 0, 1 or many elements mapped to it. To every diagonal, one can associate a pair $(S, \overline{S})$ of complementary sets corresponding to those elements mapped to either side of the diagonal. Furthermore, we will focus on the situation when the existence of the mapping of V to the cells of the polygon does not depend on the exact (convex) location of the polygon vertices; we call such polygons *deformable,* see Section 2. Our representation links deformable polygons in a tree fashion, in an almost identical way as the cactus links cycles. To derive this tree structure, we consider the *cross graph* associated with a symmetric family of sets: Its vertex set has a representative for each pair of complementary sets in $\mathcal{F}$ and two such pairs $(S, \overline{S})$ and $(T, \overline{T})$ are joined by an edge if S and T cross.

[3] $A \in \mathcal{F}$ iff $\overline{A} \in \mathcal{F}$.

Recall that two sets S and T are said to cross if $S \cap T$, $\overline{S \cup T}$, $S \setminus T$ and $T \setminus S$ are all non-empty. We show that representations for each connected component of the cross graph can be linked in a tree structure of linear size; this is explained in Section 5, where we provide a geometric proof of a (slightly modified) cactus structure of mincuts.

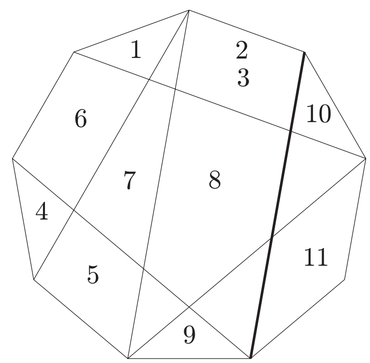

Fig. 2. A polygon representation. The bold diagonal corresponds to the sets $\{1, 2, \ldots, 9\}$ and $\{10, 11\}$. This polygon is deformable; for example, this will follow from our characterization.

The main result of this paper is to give a necessary and sufficient condition for a symmetric family to admit a representation as a tree of deformable polygons. This characterization is in terms of excluded configurations. We show that there are 3 families consisting of 4 pairs of complementary sets each ($\binom{4}{[2]}$, $\mathcal{C}_1$ and $\mathcal{C}_2$, see the forthcoming Figure 5), and the family can be represented by a tree of deformable polygons if and only if none of the three families appear as an induced subfamily. This is stated in Section 2 and proved in Sections 3 and 5. We also show in Section 4 that for any (weighted) undirected graph, the family of 6/5-near-mincuts — those cuts of value less than 6/5 times edge-connectivity λ — satisfies the condition of our main result, and therefore can be represented by a tree of deformable polygons. This, for example, implies that there are at most $\binom{n}{2}$ 6/5-near-mincuts, see Section 6; this is already known even for 4/3-near-mincuts [19, 9].

This paper focuses on a characterization of those families with a tree of deformable polygons, and does not consider efficient algorithms for its construction or implications for connectivity problems, such as speeding up algorithms for graph augmentation problems or the existence of splitting-off which maintains near-mincuts. This is covered in the Ph.D. dissertation [3] of the first author, see also [1, 2].

Many proofs of the existence of the cactus have appeared. The approach that is most useful in the context of this paper is due to Lehel, Maffray and Preissmann [12] who show using characterizations of interval hypergraphs that, for any undirected graph $G = (V, E)$, there exists a cyclic ordering of V such that every mincut corresponds to a partition of this cyclic ordering into two (cyclic) intervals. This will be the basis of our construction of deformable polygons, as we will first place a carefully selected subset of the elements along the sides of our polygon in a circular way and then add the remaining elements in cells farther inside the polygon (using Helly's theorem).

2. Deformable Polygon Representation

Given an arrangement of lines in $\mathbb{R}^2$ and a set V of points in $\mathbb{R}^2$ none of them being on any of the lines, each line partitions $\mathbb{R}^2$ into two halfplanes and hence the set V into two sets S and $\bar{S}$. We associate to this arrangement of lines the symmetric family $\mathcal{F}$ of all sets defined by these lines. We say that this family is *representable as an arrangement of lines*. For example, in Figure 3:(a), a representation of $\binom{4}{[2]} = \big\{\{1,2\},\{1,3\},\{1,4\},\{2,3\},\{2,4\},\{3,4\}\big\}$, the family of all subsets of $\{1,2,3,4\}$ of cardinality 2, is given.

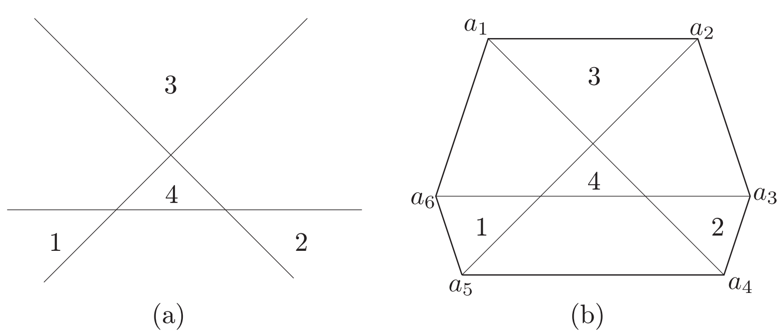

Fig. 3. (a): An arrangement of lines representing the family $\mathcal{F} = \binom{4}{[2]} = \{\{1,2\},\{1,3\},\{1,4\},\{2,3\},\{2,4\},\{3,4\}\}$. (b): An equivalent polygon representation.

A classical result of Schläfli [22] says that the maximum number of partitions of an n-element set in $\mathbb{R}^2$ by lines is $\binom{n}{2}+1$, and thus at most $\binom{n}{2}$ once we don't count the trivial partition for which $S = \emptyset$ or $S = V$. This provides a bound on the size of any family representable as an arrangement of lines.

Instead of considering representations as arrangements of lines, we consider a bounded variant of it. In this case, we have a convex polygon P with vertices $a_1, a_2, \ldots, a_k$ (for some suitable k), a subset of diagonals $D \subseteq \left\{ [a_i, a_j] \colon 1 \le i < j \le k \right\}$, and the elements of V are placed in the cells of $P \setminus D$ defined by the diagonals within the polygon. The sets represented correspond to the sets of elements on either side of a diagonal. We refer to such a representation as a *polygon representation.* Clearly, a polygon representation can be transformed into an arrangement of lines representation, and vice versa, see Figure 3 for a simple example.

For certain polygon representations, the polygon can be arbitrarily deformed in a convex manner without the cells containing elements of V vanishing. In this case, the actual positions of the polygon vertices $a_1, \cdots, a_k$ are irrelevant, provided they are in convex position. We refer to this as a *deformable polygon representation.* This is not the case for the polygon representation of $\binom{4}{[2]}$ given in Figure 3:(b), as shown in Figure 4. In fact, it is easy to show that the family $\binom{4}{[2]}$ does not admit a deformable polygon representation.

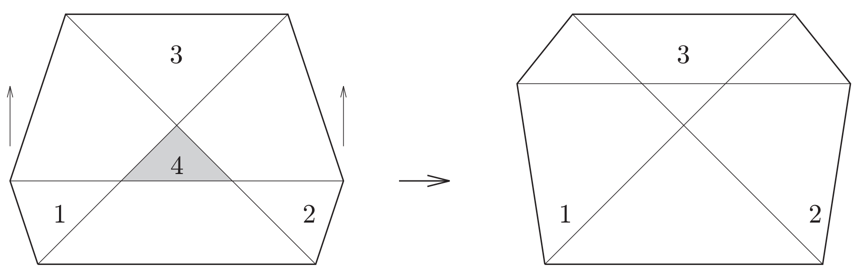

Fig. 4. If the convex polygon is deformed, the cell containing the element 4 might disappear.

In this paper, we provide a characterization in terms of excluded configurations of those symmetric families that admit a deformable polygon representation. To state our result, we first need to introduce two families that do not have a polygon representation (or a representation by an arrangement of lines). In fact, they cannot be represented by convex sets, as any such representation for a symmetric family can be transformed into a representation as an arrangement of lines by simply considering the line separating the convex sets assigned to S and $\overline{S}$.

Lemma 1. *The following two families of subsets of* $\{1,2,3,4,5,6\}$ *(see Figure 5) do not have a polygon representation.*

1. $\mathcal{C}_1 = \big\{\{1,2,3\}, \overline{\{1,2,3\}}, \{1,4\}, \overline{\{1,4\}}, \{2,5\}, \overline{\{2,5\}}, \{3,6\}, \overline{\{3,6\}}\big\}$,
2. $\mathcal{C}_2 = \big\{\{1,2,3\}, \overline{\{1,2,3\}}, \{1,5,6\}, \overline{\{1,5,6\}}, \{2,4,6\}, \overline{\{2,4,6\}}, \{3,4,5\}, \overline{\{3,4,5\}}\big\}$
$= \big\{\{1,2,3\}, \{2,3,4\}, \{3,4,5\}, \{4,5,6\}, \{5,6,1\}, \{6,1,2\}, \{1,3,5\}, \{2,4,6\}\big\}$.

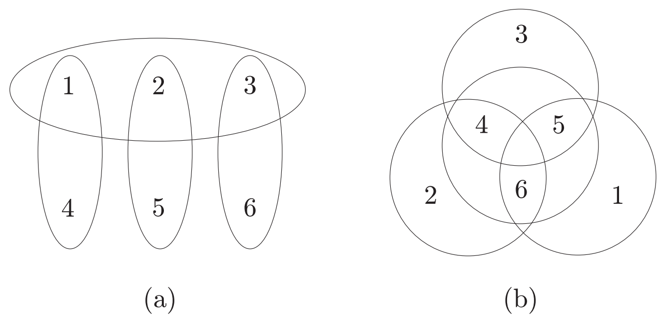

Fig. 5. Half of the sets (a) in $\mathcal{C}_1$ and (b) in $\mathcal{C}_2$. Combs, to be defined in Definition 2, have either an induced $\mathcal{C}_1$ or an induced $\mathcal{C}_2$.

Proof. Assume that $\mathcal{C}_k$ for either $k=1$ or $k=2$ has a representation by an arrangement of lines. We start with some notation for both cases. We simply let $1,2,\dots,6$ denote the points in $\mathbb{R}^2$ in the cells of the arrangement of lines corresponding to the elements of the ground set. For a pair of complementary sets $(U,\overline{U})$ in $\mathcal{C}_k$, let $\ell(U)=\ell(\overline{U})$ be the line in the arrangement for $\mathcal{C}_k$ separating U from $\overline{U}$; furthermore let $\ell^+(U)$ and $\ell^-(U)$ denote the open halfplanes containing U and $\overline{U}$, respectively. For $i \in \{1,2,3\}$ and $j \in \{4,5,6\}$, let $t_{ij} \in \ell\big(\{1,2,3\}\big)$ be the intersection point between $\ell\big(\{1,2,3\}\big)$ and the line segment extending between points i and j; t_{ij} is well-defined as i and j are separated by $\{1,2,3\} \in \mathcal{C}_k$.

Consider first the arrangement for $\mathcal{C}_1$. Consider the points t_{14}, t_{25} and t_{36}. As $t_{14} \in \ell^+\big(\{1,4\}\big)$ while $t_{25}, t_{36} \in \ell^-\big(\{1,4\}\big)$ and similarly for 14 replaced by 25 or 36, the points t_{14}, t_{25} and t_{36} are distinct. Without loss of generality, we can assume that t_{25} is between t_{14} and t_{36}. But, by convexity, $t_{14}, t_{36} \in \ell^-\big(\{2,5\}\big)$ implies that $t_{25} \in \ell^-\big(\{2,5\}\big)$ contradicting $t_{25} \in \ell^+\big(\{2,5\}\big)$.

Assume now we have a representation for $\mathcal{C}_2$. Consider two complementary sets in our family, say $U=\{1,3,5\}$ and $\overline{U}=\{2,4,6\}$. The fact that

$U \subset \ell^+(U)$ and $\overline{U} \subset \ell^-(U)$ implies that the line segments $[t_{15}, t_{35}]$ and $[t_{24}, t_{26}]$ do not intersect, and in particular that $t_{15} \neq t_{26}$. We can similarly deduce that $t_{15} \neq t_{34} \neq t_{26}$. Without loss of generality, let us assume that t_{26} is between t_{15} and t_{34}. But the disjointness of $[t_{15}, t_{35}]$ and $[t_{24}, t_{26}]$ implies that $t_{26} \notin [t_{15}, t_{35}]$, while the same argument with the complementary sets $\{1, 2, 6\}$ and $\{3, 4, 5\}$ implies that $t_{26} \notin [t_{34}, t_{35}]$, which means that $t_{26} \notin [t_{15}, t_{34}]$, a contradiction. ■

Given a family $\mathcal{F}$ of sets on V and a subset $S \subseteq V$, we define the family $\mathcal{F}|_S$ to be the family $\{F \cap S \colon F \in \mathcal{F}\}$. We say that a family $\mathcal{F}$ of V *contains* a family $\mathcal{G}$ as an *induced family* or simply contains the *induced family* $\mathcal{G}$ if there exists S such that $\mathcal{F}|_S$ contains $\mathcal{G}$, i.e. $\mathcal{F}|_S \supseteq \mathcal{G}$. In this case, if $\mathcal{G}$ does not have a polygon representation then neither does $\mathcal{F}$, and similarly for deformable polygon representations. Thus, any family $\mathcal{F}$ which contains $\binom{4}{[2]}$, $\mathcal{C}_1$ or $\mathcal{C}_2$ as an induced family does not have a deformable polygon representation. We will show that the converse to that statement is also true for families with connected cross graph.

Theorem 2. *Let $\mathcal{F}$ be a symmetric family of sets with connected cross graph. Then $\mathcal{F}$ admits a deformable polygon representation if and only if $\mathcal{F}$ does not contain $\binom{4}{[2]}$, $\mathcal{C}_1$ or $\mathcal{C}_2$ as an induced family.*

The "only if" part follows from Lemma 1 and the fact that $\binom{4}{[2]}$ does not admit a deformable polygon representation. The proof of the "if" part is constructive, and will be the focus of the next section. Here is a brief sketch of the construction proving the existence of the polygon representation. First we identify a set $O \subseteq V$ of elements that we would like to place along the sides of the polygon; we refer to the elements of O as *outside* elements. One non-trivial property of this set O is that every set $S \in \mathcal{F}$ contains at least one outside element but not all of them; this is indeed a property to expect if we place these outside elements along the sides of the polygon and represent sets by diagonals. We then use Tucker's characterization [23] of interval hypergraphs to show that there exists a circular ordering of the outside elements such that any set in $\mathcal{F}|_O$ corresponds to an interval in that circular ordering. As a first trial, one could take a k-gon P where $k = |O|$, place the outside elements in their circular order along the sides of P, and for each set $S \in \mathcal{F}$ add the diagonal that separates $S \cap O$ from $O \setminus S$. This already provides a polygon representation of $\mathcal{F}|_O$. However, placing the remaining elements of $\overline{O}$ appropriately inside the polygon in order to represent $\mathcal{F}$ is not always possible for the following reason. Several sets of $\mathcal{F}$ could have the same intersection with O and thus would be mapped to the same diagonal. We can, however, show that sets having the same intersection with O form a chain, and we can then add vertices to the polygon P so as to make the diagonals distinct (and non-crossing), and still represent $\mathcal{F}_O$. We then need to consider the placement of the remaining *inside* elements, those of $\overline{O}$. For each inside element v, we know on which side of each diagonal we would like to place it. We show that the intersection of all these halfplanes is non-empty by proving that any two or three of them have a non-empty intersection and then using Helly's theorem. This non-empty intersection gives a non-empty cell where to place v, and this completes the construction and the proof.

In Section 5, we show that the connected components of any cross graph can be arranged in a tree structure, and when each component can be represented by a deformable polygon, we obtain a representation that we call a *tree of deformable polygons.* If a family contains $\binom{4}{[2]}$, $\mathcal{C}_1$ or $\mathcal{C}_2$ as an induced family then one of the connected components of its cross graph must contain one of these subfamilies as an induced family. Therefore, after proving Theorem 2 and deriving our tree structure, we will have shown the following theorem.

Theorem 3. *Let $\mathcal{F}$ be a symmetric family of sets. Then $\mathcal{F}$ admits a representation as a tree of deformable polygons if and only if $\mathcal{F}$ does not contain $\binom{4}{[2]}$, $\mathcal{C}_1$ or $\mathcal{C}_2$ as an induced family.*

We would like to point out that we do not know of any necessary and sufficient condition for the existence of a (not necessarily deformable) polygon representation (or a representation by an arrangement of lines).

3. Construction and its Proof

In this Section, we focus on the case where the symmetric family has a connected cross graph, or if this is not the case, we redefine $\mathcal{F}$ to be the family corresponding to a single connected component of the cross graph.

Before we start, observe that we can group together any pair or set of elements which are not separated by any set in our family. Indeed, this does not affect the existence of a polygon representation as all these equivalent elements would fall in the same cell defined by the diagonals of the polygon. More formally, define two elements u and v to be equivalent if, for every set S_i in our family, $u \in S_i$ iff $v \in S_i$. Let the equivalence classes be called *atoms*. Our representation is on the atoms of our family. For simplicity, for the rest of this section, we simply refer to them as the elements; we'll use atoms when we consider several connected components of the cross graph in Section 5.

We will first restate Theorem 2 without using any induced families. For this, we need a few definitions.

Definition 1. 3 subsets C_1, C_2 and C_3 form a 3-cycle (see figure 6) if (i) $\overline{C_1 \cup C_2 \cup C_3} \neq \emptyset$ and (ii) $(C_i \cap C_{i+1}) \setminus C_{i-1} \neq \emptyset$, for all $i \in \{1, 2, 3\}$.

Everywhere in the paper, indices should always be considered cyclic; for example, C_4 represents C_1 in the above definition.

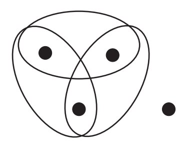

Fig. 6. A 3-cycle. Solid dots in a Venn diagram denote non-empty intersections; other intersections could be empty or not.

Observe that a symmetric family $\mathcal{F}$ contains 3 sets that form a 3-cycle if and only if it contains an *induced* $\binom{4}{[2]}$. We also need to define another configuration of sets.

Definition 2. Four subsets H, T_1, T_2 and T_3 form a comb with handle H and teeth T_1, T_2 and T_3 if (i) either $H \cap \big(T_i \setminus (T_{i-1} \cup T_{i+1})\big) \neq \emptyset$ for all $i = 1, 2, 3$ or $H \cap (T_{i-1} \cap T_{i+1} \setminus T_i) \neq \emptyset$ for all $i = 1, 2, 3$ and (ii) the same holds for H replaced by $\bar{H}$.

The above definition is such that any family of sets which contains (sets forming) a comb must contain either an induced $\mathcal{C}_1$ or an induced $\mathcal{C}_2$, and vice versa, see Figure 5; in both cases, the handle H gets restricted to $\{1, 2, 3\}$ or $\{4, 5, 6\}$. Thus, alternatively, we could reformulate Theorem 2 as follows.

Theorem 4. *Let $\mathcal{F}$ be a family of sets with connected cross graph. Then $\mathcal{F}$ admits a deformable polygon representation if and only if $\mathcal{F}$ does not contain any 3-cycle or comb.*

More generally, we define a k-cycle for $k > 3$ as follows.

Definition 3. k subsets $C_1, \ldots, C_k$ for $k > 3$ form a k-cycle (or simply cycle) if (i) $\overline{\cup_{i=1}^{k} C_i} \neq \emptyset$, (ii) $P_i = (C_i \cap C_{i+1}) \neq \emptyset$ for $i = 1, \ldots, k$, and (iii) C_i and C_j are disjoint for $j \notin \{i-1, i, i+1\}$. See Figure 7.

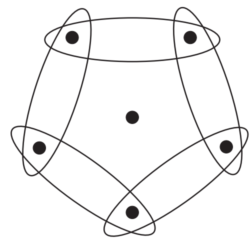

Fig. 7. A 5-cycle.

Because of condition (iii) in the definition of a k-cycle for $k > 3$ (which we didn't have for a 3-cycle), a family could contain an induced k-cycle but no k-cycle itself. However, the next proposition shows that when there are no 3-cycles and combs, we do not need to differentiate between cycles and induced cycles.

Proposition 5. *Consider a collection $\mathcal{F}$ of sets that does not contain any 3-cycle or comb. Then any induced cycle $C_1, \ldots, C_k$ contains a subcollection which forms a cycle.*

The proof is technical and can be skipped at first reading.

Proof. We assume that the collection is minimal, i.e. no subcollection of the C_i's forms an induced cycle. Let W be such that the C_i's induce a cycle in W. Define

$$P_i = W \cap \big[(C_i \cap C_{i+1})\big] = W \cap \Bigg[(C_i \cap C_{i+1}) \setminus \bigcup_{j \notin \{i, i+1\}} C_j\Bigg] \neq \emptyset$$

for $i = 1, \ldots, k$, the last equality following from the fact that $C_i \cap W$ and $C_j \cap W$ are disjoint if i and j are neither consecutive nor equal.

Assume there exist two indices i and $j \notin \{i-1, i, i+1\}$ such that $\big(C_i \setminus (C_{i+1} \cup C_{i-1})\big) \cap C_j \neq \emptyset$, i.e. $(C_i \cap C_j) \setminus (C_{i+1} \cup C_{i-1}) \neq \emptyset$. If j is $i-2$, we have a 3-cycle consisting of C_j, C_{i-1} and C_i, a contradiction. Similarly, if j is $i+2$, we have the 3-cycle (C_j, C_i, C_{i+1}). Thus we can assume $j \notin \{i-2, i-1, i, i+1, i+2\}$. We claim we have a comb with handle C_i and teeth C_{i-1}, C_{i+1} and C_j, a contradiction. This is because the following sets are all non-empty:

- $\big(C_j \setminus (C_{i+1} \cup C_{i-1})\big) \cap C_i = (C_i \cap C_j) \setminus (C_{i+1} \cup C_{i-1}) \neq \emptyset$,
- $\big(C_{i+1} \setminus (C_j \cup C_{i-1})\big) \cap C_i \supseteq P_i$ as $j \notin \{i, i+1\}$,
- $\big(C_{i-1} \setminus (C_j \cup C_{i+1})\big) \cap C_i \supseteq P_{i-1}$ as $j \notin \{i-1, i\}$,
- $\big(C_j \setminus (C_{i+1} \cup C_{i-1})\big) \cap \overline{C_i} \supseteq P_j$ as $j \notin \{i-2, i-1, i, i+1\}$,
- $\big(C_{i+1} \setminus (C_j \cup C_{i-1})\big) \cap \overline{C_i} \supseteq P_{i+1}$ as $j \notin \{i+1, i+2\}$,
- $\big(C_{i-1} \setminus (C_j \cup C_{i+1})\big) \cap \overline{C_i} \supseteq P_{i-2}$ as $j \notin \{i-2, i-1\}$.

We can thus assume that for all i and all $j \notin \{i-1, i, i+1\}$, we have

$$C_i \cap C_j \subseteq (C_{i-1} \cup C_{i+1}). \tag{1}$$

Observe that the C_i's form a cycle unless condition (iii) of Definition 3 is violated, i.e. there exists i and $j \notin \{i-1, i, i+1\}$ with $C_i \cap C_j \neq \emptyset$. Combining this with (1), we can assume that there exist i and j with $j \notin \{i-1, i, i+1\}$ such that $\emptyset \neq C_i \cap C_j \subseteq (C_{i-1} \cup C_{i+1})$. This means that either $C_i \cap C_{i+1} \cap C_j \neq \emptyset$ or $C_{i-1} \cap C_i \cap C_j \neq \emptyset$. Depending on the value of j and after possibly changing i, (i) we either have i and $j \notin \{i-2, i-1, i, i+1\}$ with $C_{i-1} \cap C_i \cap C_j \neq \emptyset$, or (ii) we have i with $C_{i-2} \cap C_{i-1} \cap C_i \neq \emptyset$. We consider both cases separately.

1. Assume that $C_{i-1} \cap C_i \cap C_j \neq \emptyset$ for some i and j with $j \notin \{i-2, i-1, i, i+1\}$. We claim that C_{i-1}, C_i and $\overline{C_j}$ form a 3-cycle, again a contradiction. Indeed $C_{i-1} \cap C_i \cap C_j \neq \emptyset$, $C_{i-1} \setminus (C_i \cup C_j) \supseteq P_{i-2} \neq \emptyset$ as $j \notin \{i-2, i-1\}$, $C_i \setminus (C_{i-1} \cup C_j) \supseteq P_i \neq \emptyset$ as $j \notin \{i, i+1\}$, and $C_j \setminus (C_{i-1} \cup C_i) \supseteq P_{j-1} \neq \emptyset$ as $j \notin \{i-1, i, i+1\}$.

2. Assume that $C_{i-2} \cap C_{i-1} \cap C_i \neq \emptyset$ for some i. If $C_{i-1} \setminus (C_{i-2} \cup C_i) \neq \emptyset$, we have a 3-cycle $(C_{i-2}, \overline{C_{i-1}}, C_i)$. Thus assume that $C_{i-1} \subseteq C_{i-2} \cup C_i$. Furthermore, we can assume that $C_{i-2} \cap C_i = C_{i-2} \cap C_{i-1} \cap C_i$ since, otherwise, C_{i-2}, C_{i-1} and C_i would form a 3-cycle. By our minimality assumption, we cannot remove C_{i-1} from our collection and still have an induced cycle; this implies that $C_{i-2} \cap C_{i-1} \cap C_i = C_{i-2} \cap C_i \subseteq \big[\bigcup_{j \notin \{i-2, i-1, i\}} C_j\big]$. Thus, let $l \notin \{i-2, i-1, i\}$ be such that $C_l \cap C_{i-2} \cap C_{i-1} \cap C_i \neq \emptyset$. If $k > 4$ then C_{i-2}, C_i and $\overline{C_l}$ can be seen to form a 3-cycle (since we have $(C_{i-2} \cap C_i) \setminus \overline{C_l} = C_{i-2} \cap C_i \cap C_l \neq \emptyset$, $(C_{i-2} \cap \overline{C_l}) \setminus C_i \supseteq P_{i-2}$, $(C_i \cap \overline{C_l}) \setminus C_{i-2} \supseteq P_{i-1}$, and $\overline{C_i \cup \overline{C_l} \cup C_{i-2}}$ contains either P_l or P_{l-1}). On the other hand, if $k = 4$ and thus $l = i+1$ then $\overline{C_{i-2}}$, C_{i-1}, $\overline{C_i}$ and C_{i+1} form a comb with C_{i-1} as handle. Indeed, we have

- $\big(\overline{C_{i-2}} \setminus (\overline{C_i} \cup C_{i+1})\big) \cap C_{i-1} \supseteq P_{i-1}$,
- $\big(\overline{C_i} \setminus (\overline{C_{i-2}} \cup C_{i+1})\big) \cap C_{i-1} \supseteq P_{i-2}$,
- $\big(C_{i+1} \setminus (\overline{C_{i-2}} \cup \overline{C_i})\big) \cap C_{i-1} = C_{i-2} \cap C_{i-1} \cap C_i \cap C_{i+1} \neq \emptyset$,
- $\big((\overline{C_{i-2}} \cap \overline{C_i}) \setminus C_{i+1}\big) \cap \overline{C_{i-1}} = \overline{C_{i-2} \cup C_{i-1} \cup C_i \cup C_{i+1}} \neq \emptyset$,
- $\big((\overline{C_i} \cap C_{i+1}) \setminus \overline{C_{i-2}}\big) \cap \overline{C_{i-1}} \supseteq P_{i+1}$,
- $\big((\overline{C_{i-2}} \cap C_{i+1}) \setminus \overline{C_i}\big) \cap \overline{C_{i-1}} \supseteq P_i$. ■

Definition 4. An element $v \in V$ of the family $\mathcal{F} \subset 2^V$ is said to be *inside* if there exists a cycle $C_1, \ldots, C_k$ of $\mathcal{F}$ such that $v \notin \cup_i C_i$. Otherwise, v is said to be outside. The set of outside elements is denoted by O.

A few remarks are in order. Given Proposition 5, we can replace "cycle" by "induced cycle" in the above definition, provided that $\mathcal{F}$ does not contain any comb or 3-cycle. For the rest of this section, we assume throughout that $\mathcal{F}$ has no 3-cycle or comb, even if it is not explicitly stated. Also, at this point, it is not obvious that O is non-empty; this will follow from Proposition 10. This would not be true if our family could have 3-cycles; $\binom{4}{[2]}$ for example has no outside elements. In fact, we will deduce from Corollary 14 that either $|O| \geq 4$ or our family with its connected cross graph consists of only one pair of complementary sets $(S, \overline{S})$ (thus separating two outside elements/atoms from each other). The latter case trivially gives rise to a deformable polygon; just take a 4-gon with one diagonal. Therefore, we will often implicitly assume in this Section that our family consists of more than one complementary pair. Observe also that if our family does not contain any cycles (as is the case for the family of mincuts in a graph, see Section 4) then all elements are outside, i.e. $O = V$.

As a first step towards the contruction of the polygon representation, we show now that the family restricted to the outside elements, $\mathcal{F}|_O$, is a *circular representable hypergraph* [21] or a *circular arc hypergraph*, i.e. there exists a circular ordering of O such that all sets in $\mathcal{F}|_O$ correspond to arcs of the circle. This is similar to the proof of the existence of the cactus representation of all minimum cuts due to Lehel et al. [12].

Proposition 6. *Consider a symmetric family $\mathcal{F}$ of sets with no cycles or combs. Then $\mathcal{F}$ is a circular representable hypergraph.*

By definition of outside elements, we can then derive the following Corollary.

Corollary 7. *Consider a symmetric family $\mathcal{F}$ of sets that does not contain any 3-cycle or comb, and let O be its set of outside elements. Then $\mathcal{F}|_O$ is a circular representable hypergraph.*

Proposition 6 and Corollary 7 follow easily from Tucker's characterization [23] of *interval hypergraphs*, i.e. hypergraphs (or families of sets) for which there exists a total ordering of the elements of the ground set such that every hyperedge (set) corresponds to an interval in the ordering. Tucker gives a necessary and sufficient condition for a 0-1 matrix to have the consecutive 1's property, see Duchet [6] for a short proof of Tucker's result in terms of hypergraphs.

Theorem 8 (Tucker [23])**.** *A family of sets define an interval hypergraph if and only if it does not contain any of the families listed in Figure 8 as an induced subfamily.*

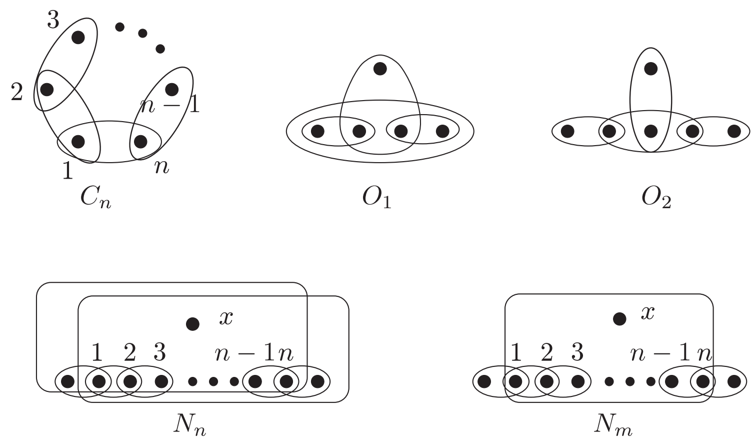

Fig. 8. List of excluded subhypergraphs for interval hypergraphs: C_n for $n \geq 3$, O_1, O_2, N_n for $n \geq 1$ and M_n for $n \geq 1$. Solid dots represent non-empty intersections.

Proof of Proposition 6. Select an element v_0 arbitrarily and consider the family $\mathcal{C} = \{S \in \mathcal{F} : v_0 \notin S\}$. We need to show that $\mathcal{C}$ is an interval hypergraph. By Tucker's Theorem, we only need to show that $\mathcal{C}$ does not contain any of the subfamilies of Figure 8. Observe that C_n is an induced cycle not containing v_0, and therefore is not present by assumption. Similarly, an induced cycle can be obtained from M_n and N_n by complementing the large sets; these cycles do not contain the element marked X in Figure 8. O_2 is a comb, and a comb can be obtained from O_1 by complementing the 4-element set which would then contain the special element v_0. As we have no comb in our family, O_1 and O_2 cannot arise. Therefore, $\mathcal{C}$ is an interval hypergraph, and $\mathcal{F}$ is a circular representable hypergraph. ■

To construct the polygon representation, we start with a convex k-gon with vertices $a_1, a_2, \ldots, a_k$ (clockwise) where $k = |O|$, and place the outside elements along the k sides in the circular order given by Proposition 7 so that any diagonal of the k-gon partitions the outside elements into 2 intervals in the circular ordering. From Proposition 10 stated and proved below, we can derive that any set $S \in \mathcal{F}$ partitions O non-trivially, i.e. $S \cap O \neq \emptyset$ and $O \setminus S \neq \emptyset$. This means that we can associate with S one of the diagonals of our polygon, corresponding to the way it partitions the outside elements. Since we were assuming that any two elements (including consecutive elements of O in the circular ordering) were separated by a set $S \in \mathcal{F}$ (by our definition of atoms at the beginning of this Section), we have at least one diagonal incident to every vertex of our k-gon.

The trouble though is that several pairs of complementary sets might be assigned to the same diagonal. As an example, consider $\mathcal{F}_4$ consisting of a 4-cycle together with the complementary sets:

$$\mathcal{F} = \big\{ \{1,2\}, \{2,3\}, \{3,4\}, \{1,4\}, \{3,4,5\}, \{1,4,5\}, \{1,2,5\}, \{2,3,5\} \big\}.$$

Element 5 is inside, while 1,2,3 and 4 are outside; however $\{1,2\}$, $\{3,4,5\}$ and $\{1,2,5\}$, $\{3,4\}$ separate $O = \{1,2,3,4\}$ in the same way, and there is no space between the corresponding diagonals to place element 5. We will prove in Proposition 15 that all sets S having the same intersection with O, say $S \cap O = A$, form a chain $S_1 \supseteq S_2 \supseteq \cdots \supseteq S_p$. So far, all these sets correspond to the same diagonal, say from a_u to a_v. We modify the polygon by replacing a_u by p vertices, say $a_u^1, \ldots, a_u^p$ clockwise, by replacing a_v by p vertices, say $a_v^1, \ldots, a_v^p$ anticlockwise and by assigning S_i to the diagonal $[a_u^i, a_v^i]$ for $i = 1, \ldots, p$, see Figure 9. The sets $S_1, \ldots, S_p$ are now assigned to p non-crossing diagonals (i.e. which do not intersect in the interior of the polygon). Other diagonals incident to a_u, like $[a_u, a_w]$, are moved to be incident to either a_u^1 if w is between v and u (clockwise) or to a_u^p if w is between u and v (clockwise), see Figure 9. The same process is repeated for every diagonal which corresponds to more than one set. Overall, this creates a polygon with $|O| + |\mathcal{F}| - \big|\mathcal{F}|_O\big|$ vertices. To reduce this number of vertices, we could have replaced only a_u (or a_v) by p vertices instead of replacing both; the proofs also carry through in that case.

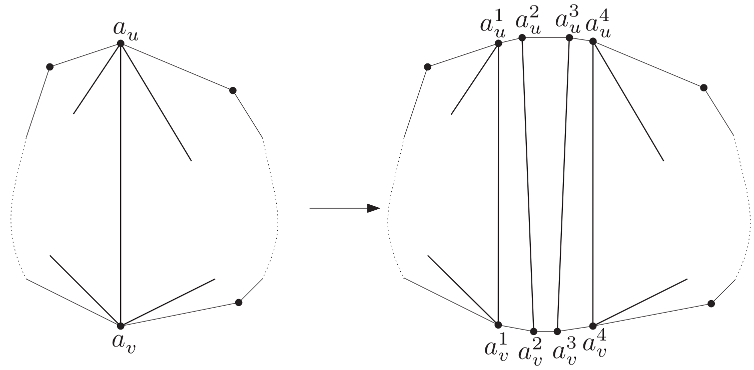

Fig. 9. For sets S having the same intersection with O, we expand the polygon and introduce non-crossing diagonals.

In the example of $\mathcal{F}_4$ (4-cycle plus the complementary sets), we first take a (convex) quadrilateral as there are 4 outside elements and then add 4 additional vertices to duplicate the two diagonals, see Figure 10.

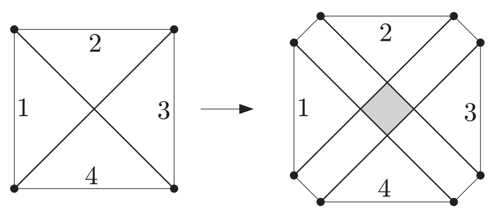

Fig. 10. Construction of the polygon and placement of the outside elements for the symmetrized 4-cycle; observe that there is a shaded cell to correctly place element 5.

At this point, we have constructed a convex q-gon $a_1, \ldots, a_q$, placed all the outside elements, and assigned every set $S \in \mathcal{F}$ to one of its diagonal, say $\ell(S) = \left[a_{i(S)}, a_{i(\overline{S})}\right]$ for some indices $i(S), i(\overline{S})$. Observe that this polygon has the following important property (in addition to representing $\mathcal{F}|_O$):

Proposition 9. *Let $S_1, S_2 \in \mathcal{F}$. Then the corresponding diagonals $\ell(S_1)$ and $\ell(S_2)$ do not cross (i.e. do not intersect in the interior of P) if and only if the sets $S_1 \cap O$ and $S_2 \cap O$ do not cross in O (i.e. $S_1 \cap O \subseteq S_2 \cap O$, $S_2 \cap O \subseteq S_1 \cap O$, $S_1 \cap S_2 \cap O = \emptyset$, or $S_1 \cup S_2 \supseteq O$).*

Indeed, this is true for our initial k-gon with $k = |O|$, and remains true as we introduce additional vertices and diagonals.

We now need to prove (i) every $S \in \mathcal{F}$ partitions O non-trivially and (ii) that sets that have the same intersection with O form a chain, and then prove that (iii) our (deformable) polygon representation can be correctly completed, i.e. that inside elements can be placed appropriately within cells of the arrangement of diagonals of our polygon and this should be true independently of the position of the vertices $a_1, \ldots, a_q$. We start with a stronger statement than (i) which will be also useful for (ii) and (iii).

Proposition 10. *Consider a family $\mathcal{F}$ of sets with a connected cross graph that does not contain any 3-cycle or comb. Let S be a minimal set in $\mathcal{F}$. Then $S \subseteq O$.*

This is the first time we require that the family has a connected cross graph. This proposition implies that any (not necessarily minimal) set $S \in \mathcal{F}$ must contain outside elements since all elements in any minimal subset of S will be outside. Applying the same proposition to $\bar{S}$, we see that every set $S \in \mathcal{F}$ partitions O non trivially.

As a first (and main) step in the proof, we show the following lemma.

Lemma 11. *Consider a family $\mathcal{F}$ that does not contain any 3-cycle or comb. Let $S \in \mathcal{F}$ contain an inside element v, and let $C_1, \ldots, C_k$ be a cycle for v (i.e. $v \notin (\cup C_i)$). Then either*

1. *there exists i such that $C_i \subset S$, or*
2. *S is disjoint from $\cup_i C_i$.*

Observe that we did not impose that the cross graph was connected. In fact, in Lemma 12, we will show that 2. can only happen if S and the cycle belong to different connected components of the cross graph.

Proof. Let us assume that $C_i \not\subset S$ for all $i = 1, \ldots, k$. We proceed in several steps.

Claim 1. For all i, $(C_i \cap C_{i+1}) \setminus S \neq \emptyset$. If not (see Figure 11, (a)), there exists i such that $C_i \cap C_{i+1} \subset S$ and we claim that C_i, C_{i+1} and $\bar{S}$ form a 3-cycle. Indeed $C_i \cap C_{i+1} \cap S = C_i \cap C_{i+1} \neq \emptyset$, $v \in S \setminus (C_i \cup C_{i+1})$, $C_i \setminus (C_{i+1} \cup S) = C_i \setminus S \neq \emptyset$ by our assumption, and similarly $C_{i+1} \setminus (C_i \cup S) \neq \emptyset$.

Claim 2. For all i, $\big(S \setminus (C_{i-1} \cup C_{i+1})\big) \cap C_i = \emptyset$. If not (see Figure 11, (b)), we have a comb with handle C_i and teeth C_{i-1}, C_{i+1} and S. Indeed, $\big[C_{i\mp1} \setminus (C_{i\pm1} \cup S)\big] \cap C_i = [C_{i\mp1} \setminus S] \cap C_i \neq \emptyset$ by claim 1, $\big[S \setminus (C_{i-1} \cup C_{i+1})\big] \cap C_i \neq \emptyset$ by assumption, $\big[C_{i\mp1} \setminus (C_{i\pm1} \cup S)\big] \cap \overline{C_i} = [C_{i\mp1} \setminus S] \cap \overline{C_i} \supseteq (C_{i\mp1} \cap C_{i\mp2}) \setminus S \neq \emptyset$ by claim 1, and $v \in \big[S \setminus (C_{i-1} \cup C_{i+1})\big] \cap \overline{C_i}$.

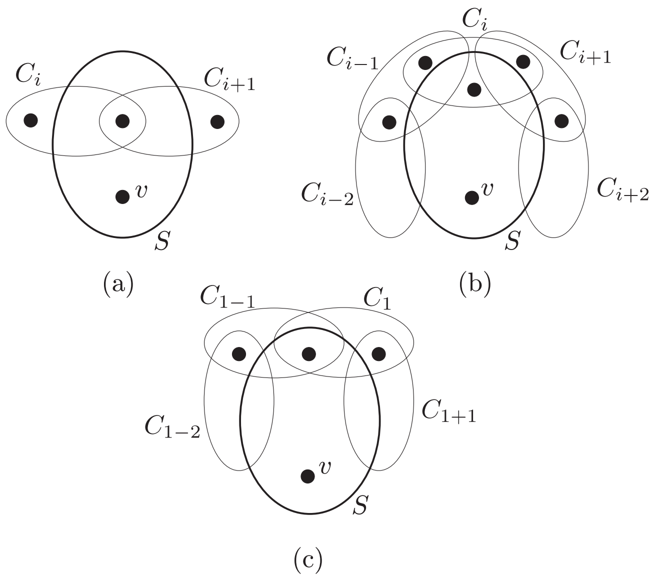

Fig. 11. Cases in the proof of Lemma 11.

Claim 3. For all i, S is disjoint from C_i. If not, there exists i with $S \cap C_i \neq \emptyset$. By Claim 2, $\emptyset \neq S \cap C_i = S \cap C_i \cap (C_{i-1} \cup C_{i+1}) = (S \cap C_i \cap C_{i-1}) \cup (S \cap C_i \cap C_{i+1})$. Thus, $S \cap C_{l-1} \cap C_l \neq \emptyset$ for $l = i$ or $l = i+1$. This implies that C_{l-1}, C_l and $\bar{S}$ form a 3-cycle (see Figure 11, (c)) since $v \in S \backslash (C_{l-1} \cup C_l)$, $C_l \backslash (C_{l-1} \cup S) \supseteq (C_l \cap C_{l+1}) \backslash S \neq \emptyset$ by claim 1 and $C_{l-1} \setminus (C_l \cup S) \supseteq (C_{l-1} \cap C_{l-2}) \setminus S \neq \emptyset$ also by claim 1.

This completes the proof of the lemma. ■

Lemma 12. *Consider a family $\mathcal{F}$ of sets with a connected cross graph that does not contain any 3-cycle or comb. Let $S \in \mathcal{F}$ contain an inside element v, and let $C_1, \ldots, C_k$ be a cycle for v. Then there exists i such that $C_i \subset S$.*

Proof. From Lemma 11, assume that S is disjoint from all C_i's. Since the cross graph is connected, there exists a path in the cross graph from S to one of the C_i's. Take a shortest path from S to one of the C_i's and consider the last two sets P and Q on it (P might be S), see Figure 12. We therefore have P disjoint from all C_i's, Q crossing one of them, and P and Q crossing. By Lemma 11 applied to Q and $\bar{Q}$, we derive that $Q \supset C_s$ for some s and that Q is disjoint from C_t for some t. Therefore, we can find two non-consecutive (and hence disjoint) sets C_i and C_j which both cross Q. However, this is a contradiction since we now have three sets, C_i, C_j and P, all disjoint and all crossing Q, which therefore form a comb. ∎

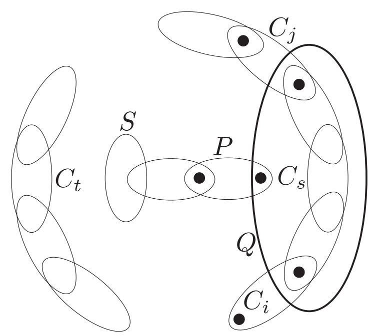

Fig. 12. Setting in the proof of Lemma 12.

Proposition 10 now follows straightforwardly from Lemma 12. We now need to consider the intersection of two sets in our family; this will be useful both for showing (ii) that sets having the same intersection with O form a chain, and (iii) that inside elements can be placed appropriately in the interior of the polygon. Throughout the rest of this section, we assume that $\mathcal{F} \subset 2^V$ is a symmetric family of sets with a connected cross graph that does not contain any 3-cycle or comb. For brevity, the assumption will not be stated in every statement.

Proposition 13. *Let S_1, S_2 be two sets in $\mathcal{F}$ with $S_1 \cap S_2 \neq \emptyset$ and minimal in the following sense: there are no $S_3, S_4 \in \mathcal{F}$ with $S_3 \subseteq S_1$, $S_4 \subseteq S_2$ and $\emptyset \neq S_3 \cap S_4 \neq S_1 \cap S_2$. Then either $S_1 \cup S_2 = V$ or $S_1 \cap S_2 \subseteq O$, i.e. $S_1 \cap S_2$ only contains outside elements.*

Proof. Assume that $S_1 \cap S_2$ contains an inside element v, corresponding to a cycle $C_1, \ldots, C_k$. We would like to show that $S_1 \cup S_2 = V$. By Lemma 12 applied to S_1 and S_2, we know the existence of $C_s \subset S_1$ and $C_t \subset S_2$. We claim that $C_s \subseteq S_1 \setminus S_2$; if not, replacing S_1 by C_s would contradict the minimality of S_1, S_2 as $C_s \cap S_2 \subsetneq S_1 \cap S_2$. Similarly, $C_t \subseteq S_2 \setminus S_1$.

Let p and q be such that $C_{p+1}, \ldots, C_s, \ldots, C_{q-1} \subseteq S_1 \setminus S_2$, but $C_p \nsubseteq S_1 \setminus S_2$ and $C_q \nsubseteq S_1 \setminus S_2$. Since p and q cannot be consecutive (because of the existence of t), we have that C_p and C_q are disjoint. Moreover, because of minimality, we also have that $C_p \not\subset S_1$ and $C_q \not\subset S_1$. The fact that C_p crosses $C_{p+1} \subseteq S_1 \setminus S_2$ implies that (a) $C_p \cap (S_1 \setminus S_2) \neq \emptyset$ and (b)

$S_1 \setminus (S_2 \cup C_p) \neq \emptyset$, and the same holds for C_p replaced with C_q (and C_{p+1} replaced with C_{q-1}).

We consider two cases.

1. Assume first that $C_p \setminus (S_1 \cup S_2) \neq \emptyset$. If $C_q \cap (S_2 \setminus S_1) \neq \emptyset$ then C_q, S_1 and S_2 would form a 3-cycle, a contradiction: $(S_1 \cap C_q) \setminus S_2 \neq \emptyset$ by (a) for C_q, $(S_2 \cap C_q) \setminus S_1 \neq \emptyset$ by assumption, $(S_1 \cap S_2) \setminus C_q \neq \emptyset$ as it contains v, and $\overline{S_1 \cup S_2 \cup C_q} \supseteq C_p \setminus (S_1 \cup S_2) \neq \emptyset$ since we assumed it.

   Thus we can assume that $C_q \cap (S_2 \setminus S_1) = \emptyset$, which implies that $C_q \setminus (S_1 \cup S_2) = C_q \setminus S_1 \neq \emptyset$. We now claim that the teeth C_p, C_q and S_2 together with the handle S_1 form a comb, a contradiction. Indeed, the following six sets are non-empty:

   - $\big(C_p \setminus (C_q \cup S_2)\big) \cap S_1 = C_p \cap (S_1 \setminus S_2) \neq \emptyset$ by (a),
   - similarly for $\big(C_q \setminus (C_p \cup S_2)\big) \cap S_1$,
   - $\big(S_2 \setminus (C_p \cup C_q)\big) \cap S_1$ as it contains v,
   - $\big(C_p \setminus (C_q \cup S_2)\big) \cap \overline{S_1} = C_p \setminus (S_1 \cup S_2) \neq \emptyset$ by our assumption,
   - $\big(C_q \setminus (C_p \cup S_2)\big) \cap \overline{S_1} = C_q \setminus (S_1 \cup S_2) \neq \emptyset$ as we have derived,
   - $\big(S_2 \setminus (C_p \cup C_q)\big) \cap \overline{S_1} = S_2 \setminus (S_1 \cup C_p \cup C_q) = S_2 \setminus (S_1 \cup C_p) \supseteq C_t \setminus C_p \neq \emptyset$ (the second equality following from $C_q \cap (S_2 \setminus S_1) = \emptyset$).

2. We can therefore assume that $C_p \subseteq (S_1 \cup S_2)$ and $C_q \subseteq (S_1 \cup S_2)$. Now, S_1, S_2 and C_p (or C_q) form a 3-cycle ($v \in S_1 \cap S_2 \setminus C_p$, $C_p \cap S_1 \setminus S_2 \neq \emptyset$ by (a), $C_p \cap S_2 \setminus S_1 = C_p \setminus S_1 \neq \emptyset$) unless $S_1 \cup S_2 = V$, proving the result. ■

Corollary 14. *Let S_1, S_2 be two sets in $\mathcal{F}$ with $S_1 \cap S_2 \neq \emptyset$. Then either $S_1 \cup S_2 = V$ or $S_1 \cap S_2 \cap O \neq \emptyset$.*

The corollary follows from Proposition 13 by considering a minimal pair (S_3, S_4) within (S_1, S_2).

As a side remark, Corollary 14 implies that if there exists an inside element v then we must have at least 4 outside elements; a cycle $C_1, \dots, C_k$ for v (where $k \geq 4$) shows the existence of outside elements in each $C_i \cap C_{i+1}$. On the other hand, if we have no inside element and fewer than 4 outside elements then we must have $|V| = |O| = 2$ as a family on 3 elements could not have a connected cross graph.

We can also deduce from Corollary 14 that sets with the same intersection with O form a chain.

Proposition 15. *Let $A \subseteq O$, and let $\mathcal{C} = \{S \in \mathcal{F} \colon S \cap O = A\}$. Then $\mathcal{C}$ is a chain, i.e. the members of $\mathcal{C}$ can be ordered such that $S_1 \subset S_2 \subset \cdots \subset S_p$.*

Proof. Consider any two members S and T in $\mathcal{C}$, and let $S \setminus T \neq \emptyset$ (if $S \subseteq T$ then simply exchange S and T). By applying Corollary 14 to S and $\overline{T}$, we obtain that $S \cup \overline{T} = V$, i.e. $T \subseteq S$. As this is true for any two sets in $\mathcal{C}$, we have that $\mathcal{C}$ is a chain. ■

What remains now is to show that any inside element v can be placed in one of the cells of $P \setminus \{\ell(S) \colon S \in \mathcal{F}\}$. Fix an inside element v, and let $\mathcal{F}_v = \{S \in \mathcal{F} \colon v \in S\}$. For any $S \in \mathcal{F}$, let $R(S)$ be the intersection of the interior of the polygon with the open halfplane on the left of the line through a_S and $a_{\overline{S}}$ (i.e. the halfplane already containing the elements $S \cap O$). To prove that v can be placed adequately, we need to prove that $\cap_{S \in \mathcal{F}_v} R(S) \neq \emptyset$. Interestingly, this is an implication of Helly's theorem in 2 dimensions (see e.g. [14]):

Theorem 16 (Helly)**.** *Let a collection of convex subsets of $\mathbb{R}^d$ have the property that any collection of up to $d+1$ of them have a non-empty intersection. Then all of them have a common intersection.*

To apply Helly's theorem, we first consider the intersection of two such regions, and then show that the intersection of three regions essentially reduces to the intersection of two regions. In the proofs below, $cl(\cdot)$ denotes the closure operator.

Proposition 17. *Let $S_1, S_2 \in \mathcal{F}_v$. Then $R(S_1) \cap R(S_2) \neq \emptyset$.*

Proof. If $S_1 \cup S_2 = V$ then, by Proposition 9, the diagonals $\ell(S_1)$ and $\ell(S_2)$ do not cross. If $R(S_1)$ and $R(S_2)$ were disjoint then the fact that $S_1 \cup S_2$ contains all outside elements would imply that S_1 and $\overline{S_2}$ have identical intersections with O. But this case was taken care of when we introduced additional polygon vertices, as we made sure that if $S_1 \cap O = \overline{S_2} \cap O$ and $S_1 \cap S_2 \neq \emptyset$ then $R(S_1) \cap R(S_2) \neq \emptyset$. In fact, $R(S_1) \cap R(S_2)$ is a strip extending between $\ell(S_1)$ and $\ell(S_2)$; more formally, $\ell(S_i) \subset cl\big(R(S_1) \cap R(S_2)\big)$ for $i = 1, 2$.

If $S_1 \cup S_2 \neq V$ then Corollary 14 implies that $S_1 \cap S_2$ contains an outside element w, and the result follows trivially as the cell containing w will be in $R(S_1) \cap R(S_2)$. ■

Finally, we consider the intersection of 3 regions defined by sets in $\mathcal{F}_v$.

Proposition 18. *Let $S_1, S_2, S_3 \in \mathcal{F}_v$. Then $R(S_1) \cap R(S_2) \cap R(S_3) \neq \emptyset$.*

Proof. If $S_1 \cap S_2 \cap S_3$ contains an outside element, we are done. Thus, we assume that $S_1 \cap S_2 \cap S_3 \cap O = \emptyset$. We consider two cases.

Case 1. For every pair (i, j), $S_i \cup S_j \neq V$. In this case, Corollary 14 implies that $T_{ij} = S_i \cap S_j \cap O \neq \emptyset$ for $1 \leq i < j \leq 3$. As the intersection of any two of the T_{ij}'s is empty ($S_1 \cap S_2 \cap S_3 \cap O = \emptyset$), we have that the three sets $S_{i-1} \cap S_{i+1} \setminus S_i$ are non-empty for $i \in \{1, 2, 3\}$. As S_1, S_2, S_3 do not form a 3-cycle, we must have that $S_1 \cup S_2 \cup S_3 = V$. We also know that $\overline{S_1}, \overline{S_2}, \overline{S_3}$ do not form a 3-cycle, and as $v \notin \bigcup_i \overline{S_i}$, we have that $\overline{S_i} \cap \overline{S_j} \setminus \overline{S_k} = \emptyset$ for some permutation i, j, k of $\{1, 2, 3\}$. The fact that $\bigcap_l \overline{S_l} = \emptyset$ then implies that $\overline{S_i} \cap \overline{S_j} = \emptyset$, i.e. that $S_i \cup S_j = V$ contradicting our assumption.

Case 2. $S_i \cup S_j = V$ for some $i, j \in \{1, 2, 3\}$, $i \neq j$. As in the first part of the proof of Proposition 17, we derive that both diagonals $\ell(S_i)$ and $\ell(S_j)$ are in $cl\big(R(S_i) \cap R(S_j)\big)$. Let k be the remaining index. If $\ell(S_k)$ crosses either $\ell(S_i)$ or $\ell(S_j)$ then a segment of $\ell(S_k)$ is in $cl\big(R(S_1) \cap R(S_2) \cap R(S_3)\big)$ showing non-emptyness of the intersection. On the other hand, if $\ell(S_k)$ crosses neither $\ell(S_i)$ nor $\ell(S_j)$ then either $R(S_k) \supseteq R(S_i) \cap R(S_j)$ (and we are done) or $R(S_k) \cap R(S_i) = \emptyset$ or $R(S_k) \cap R(S_j) = \emptyset$. In these latter cases, we obtain a contradiction from Proposition 17. ■

Helly's theorem thus shows that every inside element can be placed in one of the cells. As we did not make any assumptions on the position of the vertices $a_1, \ldots, a_q$ (except that they are in convex position), the polygon representation obtained is deformable, and this completes our proof of Theorem 2.

We now state a few more properties of the polygon representation. First, we can strengthen Proposition 9.

Proposition 19. *Consider a symmetric family with a connected cross graph and no 3-cycles or combs, and consider its deformable polygon representation. Let $S_1, S_2 \in \mathcal{F}$. Then the following are equivalent:*

1. *the diagonals $\ell(S_1)$ and $\ell(S_2)$ do not cross,*
2. *the sets $S_1 \cap O$ and $S_2 \cap O$ do not cross in O,*
3. *the sets S_1 and S_2 do not cross.*

Indeed, Proposition 9 says that 1. and 2. are equivalent, 3. always implies 2., and 1. implies 3. simply by the existence of the polygon representation.

In Proposition 15, we have shown that sets with the same intersection with O form a chain, and this was the basis for introducing new polygon vertices after having placed the outside elements. We now show that this can only happen if we have 4-cycles.

Proposition 20. *If $S_1, S_2 \in \mathcal{F}$ with $S_1 \neq S_2$ and $S_1 \cap O = S_2 \cap O$ then $\mathcal{F}$ contains a 4-cycle.*

Proof. By Proposition 15, we can assume that $S_1 \subset S_2$. Let $v \in S_2 \setminus S_1$; v must be inside as $(S_2 \setminus S_1) \cap O = \emptyset$. Let $C_1, \ldots, C_k$ be a cycle for v. Remember that $O \subseteq \cup_i C_i$ and that $C_i \cap C_{i+1} \cap O \neq \emptyset$ (by Corollary 14). Thus, there is an index i so that C_i contains elements of both A and $O \setminus A$, where $A = S_1 \cap O = S_2 \cap O$. If C_i contains all elements of A then C_i and S_2 do not cross in O (as $S_2 \cap O = A$), and by Proposition 19, they do not cross (in V); thus either $S_2 \subseteq C_i$ or $C_i \subseteq S_2$. This is a contradiction as $v \in S_2 \setminus C_i$ and $S_2 \cap O \subsetneq C_i \cap O$. Therefore, C_i crosses S_1 and S_2. Similarly, this implies that there exists another index j such that C_j crosses S_1 and S_2. If i and j are consecutive then we have a 3-cycle C_i, C_j and S_1, a contradiction. Otherwise, we have a 4-cycle composed of C_i, S_1, C_j and $\overline{S_2}$. ∎

If we have a chain $S_1 \subset \cdots \subset S_k$ with $S_1 \cap O = S_k \cap O$, we say that the inside elements in $S_k \setminus S_1$ are *sandwiched* by this chain. The next proposition shows that any inside element can be sandwiched by at most 2 chains; otherwise, we would have a 3-cycle, a contradiction. This will be useful when establishing a tight bound on the size of families with no 3-cycles or combs in Proposition 26.

Proposition 21. *For $v \in V$, let $\mathcal{A} = \big\{ \{A, O \setminus A\} \colon \exists\, S_1, S_2 \in \mathcal{F}$ s.t. $v \notin S_1$, $v \notin S_2$, $S_1 \cap O = A$, and $S_2 \cap O = O \setminus A \big\}$. Then $|\mathcal{A}| \leq 2$.*

Proof. Assume on the contrary that a given $v \in V$ is sandwiched by 3 chains. Thus, we have $S_{ik} \in \mathcal{F}$ for $i = 1, 2, 3$ and $k = 1, 2$ with $v \notin S_{ik}$ for all i and k, $S_{i1} \cap O = A_i$ for all i, $S_{i2} \cap O = O \setminus A_i$ for all i, and the 6 sets A_i's and $O \setminus A_i$'s are all distinct. Furthermore, possibly exchanging S_{i1} and S_{i2} (replacing A_i by $O \setminus A_i$), we can assume that $w \in A_1 \cap A_2 \cap A_3$ for some $w \in O$.

We claim that, for any $i, j \in \{1,2,3\}$ with $i \neq j$, A_i and A_j cross in O. Otherwise, suppose that $A_i \subset A_j$. Applying Proposition 19 to $\overline{S_{i2}}$ and S_{j1} (recall that $\overline{S_{i2}} \cap O = A_i$ and $S_{j1} \cap O = A_j$, we get that $\overline{S_{i2}} \subset S_{j1}$, a contradiction as $v \in \overline{S_{i2}} \setminus S_{j1}$.

Therefore, the A_i's are pairwise mutually crossing in O and correspond all to circular intervals of O containing w. This means that their start elements (of their circular interval) in O are distinct, so are their end elements, and these elements are ordered in the same way among the 3 sets. Say that A_2 corresponds to the middle interval. Then A_1, $O \setminus A_2$ and A_3 are such that the intersection of any two of them minus the third one is non-empty. This implies that S_{11}, S_{22} and S_{31} form a 3-cycle for v, a contradiction. ∎

4. Mincuts and Near-Mincuts

In this section, we show that the main configurations discussed in the previous sections do not exist for families of cuts of sufficiently small value compared to the edge-connectivity. In particular, we show that each connected component of the cross graph of 6/5-near-mincuts has a deformable polygon representation. Recall that λ denotes the edge-connectivity and an α-near-mincut is a cut whose value is (strictly) less than $\alpha\lambda$.

In what follows, let $d(X, Y)$ denote the total weight of edges connecting $X, Y \subset V$ in a weighted graph $G = (V, E)$. We first show that sufficiently near-mincuts have no short cycles as defined in Definitions 1 and 3.

Lemma 22 (Excluded Cycles)**.** *Let k δ-near-mincut sides C_i for $i \leq k$ form a k-cycle. Then $\delta > 1 + 1/k$.*

Proof. For $k = 3$ the result follows from 3-way submodularity of the cut function, see Lovász [13], exercise 6.48 (c):

$$\begin{aligned} 3\delta\lambda > d(C_1) + d(C_2) + d(C_3) &\geq d(C_1 \cap C_2 \setminus C_3) + d(C_1 \cap C_3 \setminus C_2) \\ &\quad + d(C_2 \cap C_3 \setminus C_1) + d(\overline{C_1 \cup C_2 \cup C_3}) \geq 4\lambda. \end{aligned}$$

3-way submodularity can be established by observing that the contribution of any edge is at least as large on the left-hand-side as on the right-hand-side.

For $k \geq 4$, a similar edge counting argument gives:

$$d(\cup_i C_i) + \sum_{i \leq k} d(C_i \cap C_{i+1}) \leq \sum_{i \leq k} d(C_i).$$

Indeed, (i) an edge that contributes to the left-hand-side also contributes to the right-hand-side,(ii) no edge contributes more than twice to the left-hand-side, and (iii) an edge that contributes twice to the left-hand-side must be either between some $C_i \cap C_{i+1}$ and $\overline{\cup_j C_j}$ or between $C_i \cap C_{i+1}$ and $C_j \cap C_{j+1}$ for $i \neq j$; in all cases, it contributes at least twice to the right-hand-side. Since $d(C_i) < \delta\lambda$, $d(C_i \cap C_{i+1}) \geq \lambda$ for $i \leq k$ and $d(\cup_i C_i) \geq \lambda$, the above inequality implies that $\delta > 1 + \frac{1}{k}$ as required. ■

In particular, 6/5-near-mincuts may not contain k-cycles for $k \leq 5$, 4/3-near-mincuts have no 3-cycles, and mincuts contain no cycles at all.

Lemma 23 (Excluded Combs)**.** *There are no four* 6/5*-near-mincut sides* T_1*,* T_2*,* T_3 *and* H *which form a comb as defined in Definition* 2*.*

Proof. Definition 2 gives 3 disjoint and nonempty subsets C_1, C_2 and C_3 of H: Either C_i is $H \cap \big(T_i \setminus (T_{i-1} \cup T_{i+1})\big)$ or C_i is $H \cap (T_{i-1} \cap T_{i+1} \setminus T_i)$. Similarly, it gives 3 disjoint nonempty subsets D_1, D_2 and D_3 of $\overline{H}$. By a counting argument, one observes that

$$(2) \quad d(T_1) + d(T_2) + d(T_3) + 2d(H) \geq d(C_1) + d(C_2) + d(C_3) + d(D_1) + d(D_2) + d(D_3).$$

Indeed, either an edge does not cross H and the counting argument is similar to the derivation of 3-way submodularity, or the edge crosses H in which case its multiplicity is at most 2 on the right-hand-side and at least 2 on the left-hand-side. Inequality (2) now implies that $5\delta\lambda > 6\lambda$. ■

We have thus derived that the symmetric family of 6/5-near-mincut sides does not contain any comb and any k-cycles for $k \leq 5$. Thus each connected component of its cross graph admits a deformable polygon representation.

5. Tree Hierarchy

In this section, we show that the connected components of the cross graph of any symmetric family $\mathcal{F} \subseteq V$ of sets can be arranged in a tree structure. We first derive it for families with no combs and cycles, and in so doing, rederive the tree structure of a (slightly modified) cactus representation. Then we show that the tree structure depends only on the connected components of the cross graph and thus can be applied to arbitrary symmetric families, including those having deformable polygon representations for each connected component of the cross graph. This tree structure will then be used to derive bounds on the size of the representation and (in the next section) on the cardinality of symmetric families with no 3-cycles or combs.

To fix notation, let $\mathcal{F}_1, \ldots, \mathcal{F}_q$ represent the sets in each of the connected components of the cross graph. Also, for any $1 \leq i \leq q$, let P_i denote the partition of V induced by $\mathcal{F}_i$, i.e. the members of P_i correspond precisely to the atoms of $\mathcal{F}_i$ as defined at the beginning of Section 3. Observe that, for any i, the number of atoms in P_i can either be 2 or greater or equal to 4. Indeed, if $|P_i|$ was 3 the corresponding pairs of sets would not be crossing.

The tree structure on the connected components essentially follows from laminarity. It is similar to the usual cactus representation for the mincuts of a graph, except that, there, the cuts represented by a cycle of the cactus do not quite form a single connected component of the cross graph. Indeed the cut corresponding to a single atom (i.e. two consecutive edges of the cycle) does not cross any of the other cuts represented by the cycle and thus do not belong to the same connected component of the cross graph. Here, however, we redefine the cactus representation and assume that the cuts represented by a cycle of length k with $k \neq 2$ of the cactus are the cuts obtained by removing two *non-consecutive* edges of the cycle. The rest of the definition is unchanged. It is rather easy to see using classical arguments that the mincuts of a graph admit a cactus representation for this slightly modified notion of a cactus representation. For completeness, we provide here a somewhat different proof, of a geometric nature as in the rest of this paper. For generality and to be able to later apply it to our deformable polygons, we state it in terms of symmetric families with no cycles or combs (recall that from Lemmas 22 and 23, mincut sides do not have any cycles or combs).

Proposition 24. *Let $\mathcal{F} \subseteq 2^V$ be a symmetric family of sets with no cycles and no combs. Let $\mathcal{F}_i$ be the connected components of its cross graph and let P_i be the partition of V corresponding to the atoms of $\mathcal{F}_i$. Then there exists a cactus $H = (N, A)$ and a mapping $\phi \colon V \to N$ such that*

1. *H has no cycle of length 3,*
2. *there is a 1-to-1 correspondence between the connected components $\mathcal{F}_i$ of the cross graph and the cycles C_i of H,*
3. *the removal of the edges of $C_i = u_1 - u_2 \cdots - u_k - u_1$ break H into k (depending on i) connected components, $A_1, \ldots, A_k \subset N$ where $u_j \in A_j$ such that $P_i = \left\{ \phi^{-1}(A_j) \colon 1 \leq j \leq k \right\}$,*
4. *for each set $S \in \mathcal{F}$, there is a unique cycle C_i in H and two edges of C_i which are non-consecutive if the cycle is not of length 2, whose removal partitions N into U and $N \setminus U$ with $S = \phi^{-1}(N)$.*

To complete the proof that mincuts admit a (modified) cactus representation, it remains to show that the removal of *any* 2 edges of a cycle – non-consecutive if the cycle has length different from 2 – gives a mincut of G. This follows from submodularity of the cut function (union or intersection of crossing mincuts is a mincut); this is left to the reader.

Proof. From Proposition 6 and Lemmas 22 and 23, we know that $\mathcal{F}$ is a circular representable hypergraph. Consider one such circular ordering $v_1, \ldots, v_n$ where $n = |V|$. If there are several, take one in which the elements in each atom of $\mathcal{F}$ appear consecutively in the ordering (one could for example first shrink the atoms of $\mathcal{F}$).

We provide a geometric construction of the cactus. Take a circle and divide it into n arcs representing the vertices in the circular ordering. For each pair of complementary sets $S, \overline{S}$, draw the corresponding chord that has S on one side and $\overline{S}$ on the other. Observe that two chords corresponding to S_1 and S_2 will cross in the geometric sense (i.e. will have an intersection in their relative interior) if and only if the sets S_1 and S_2 cross.

Consider any of the connected components of the cross graph, say $\mathcal{F}_i$ with partition P_i. The chords defining $\mathcal{F}_i$ connect $k = |P_i|$ points on the circle, say p_1, $p_2, \ldots, p_k$, and the arcs between these points correspond to the sets in P_i. See Figure 13. Let R_i be the convex hull of these points. The fact that sets from different components do not cross imply that the chord for a set of a different component $\mathcal{F}_j$ can only intersect R_i on its boundary. This means that the relative interiors of any two such regions, R_i and R_j, will be disjoint. Replace all the chords corresponding to $\mathcal{F}_i$ by a star with root r_i and k spokes connected to $p_1, \ldots, p_k$. To make sure that the stars for different components do not cross, place r_i in the relative interior of R_i. This gives a plane graph D with the outside region delimited by our circle. Now define H to be its dual graph except that we do not create a node of H for the outside face of D. The nodes of H corresponding to the inside faces of D along the circle are labelled with the atoms of $\mathcal{F}$. The node set N of H is thus the set of *bounded* faces of D. Observe that H will be the union of cycles, one for each vertex r_i. Furthermore, every edge is in precisely one cycle; thus, H is a cactus. All the claims in the statement of the Proposition follow easily from the construction itself. ■

This construction of this modified cactus representation has slightly more empty nodes than the usual constructions; the size is still linear as shown below.

Proposition 25. *The modified cactus representation for a symmetric family $\mathcal{F} \subseteq 2^V$ with no cycles and no combs has at most $3n-4$ nodes and at most $5n-8$ edges where $n = |V| \geq 2$.*

Proof. Let $N_p(n)$ and $E_p(n)$ resp. represent the maximum number of nodes and edges resp. of our cactus representation for families with at most p *non-trivial* connected components, where we define a connected component as non-trivial if it contains more than one pair of complementary sets. We proceed by induction on p.

If $\mathcal{F}$ has no crossing sets ($p = 0$) then the number of chords in our construction is at most $2n-3$ (the maximum cardinality of a laminar family with no complementary sets). Each chord will lead to 2 edges of the cactus and the cactus will be a tree of cycles of length 2. Thus, its number of edges will be at most $E_0(n) = 4n-6 \leq 5n-8$ for $n \geq 2$ and the number of nodes will be at most $N_0(n) = 2n-2 \leq 3n-4$.

Suppose $\mathcal{F}$ has p non-trivial connected components for $p > 0$. Let $\mathcal{F}_i$ be a connected component of the cross graph which induces a partition P_i with $k \geq 4$ atoms. Let $n_1, \ldots, n_k$ be the cardinalities of these atoms. Geometrically, we can partition $\mathcal{F} \setminus \mathcal{F}_i$ into k symmetric families $\mathcal{G}_1, \ldots, \mathcal{G}_k$ in a natural way: for the jth family $\mathcal{G}_j$, we only keep those chords that are on the side of $[p_j, p_{j+1}]$ opposite to R_i. Observe that $[p_j, p_{j+1}]$ might be one of the chords represented by this family. Collectively, the representations for the $\mathcal{G}_j$'s account for all nodes and edges of the representation for $\mathcal{F}$, except

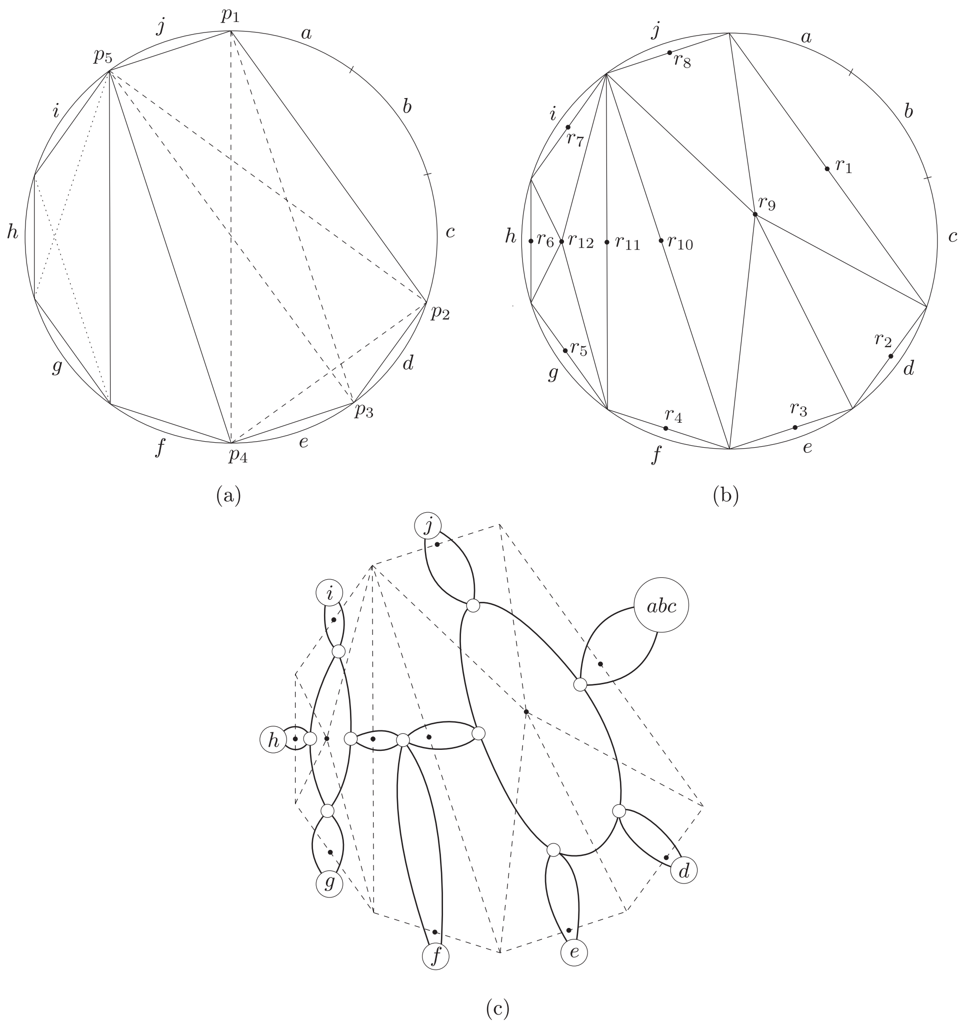

Fig. 13. Obtaining a cactus. (a): The cyclic ordering. The sets are represented by chords. There are 12 connected components, the dashed one, the dotted one, and 10 others with only one chord. The dashed component connects $p_1, \ldots, p_5$. (b): Replacing every component by a star with root r_i. (c): Inside faces along the circle are labelled with the atoms of $\mathcal{F}$, and the dual graph is the cactus.

for the k edges of the cycle corresponding to $\mathcal{F}_i$. $\mathcal{G}_j$ has at most $n_j + 1$ atoms (and fewer non-trivial components), and by induction, we therefore get that:

$$N_p(n) \leq \sum_{j=1}^{k} N_{p-1}(n_j + 1) \leq \sum_j (3n_j - 1) = 3n - k \leq 3n - 4$$

and

$$E_p(n) \leq \sum_{j=1}^{k} E_{p-1}(n_j + 1) + k \leq \sum_j (5n_j - 3) + k = 5n - 2k \leq 5n - 8,$$

since $k \geq 4$. ■

The bounds given in Proposition 25 are tight whenever n is even, and this is achieved when we have $(n-1)/2$ disjoint cycles of length 4 in the cactus linked by cycles of length 2, see Figure 14.

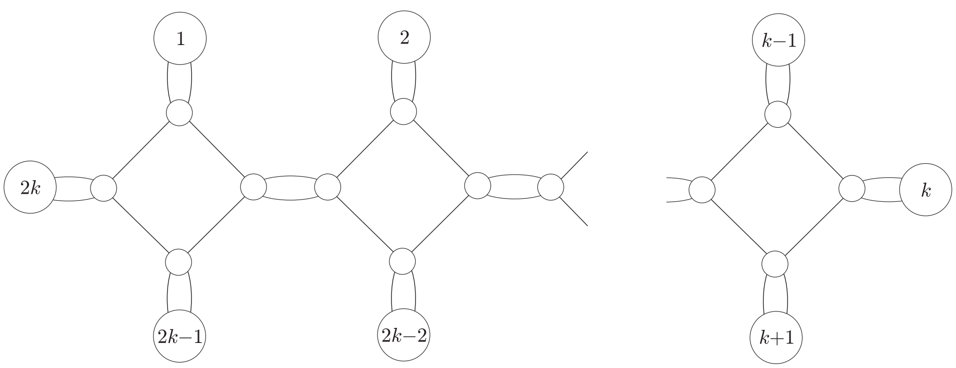

Fig. 14. A cactus with $3n - 4$ nodes and $5n - 8$ edges for a symmetric family with no cycles or combs defined on a ground set of size $n = 2k$.

This cactus representation provides a way to combine representations (i.e. cycles on partitions of V) for each connected component of the cross graph in a tree structure in which certain atoms of different connected components (i.e. cycles) are identified (those corresponding to the same node of the cactus). To highlight the tree T (and get rid of the cycles), we can replace each cycle of the cactus H by a star rooted at a new node representing this connected component of the cross graph. In other words, the nodes of T consist of (i) one node in C for each connected component of the cross graph and also (ii) one node in N for each node of the cactus; the latter ones correspond to atoms that have been identified from one or several connected components of the cross graph. T has an edge between a node c in C and a node u in N if one of the atoms of the connected component corresponding to c is associated with node u of the cactus.

This can be generalized to any symmetric family (independently of whether each connected component of the cross graph can or cannot be represented by polygons). For any symmetric family $\mathcal{F}$, consider the connected components $\mathcal{F}_i$ of the cross graph and let P_i be the atoms of $\mathcal{F}_i$. Now, for each i, arbitrarily choose a cyclic ordering on the atoms in P_i and define $\mathcal{G}_i$ to be the family of sets (with the same atoms as $\mathcal{F}_i$) corresponding to cyclic intervals containing at least 2 and at most $k_i - 2$ of the atoms in P_i, where $k_i = |P_i| \geq 4$ is the number of atoms of $\mathcal{F}_i$; if $k_i = 2$ (i.e. the component is trivial), we simply let $\mathcal{G}_i = \mathcal{F}_i$. Observe that the family $\mathcal{G} = \cup \mathcal{G}_i$ has no cycles or combs since (i) the connected components of the cross graph of $\mathcal{G}$ are still the $\mathcal{G}_i$'s, (ii) any cycle or comb would need to be contained within a connected component of the cross graph and (iii) the $\mathcal{G}_i$'s have no combs or cycles by construction. By Proposition 24, the family $\mathcal{G}$ has a cactus representation, and this means that atoms of different connected components of the cross graph of $\mathcal{G}$, and thus of $\mathcal{F}$, are identified. This gives a tree T for $\mathcal{G}$ and thus also for $\mathcal{F}$. If each connected component of the cross graph has a deformable polygon representation, these polygon representations together with the tree T form our representation as a *tree of deformable polygons.*

6. Size of the Family

In this section, we deduce from the representation as a tree of deformable polygons that any symmetric family of sets $\mathcal{F} \subset 2^V$ with no 3-cycles and no combs has at most $\binom{n}{2}$ complementary pairs where $n = |V|$. For near-mincuts, this shows that there are at most $\binom{n}{2}$ 6/5-near-mincuts, although this is known even for 4/3-near-mincuts [19, 9] and a direct proof is simpler. See also [10] for a proof that there are at most $O(n^2)$ 3/2-near-mincuts.

Proposition 26. *Let $\mathcal{F} \subseteq 2^V \setminus \{\emptyset, V\}$ be a symmetric family of sets with no 3-cycles or combs. Then $|\mathcal{F}| \leq n(n-1)$ where $n = |V|$, i.e. the number of complementary pairs is at most $\binom{n}{2}$.*

Proof. We will first focus on a connected component $\mathcal{F}_i$ of the cross graph with $k \geq 4$ atoms, and show that our construction of a deformable polygon gives rise to at most $k(k-3)/2$ diagonals. Let $k_{\text{out}} \geq 4$ and k_{in} be the number of outside and inside atoms respectively; thus $k = k_{\text{out}} + k_{\text{in}}$. If $\mathcal{F}_i$ has no 4-cycles then our polygon has k_{out} sides (see Proposition 20) and its number of diagonals is at most $k_{\text{out}}(k_{\text{out}}-3)/2 \leq k(k-3)/2$. If we have 4-cycles then by Proposition 21, the number of diagonals is at most $k_{\text{out}}(k_{\text{out}}-3)/2+2k_{\text{in}} \leq k(k-3)/2$ as $(k_{\text{out}}+1)(k_{\text{out}}-2)/2-k_{\text{out}}(k_{\text{out}}-3)/2 \geq 2$ for $k_{\text{out}} \geq 3$. Thus, a connected component of the cross graph on $k \geq 4$ atoms has at most $k(k-3)/2$ diagonals.

Consider now the various connected components of the cross graph and, as in the proof of Proposition 25, we proceed by induction on p, the number of non-trivial connected components. Let $S_p(n)$ denote the maximum number of complementary pairs for families with at most p non-trivial components. If $p = 0$ then we have $S_0(n) \leq 2n-3 \leq \binom{n}{2}$ complementary pairs, see Proposition 25. If $p > 0$, let $\mathcal{F}_i$ be one of those non-trivial components on $k \geq 4$ atoms with cardinalities $n_1, \ldots, n_k$ where $n = \sum_{j=1}^{k} n_j$. We use the same notation as in Proposition 25. From the tree structure, we get that:

$$\begin{aligned}
S_p(n) &\leq \sum_{j=1}^{k} S_{p-1}(n_j+1) + \frac{k(k-3)}{2} \leq \sum_{j=1}^{k} \binom{n_j+1}{2} + \frac{k(k-3)}{2} \\
&\leq (k-1) + \frac{(n-k+2)(n-k+1)}{2} + \frac{k(k-3)}{2} \\
&= \binom{n}{2} - (n-k)(k-2) \leq \binom{n}{2},
\end{aligned}$$

the third inequality follows from the fact that the maximum of a convex function over $n_j \geq 1$ for $j = 1, \ldots, k$ and $\sum_j n_j = n$ is attained for all but one n_j equal to 1. ■

We should point out that the existence of a (non-deformable) polygon representation is not enough to prove the bound of $\binom{n}{2}$; for example, all cuts of K_4 admit a representation as a tree of (non-deformable) polygons although there are $7 > \binom{4}{2}$ of them. One can, however, prove a slightly weaker bound by observing that a connected component of the cross graph on k atoms has at most $\binom{k}{2}$ complementary pairs (by Schläfli's result) and the same argument as in the proof above gives an overall bound of $\binom{n}{2} + 2n - 4$.

7. Conclusion

We have derived a representation for symmetric families that do not contain 3-cycles or combs, and this applies to the family of 6/5-near-mincuts in a graph. We refer the reader to the Ph.D. thesis [3] of the first author for algorithmic issues and applications of this representation. It would be interesting to find a representation for families that may contain combs but do not have 3-cycles; this would allow to represent 4/3-near-mincuts.

Acknowledgements. The second author would like to thank the hospitality of the Research Institute for Mathematical Sciences, Kyoto University, where a great deal of the writing of this paper was done.

References

[1] A. A. Benczúr, *On the structure of near-minimum edge cuts,* Technical report MIT/LCS/TR-639 (1994).

[2] A. A. Benczúr, A representation of cuts witihin 6/5 times the edge connectivity with applications, in: *Proc. 34th Annual Symp. on Found. of Comp. Sci.* (1995), pages 92–102.

[3] A. A. Benczúr, *Cut Structures and Randomized Algorithms in Edge-Connectivity Problems,* Ph.D. Thesis, M.I.T. (1997).

[4] Y. Dinitz (E. A. Dinits), A. V. Karzanov and M. L. Lomonosov, On the Structure of a Family of Minimal Weighted Cuts in a Graph, in: *Studies in Discrete Optimization* (in Russian), A. A. Fridman (Ed), Nauka, Moscow (1976), pages 290–306.

[5] Y. Dinitz and Z. Nutov, A 2-level cactus model for the minimum and minimum+1 edge-cuts in a graph and its incremental maintenance, in: *Proc. of the 27th Symposium on Theory of Computing* (1995), pages 509–518.

[6] P. Duchet, Classical Perfect Graphs: An Introduction with Emphasis on Triangulated and Interval Graphs, *Annals of Discrete Mathematics,* **21** (1984), 67–96.

[7] L. Fleischer, Building Chain and Cactus Representations of All Minimum Cuts from Hao–Orlin in the Same Asymptotic Run Time, *Journal of Algorithms,* **33** (1999), 51–72.

[8] H. N. Gabow, *Applications of a Poset Representation to Edge Connectivity and Graph Rigidity,* in: *Proc. of the 32nd Annual Symp. on Found. of Comp. Sci.* (1991), pages 812–821.

[9] M. X. Goemans and V. S. Ramakrishnan, Minimizing Submodular Functions over Families of Sets, *Combinatorica*, **15** (1995), 499–513.

[10] M. Henzinger and D. P. Williamson, On the Number of Small Cuts in a Graph, *Information Processing Letters* (1996), **59**, 41–44.

[11] A. V. Karzanov and E. A. Timofeev, Efficient Algorithms for Finding all Minimal Edge Cuts of a Nonoriented Graph, *Cybernetics* (1986), 156–162. Translated from *Kibernetika,* **2** (1986), 8–12.

[12] J. Lehel, F. Maffray and M. Preissmann, Graphs with Largest Number of Minimum Cuts, *Discrete Applied Mathematics*, **65** (1996), pages 387–407.

[13] L. Lovász, *Combinatorial Problems and Exercises,* North-Holland, Amsterdam (1979).

[14] J. Matoušek, *Lectures on Discrete Geometry,* Springer (2002).

[15] H. Nagamochi and T. Kameda, Canonical Cactus Representation for Minimum Cuts, *J. of Japan Society for Industrial and Applied Mathematics*, **11** (1994), 343–361.

[16] H. Nagamochi and T. Kameda, Constructing Cactus Representation for all Minimum Cuts in an Undirected Network, *Operations Research Society of Japan*, **39** (1996), pages 135–158.

[17] H. Nagamochi, S. Nakamura and T. Ishii, Constructing a Cactus for Minimum Cuts of a Graph in $O(mn+n^2 \log n)$ time and $O(m)$ space, *Inst. Electron. Inform. Comm. Eng. Trans. Information and Systems,* **E86-D** (2003), 179–185.

[18] H. Nagamochi, Y. Nakao and T. Ibaraki, A Fast Algorithm for Cactus Representations of Minimum Cuts, *J. of Japan Society for Industrial and Applied Mathematics,* **17** (2000), 245–264.

[19] H. Nagamochi, K. Nishimura and T. Ibaraki, Computing all small cuts in undirected networks, *SIAM Disc. Math.,* **10** (1997), 469–481.

[20] D. Naor and V. V. Vazirani, Representing and Enumerating Edge Connectivity Cuts in RNC, in: *Proc. Second Workshop on Algorithms and Data Structures,* Lecture Notes in Computer Science **519**, Springer–Verlag (1991), pages 273–285.

[21] A. Quilliot, Circular Representation Problems on Hypergraphs, *Discrete Mathematics,* **51** (1984), pages 251–264.

[22] L. Schläfli, Theorie der vielfachen Kontinuität, 1901. In: *Gesammelte Mathematische Abhandlungen,* volume I (1950), pages 167–387. Verlag Birkhäuser.

[23] A. Tucker, A Structure Theorem for the Consecutive 1's Property, *JCT B,* **12** (1972), pages 153–162,.

András A. Benczúr

Computer and Automation Research Institute of the Hungarian Academy of Sciences

`e-mail: benczur@sztaki.hu`

Michel X. Goemans

M.I.T., Department of Mathematics

`e-mail: goemans@math.mit.edu`