A 2-level cactus tree model for the system of minimum and minimum+1 edge cuts of a graph and its incremental maintenance. Part I: the odd case

Yefim Dinitz [*†]
Dept. of Computer Science
Technion, Haifa, 32000, Israel
e-mail: dinitz@cs.technion.ac.il

Zeev Nutov
Dept. of Applied Mathematics
Technion, Haifa, 32000, Israel
e-mail: nutov@mpi-sb.mpg.de

January 11, 1999

Abstract

The known cactus tree model represents the minimum edge cuts of a graph in a clear and compact way and is used in related studies. We generalize this model to represent the minimum and minimum+1 edge cuts; for this purpose, we use new tools for modeling connectivity structures. The obtained representations are different for λ odd and even; their size is linear in the number of vertices of the graph. Let λ denote the cardinality of a minimum edge cut. We suggest efficient algorithms for the maintenance of our representations, and, thus, of the $(\lambda+2)$-connectivity classes of vertices (called also "$(\lambda+2)$-components") in an arbitrary graph undergoing insertions of edges. The time complexity of those algorithms, for λ odd and even, is the same as achieved previously for the cases $\lambda = 1$ and 2, respectively. In this paper we consider the case of odd $\lambda \geq 3$. The case of even connectivity is considered in the companion paper (Part II).

1 Introduction

Connectivity is a fundamental property of graphs, which has important applications in network reliability analysis, in network design problems and in other applications. For many connectivity problems, a clear and compact representation of minimum and near minimum cuts of a graph is of much help. In this paper we consider only edge-connectivity and edge cuts of an undirected multigraph (henceforth, we omit the prefix "edge" and say "graph" instead of "multigraph"). Recently, connectivity augmentation problems and the problem

[*]Up to 1990, E. A. Dinic, Moscow.

[†]This research was supported partly by the Fund for the Promotion of Research at the Technion, Israel.

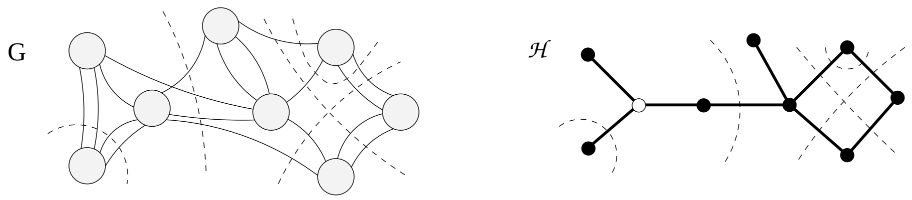

Figure 1: Cactus-tree model of a 4-connected graph. (The gray blobs represent vertex classes of 5-connectivity. Some cuts of the graph and their representing cuts of the cactus tree are shown by dashed lines.)

of maintaining the vertex classes of k-connectivity (called in literature also "k-components") of a dynamic undirected graph gave an impetus for development of connectivity models and related algorithms. Some graph structures have been discovered [7, 25, 14, 16, 5, 6, 15, 26, 13, 12, 2]; most of them serve the incremental setting, i.e., they can be efficiently updated when the graph undergoes edge insertions.

Let $G = (V, E)$ be an undirected connected graph on $n \geq 2$ vertices. A set of edges $C \subset E$ is an (*edge*)-*cut* of G if there exists a *bisection* (partition into two nonempty parts) of V such that C is the set of edges having endnodes in distinct parts. In a connected graph, there is a bijective correspondence between the cuts of the graph and the bisections of its vertex set. Let λ denote the minimum cardinality of a cut of G. When analyzing the connectivity of G, the natural first "stratum" is the set of its minimum cuts (i.e., cuts of cardinality λ) and the set of vertex classes of $(\lambda+1)$-connectivity that V is "cut into" by these cuts. Both these are represented, in a compact and simple way, by the (inclusion) minimal cuts of the cactus tree model $(\mathcal{H}, \varphi)$ [7]: $\mathcal{H}$ is a tree-of-edges-and-cycles, or cactus tree, for short (i.e., a connected graph such that every its block is an edge or a cycle) and φ is a mapping from V to the node set of $\mathcal{H}$ (see Fig. 1 for an example). The minimal cuts of a cactus tree have a simple structure: any such cut is either a bridge or a pair of edges belonging to the same cycle. The cactus tree model has the following properties:

(i) For every node $\mathcal{N}$ of $\mathcal{H}$, $\varphi^{-1}(\mathcal{N})$ is either a $(\lambda+1)$-class of G or the empty set;

(ii) The mapping φ^{-1} takes the set of bisections corresponding to the minimal cuts of $\mathcal{H}$ onto the set of bisections corresponding the set of minimum cuts of G.

(iii) The number of edges in $\mathcal{H}$ is linear in the number of $(\lambda+1)$-classes, i.e., is $O(n)$.

For odd λ, $\mathcal{H}$ is unique and has no cycles (i.e, is a tree), and this representation is bijective. For even λ, $\mathcal{H}$ is unique up to conversions of nodes $\mathcal{N}$ of degree 3 with $\varphi^{-1}(\mathcal{N}) = \emptyset$ into cycles of length three and vice versa, and this representation is almost bijective (for a formal proof and for the only case of nonbijectiveness, where a cut of G is represented by two model cuts, see [23]). In this paper, we consider the unique version of $\mathcal{H}$ in which all empty nodes of degree 3 are replaced by cycles.

There are several applications requiring the knowledge of all mincuts and the way they are structured. Note that there can be $\Omega(n^2)$ mincuts, and thus the space required to list all of them can be $\Omega(n^3)$ if every cut is described as a bisection of V, and $\Omega(\lambda n^2)$ if a cut is described as a set of edges. The cactus tree-model represents the mincuts not only in a compact way, but also establishes reductions for several connectivity problems from arbitrary λ to the case $\lambda = 2$, and, in the case of λ odd, to $\lambda = 1$. [1] In particular, the cactus tree model was used for incremental maintenance of the classes of 4-connectivity and of the classes of $(\lambda + 1)$-connectivity [13] and for the edge-connectivity augmentation problems [22, 20, 1].

More general models were suggested by Gabow in [14, 15]. These models represent minimum directed cuts in a directed graph, and have size $O(n^2)$. These representations can be applied to represent minimum cuts of an undirected graph, via replacing every edge by a pair of antiparallel arcs between its endvertices; they can also be converted to the cactus-tree model. However, these models are more complicated than the cactus-tree model, and we are not aware that they directly lead to the reductions mentioned above. An algorithm for the construction of the cactus tree model with the best known time complexity $O(|E|+\lambda^2 n \log(|E|/n))$ is presented in [14] (see also algorithm [19] with complexity $O(\lambda n^2)$).

Another related problem is representation and dynamic maintenance of the set of all minimum S-cuts (i.e., cuts partitioning a subset S of V that have the minimum cardinality among such cuts) and vertex subsets that S is cut into by those cuts. Such a representation, called the connectivity carcass of S, and an efficient algorithm for its incremental maintenance are suggested in [12]. This structure generalizes the cactus tree model (the case $S = V$) and the representation [25] for the minimum cuts between two given vertices (the case $|S| = 2$).

Recently, there has been a growing interest in analysis of near minimum cuts. The following results were obtained for an arbitrary nonnegatively weighted graph. Karger [18] proves that the number of cuts of weight within $\alpha\lambda$ is at most $O(n^{2\alpha})$. This bound was improved for $\alpha < \frac{4}{3}$ to $\binom{n}{2}$ in [24] and for $\frac{4}{3} \leq \alpha < \frac{3}{2}$ to $O(n^2)$ in [17] (all these papers do not provide any representation of such cuts). Benczur in [2] gives a geometric representation of the cuts of weight less than $\frac{6}{5}\lambda$. This model is less compact in comparison with the cactus

[1]The companion paper [9] presents a concept of an "r-skeleton" that generalizes and formalizes these reductions.

tree model: its size can be $\Omega(n^2)$.

In this paper for the case λ odd and in the companion paper [9] for the case λ even, we suggest an extension of the cactus tree model, called the 2-level cactus tree model. Our models represent the system of the λ- and $(\lambda+1)$-cuts and of the vertex classes of $(\lambda+2)$-connectivity of a graph; their sizes are $O(n)$. We give an algorithm that constructs our model in $O(\lambda^2|V|^2)$ time.

For comparison of the range of cut cardinalities covered by our models and by the model of Benczur [2] for multigraphs, observe that: for $\lambda \leq 5$ our model is stronger, since $\lambda+1 \geq \frac{6}{5}\lambda$; in the range $6 \leq \lambda \leq 10$ they represent the same cuts; starting from $\lambda = 11$, when $\lambda+2 < \frac{6}{5}\lambda$, the model of [2] is stronger than our one.

Previous results are as follows: Galil and Italiano [16] and La Poutré et al. [21] suggested a structure for the case $\lambda = 1$, Dinitz [5] and Westbrook [26] suggested another structure for the case $\lambda = 2$. Our 2-level cactus tree models for the cases λ odd and even are not of the same kind: the odd case generalizes the model for the case $\lambda = 1$ and the even case the model for the case $\lambda = 2$.

The model of [16, 21] for 1- and 2-cuts is obtained by shrinking every 3-class of G into a single node; the resulting graph is a cactus tree, possibly with cycles of length two. The 1-cuts of G correspond to the bridges of the model; the 2-cuts are modeled by a pair of edges belonging to the same cycle or by a pair of bridges. Our model for λ- and $(\lambda+1)$-cuts, $\lambda \geq 3$ odd, is almost as simple as this model for 1- and 2-cuts with the following main exception: *not every* pair of bridges of the model corresponds to a $(\lambda+1)$-cut of G, and pairs that do are specified in a compact way. In the case of λ even, our model in [9] is a 2-connected graph, with λ-cuts being modeled by 2-cuts, and $(\lambda+1)$-cuts by 3-cuts of the model.

Similarly to the way that the cactus-tree model represents the λ-cuts of G, we represent the family of λ- and $(\lambda+1)$-cuts of G by a model for the family of bisections of V corresponding to all those cuts. The paper [10] (see also [8]) provides a constructive generic "2-level" approach for modeling bisection families of a set, and gives a simple characterization of bisection families that can be modeled by the minimal cuts of a cactus tree (extending and widely generalizing [7]). This approach is as follows. Two bisections are called parallel if they collectively partition V into 3 parts. Start by choosing a certain subfamily F^{bas} of F consisting of mutually parallel bisections; henceforth, let us call them "basic". Such a family can be always modeled by bridges of a tree. Then the bisections in $F \setminus F^{bas}$ are partitioned into two main groups: those that are parallel to all basic ones are called "local", and "global" otherwise. The local bisections are further decomposed w.r.t. F^{bas} into smaller subgroups, and for each subgroup a so called "local model" is constructed;

[8] shows a way to combine these models with tree model for the basic family into one model that represents all basic and local bisections. The modeling family of this model can be extended to represent also the global bisections, if F^{bas} is chosen such that the partition of V by the basic and local bisections coincides with the partition of V by the whole family F. In general, choosing a basic family that allows "good" modeling of both local and global bisections is nontrivial, or such a family may even not exist.

A bisection in F is called F-separating if it does not cross any other bisection in F. Let F^{sep} denote all members in F that are F-separating. For F being the family of bisections corresponding to mincuts, [7] chooses as the basic family the family F^{sep}. Clearly, this family is parallel, and this choice implies that there are no global bisections. In the case of λ odd, there are also no local bisections. In the case of λ even, the local models are cycles, and combining these with the tree model for the basic family results in the cactus-tree model.

For the family of λ- and $(\lambda+1)$-cuts, all these steps are much more complicated than in the prototype [7]: (i) the choice of an appropriate basic parallel family is not evident; (ii) the local models are far not only cycles, as in [7]; in the case of λ even, their construction goes through reductions, while for λ odd, we derive local models by using the aforementioned characterization [8, 10] of bisection families that can be modeled by a cactus tree; (iii) there exist global bisections, and these are compactly modeled by specific techniques.

As in [16, 21] ($\lambda = 1$) and in [13] ($\lambda = 2$), we use our models for the incremental maintenance of the classes of (λ_0+2)-connectivity, $\lambda_0 \geq 3$, where λ_0 is the connectivity of the initial graph. This means that we support our structure under a sequence of update operations

Insert-Edge(x, y): Insert a new edge between the two given vertices x and y;

and at any time are able to answer the query

Same-(λ_0+2)*-Class*(x, y)*?*: Return "true" if two given vertices x and y belong to the same (λ_0+2)-class of G, and "false" otherwise.

In the case λ odd, our 2-level cactus tree model has a structure similar to the one suggested in [16] for the case $\lambda_0 = 1$. Though there is no immediate reduction of the incremental maintenance problem to the case $\lambda_0 = 1$, we extend for this model the algorithm of [16], preserving the complexity. For an arbitrary sequence of u updates *Insert-Edge* and q queries *Same-*(λ_0+2)*-Class*(x, y)*?*, total time required is $O((u+q+n)\alpha(u+q,n))$, where α is the inverse of the Ackerman function (which grows extremely slow, see, for example, [4]).

This paper is organized as follows. Section 2 brings basic definitions and notations. Section 3 introduces our tools: 2-level cut modeling. In Section 4 we give some properties of λ- and $(\lambda+1)$-cuts. Section 5 deals with both statics and dynamics for the case of odd λ. Section 6 contains concluding remarks.

The preliminary version of this and the companion papers is Extended Abstract [8].

2 Preliminaries and Notations

Let $G=(V,E)$ be an undirected connected (multi)graph with vertex set V and edge set E, where $|V|=n\geq 2$, and $|E|=m$. For any graph H, let $V(H)$ and $E(H)$ denote the vertex and edge sets of H, respectively.

To **shrink** a subset of vertices $S\subseteq V$ means to replace all vertices in S by a single vertex s, to delete all edges with both endvertices in S, and, for every edge with one endvertex in S, to replace this endvertex by s; an edge of a new graph is identified with its corresponding edge of G. For a given partition of V, the **quotient graph** is defined to be the result of shrinking each part into a single node (a **quotient set** of a set is defined similarly).

For $X,Y\subset V$ we denote by $\delta(X,Y)$ the set of edges with one end in X and the other end in Y (clearly, $\delta(X,Y)=\delta(Y,X)$). For brevity, let us use the notations $\bar{X}=V\setminus X$, $\delta(X)=\delta(X,\bar{X})$, $d(X,Y)=|\delta(X,Y)|$, and $d(X)=|\delta(X)|$; $d(X)$ is called the **degree** of X.

A partition of a set into two nonempty parts is called its **bisection**. For a proper subset X of a set U, we denote by $B(X)$ the bisection $\{X,\bar{X}\}$; evidently, $B(X)=B(\bar{X})$. Any bisection $\{X,\bar{X}\}$ of V defines the **edge cut** $C=\delta(X,\bar{X})$; each of $X,\bar{X}$ is called a **side** of C (and, in fact, defines C). The following statement shows that in a connected graph the correspondence between the cuts of G and the bisections of V is bijective (therefore, it is legal to study cuts as vertex bisections).

Proposition 2.1 *For every cut of a connected graph, there is a unique bisection of the vertex set defining it.*

Proof: Let C be a cut of a connected graph $G=(V,E)$ defined by a bisection $\{X,\bar{X}\}$. Deletion of the edges of C partitions G into two nonempty sets of connected components: components in the first set have all vertices in X, and in the other have all vertices in $\bar{X}$. Clearly, there are no edges of G between components belonging to the same set. Thus, shrinking in G every component into a single node results in a connected bipartite graph H, whose parts correspond to the above two sets. Observe that H is defined by C only, independently of any bisection defining it. Thus, if C is defined by another bisection of V, then the connected graph H has two distinct bipartite representations, which is impossible. Indeed, let us fix an arbitrary node v of H; evidently, all the nodes with even distance from v must be on the side containing v, while all the other nodes—with odd distance from v—must be on the other side of any bipartite representation of H. $\square$

A cut C is said to be **minimal** if no its proper subset is a cut. It is well known that $C = \delta(X, \bar{X})$ is a minimal cut of a connected graph G if and only if each of the subgraphs induced by X and $\bar{X}$ is connected. If $|C| = k$ then C is said to be a k-cut; 1-cuts are referred also as **bridges**. The family of all k-cuts of G is denoted by F^k.

We say that a cut $C = \delta(X, \bar{X})$ **divides** a subset S of V (or that C is an $\boldsymbol{S}$**-cut**) if both $X \cap S$ and $\bar{X} \cap S$ are nonempty. We say that a cut divides a subgraph if it divides its vertex set. A subset S of V is called $\boldsymbol{k}$**-connected** if there are no S-cuts of cardinality less than k. The **connectivity** $\lambda(S)$ of a subset S of V is defined to be the maximum k for which S is k-connected (equivalently: $\lambda(S)$ is the minimum number of edges in an S-cut in G). The connectivity λ of G is defined to be $\lambda(V)$. It is easy to see that the relation on vertices "$\{x, y\}$ is k-connected" is an equivalence. Its equivalence classes are called **classes of** $\boldsymbol{k}$**-connectivity**, or, for simplicity, $\boldsymbol{k}$**-classes** (they are often called in literature "k-components"); let $\boldsymbol{n_k}$ denote the number of k-classes. Obviously, the partition of V into $(k+1)$-classes is a subdivision of its partition into k-classes.

For an edge $e = (v, v')$ of a tree, the **branch** that hangs on v via e is the connected component of $T \setminus e$ not containing v. The **bridge-tree** of a graph is the model obtained by shrinking every its 2-class into a single node. We call an ordered sequence of bridges of a graph a **bridge-path** if it forms a path in its bridge-tree.

Following are some definitions concerning bisections and relations between them (see Fig. 2). Two distinct bisections $\{X, \bar{X}\}$ and $\{Y, \bar{Y}\}$ of a set V are called **crossing** if all the four **corner sets** $X \cap Y, X \cap \bar{Y}, \bar{X} \cap Y, \bar{X} \cap \bar{Y}$ are nonempty, and **parallel** otherwise (i.e., if exactly one of these sets is empty). For brevity, we denote these corner sets by A_1, A_2, A_3, A_4, respectively, if no ambiguity arises (see Fig. 2(a)). A bisection defined by a nonempty corner set is called a **corner bisection**. For a pair of crossing bisections $\{X, \bar{X}\}$ and $\{Y, \bar{Y}\}$, the bisection $B((X \cap Y) \cup (\bar{X} \cap \bar{Y}))$ is called their **diagonal bisection**. For simplicity of considerations, we always assume that for parallel bisections $\{X, \bar{X}\}$ and $\{Y, \bar{Y}\}$ the set $A_4 = \bar{X} \cap \bar{Y}$ is empty (as in Fig. 2(b) and Fig. 3(a)), if this does not lead to contradictions. For such two bisections, a bisection $\{Z, \bar{Z}\}$ is said to be **between** them, if $\bar{X} \subset Z$ and $\bar{Y} \subset \bar{Z}$ or $\bar{X} \subset \bar{Z}$ and $\bar{Y} \subset Z$ (see Fig. 2(b)). For a family F of bisections of V, the equivalence classes of the relation "$x, y \in V$, $\{x, y\}$ is not divided by any bisection in F" are called $\boldsymbol{F}$**-atoms**; let $\boldsymbol{n_F}$ denote the number of F-atoms. Note that k-classes are

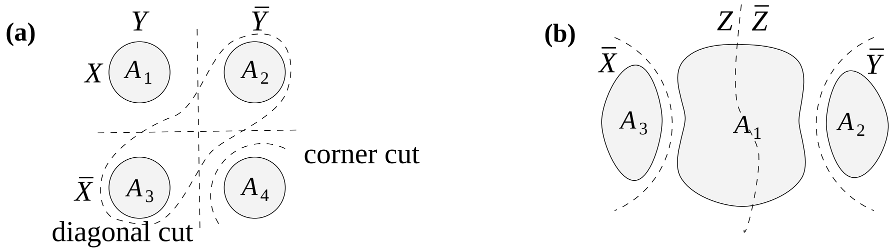

Figure 2: Relations between bisections: (a) crossing bisections; (b) parallel bisections and a bisection between them.

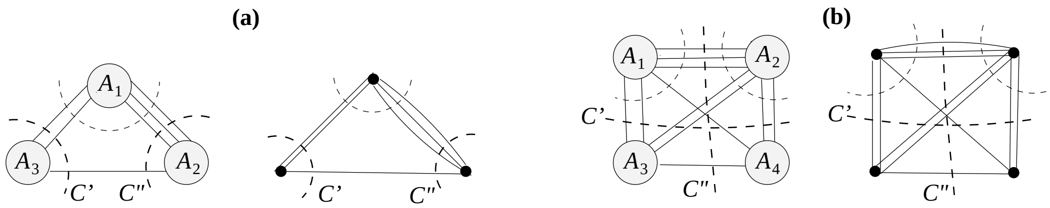

Figure 3: (a) a pair of parallel cuts and their triangle. (b) a pair of crossing cuts and their square.

just F atoms for F being the family of all l-cuts, $l \leq k-1$.

When V is the vertex set of a graph G, similar definitions are used for cuts, considering them as bisections of V (see Fig. 3). For a pair $\{C, C'\}$ of cuts, a cut defined by anyone of their corner bisections is called a **corner cut** (for example see in Fig. 3(b) the cuts defined by A_1 and A_2, of cardinality 6 and 7, respectively). The quotient graph defined by the (nonempty) corner sets is called the **$\{C, C'\}$-square**, in the case C, C' are crossing, and the **$\{C, C'\}$-triangle**, in the case they are parallel. An edge of the square belonging to both C and C' is called a **diagonal edge** (those are the edges in $\delta(A_1, A_4)$ and $\delta(A_2, A_3)$); the other edges of the square are called **side edges**. For brevity, we denote $d_{ij} = d(A_i, A_j)$, $d_i = d(A_i)$ for $i \neq j = 1, \ldots, 4$. Observe that for a pair of parallel cuts

$$d_1 = |C'| + |C''| - 2d_{23}, \tag{1}$$

and for a pair of crossing cuts

$$\begin{aligned} d_1 + d_4 &= |C'| + |C''| - 2d_{23} && (2)\\ d_2 + d_3 &= |C'| + |C''| - 2d_{14}. && (3) \end{aligned}$$

Most of our definitions and results apply to cuts of an integrally weighted graph as well, by replacing the cardinality of a set of edges by the sum of their weights. In fact, in what follows we do not distinguish between a multigraph and its corresponding weighted simple graph if this does not lead to misunderstanding ("the weight of an edge (x, y) is k" means "$d(x, y) = k$", and vice versa). We say that a multigraph is a cycle if its corresponding weighted simple graph is a cycle, and call it ***l*-uniform** if the weight of every edge in the latter is l.

3 Modeling tools??

In this Section, we introduce the hierarchic 2-level approach of [8, 10] to the construction of cut models for families of bisections.

The following concept of a model, applying to cuts of a connected graph as to bisections of its vertex set, has been used in connectivity studies since [7]. Following [10], we present this concept abstractly, for bisections of an arbitrary set (one reason for this decision is to emphasize that edges of original graph play no role in modeling, the other is that illustrating figures are much more clear without such edges).

A **cut model for a set V** (or, for short, a **model**) is a pair $(\mathcal{G}, \psi)$, where $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a connected graph and $\psi : V \longrightarrow \mathcal{V}$ is a mapping;[2] we sometimes abbreviate this notion by $\mathcal{G}$, if ψ is understood. We call ψ a **model mapping** and $\mathcal{G}$ a **structural graph**; vertices of $\mathcal{G}$ are called **nodes** and its edges **structural edges**. A node $\mathcal{N}$ of $\mathcal{V}$ is called **empty** if $\psi^{-1}(\mathcal{N}) = \emptyset$. Observe that, for any cut model, shrinking a subset of nodes of $\mathcal{G}$ implies naturally a new model: its mapping is the composition of the original mapping and the quotient one.

We say that a cut $\mathcal{C} = \delta(\mathcal{X}, \bar{\mathcal{X}})$ of $\mathcal{G}$ **ψ-induces** the bisection $\psi^{-1}(\mathcal{C}) = \{\psi^{-1}(\mathcal{X}), \psi^{-1}(\bar{\mathcal{X}})\}$ of V if both $\psi^{-1}(\mathcal{X}), \psi^{-1}(\bar{\mathcal{X}})$ are nonempty. Any bisection of V that is ψ-induced by a cut of $\mathcal{G}$ is said to be **compatible** with $\mathcal{G}$ (or with $\mathcal{V}$). For a family of cuts $\mathcal{F}$ of $\mathcal{G}$, we denote $\psi^{-1}(\mathcal{F}) = \{\psi^{-1}(\mathcal{C}) : \mathcal{C} \in \mathcal{F}\}$. For a subgraph $\mathcal{G}'$ of $\mathcal{G}$ with node set $\mathcal{V}'$, $\psi^{-1}(\mathcal{G}')$ is defined to be $\psi^{-1}(\mathcal{V}')$.

Let F be a family of bisections of V. Then a triple $(\mathcal{G}, \psi, \mathcal{F})$, where $(\mathcal{G}, \psi)$ is a model for V and $\mathcal{F}$ is a family of cuts of $\mathcal{G}$, is said to be a **cut model for F** if $\psi^{-1}(\mathcal{F}) = F$; then $\mathcal{F}$ is called a **modeling family (for F)** and its members are called **modeling cuts**. For any two models: $(\mathcal{G}, \psi, \mathcal{F})$ for F and $(\mathcal{G}', \psi', \mathcal{F}')$ for $\mathcal{F}$, the triple $(\mathcal{G}', \psi \circ \psi', \mathcal{F}')$ is, clearly, a model for F; it is called the **composition** of the former models.

For short, we say that F is **modeled by a graph $\mathcal{G}$** if there is a cut model, whose structural graph is $\mathcal{G}$ and modeling family is the family of all minimal cuts of $\mathcal{G}$. A cut

[2] In this paper, objects related to a model, which is not a quotient graph, are usually denoted by letters in their calligraphic form, for example $\mathcal{C}, \mathcal{F}, \mathcal{G}, \mathcal{V}(\mathcal{G})$.

model $(\mathcal{G}, \psi, \mathcal{F})$ is called **condensed** if, for every node $\mathcal{N}$ of $\mathcal{G}$, $\psi^{-1}(\mathcal{N})$ is an F-atom or the empty set. It is easy to see that a sufficient condition for a cut model to be condensed is that each $\mathcal{F}$-atom is a singleton.

In order to describe the modeling family in a compact way, we allow an *indirect* description of some of its parts by "bunches" (generalizing [7] and following [10]). For $\mathcal{E}_1, \ldots, \mathcal{E}_l \subseteq \mathcal{E}$, the **bunch generated by $\boldsymbol{\mathcal{E}_1, \ldots, \mathcal{E}_l}$** is $\{\{\varepsilon_1, \ldots, \varepsilon_l\} : \varepsilon_i \in \mathcal{E}_i$ and $\varepsilon_i \neq \varepsilon_j$ for $1 \leq i \neq j \leq l\}$. For example, in the cactus tree model, each bunch is generated by $\mathcal{E}_1, \mathcal{E}_2$, where $\mathcal{E}_1 = \mathcal{E}_2 = \mathcal{L}$ and $\mathcal{L}$ is a cycle of the cactus tree. In our model for λ odd, we have additional type of bunches: each bunch is generated by $\mathcal{E}_1, \mathcal{E}_2$, where $\mathcal{E}_1 = \mathcal{E}_2 = \mathcal{P}$ is a bridge path. Note that the space required to describe a bunch is $|\mathcal{E}_1| + \cdots + |\mathcal{E}_l|$ (or even less if for identical sets we just specify their multiplicity), while the number of elements in a bunch can be $|\mathcal{E}_1| \times \cdots \times |\mathcal{E}_l|$.

The **size** of a model $(\mathcal{G}, \psi, \mathcal{F})$ is the sum of sizes of its three parts: (i) of $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, that is $|\mathcal{V}| + |\mathcal{E}|$, (ii) of ψ, that is $O(|V|)$, and (iii) of the description of $\mathcal{F}$. Observe that describing $\mathcal{F}$ by bunches instead of a trivial listing of all members in $\mathcal{F}$, the size of a model can be much less than the number of bisections (or cuts) in F. Notice that the number of F-atoms n_F can serve, instead of $|V|$, as a natural parameter for measuring the size of a condensed model, excluding the size of the modeling mapping; for simplicity, we say that a model is **linear in $\boldsymbol{n_F}$** if all its parts, except for the model mapping, have size linear in n_F. As an example, the cactus tree model is linear in the number $n_{\lambda+1}$ of $(\lambda+1)$-classes of V, while there can be $\Omega((n_{\lambda+1})^2)$ modeling cuts.

Let us consider an important simple case of a cut model. A family F^p of bisections of V is called **parallel** if its members are pairwise parallel. By [19], $|F^p| = O(|V|)$ (in fact, $|F^p| \leq 2|V| - 3$). Following [7], we represent such a family by the naturally defined **tree model** $(\mathcal{T}^p, \psi^p, \text{1-cuts of } \mathcal{T}^p)$, where $\mathcal{T}^p$ is a tree (see Fig. 4, for a formal definition see [13, 10])). This model is condensed and is bijective, i.e., every bisection in F^p is ψ^p-induced by a unique 1-cut (i.e., by a structural edge) of $\mathcal{T}^p$. For a node $\mathcal{N}$ of $\mathcal{T}^p$, the family of bisections $F^p_{\mathcal{N}} = \{(\psi^p)^{-1}(\varepsilon) : \varepsilon = (\mathcal{N}, \mathcal{N}') \in \mathcal{T}^p\}$ is called the **neighbor group at $\boldsymbol{\mathcal{N}}$** (for example, in Fig. 4, $\{C_1, C_2, C_3, C_4\}$ is a neighbor group at $\mathcal{Z}$). Note that two bisections belong to the same neighbor group if and only if there is no other bisection in F^p between them. Indeed, if $C', C'' \in F^p$ are ψ^p-induced by $\varepsilon', \varepsilon''$, respectively, then $C \in F^p$ is between C' and C'' if and only if the structural edge ψ^p-inducing C belongs to the unique path between ε' and ε'' in $\mathcal{T}^p$; thus no bisection in F^p is between C_1 and C_2 if and only if $\varepsilon', \varepsilon''$ are incident to the same node of $\mathcal{T}^p$.

Given a parallel bisection family, we use the following classification of bisections w.r.t.

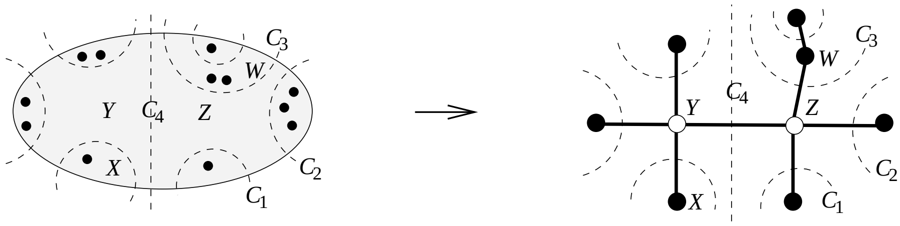

Figure 4: A parallel family and its tree model. (The nodes whose preimages are empty are shown white. Bisections and cuts are shown by dashed lines.)

that family (see Fig. 5(a)). We call that family and its members **basic**, and in what follows denote it by F^{bas} and its tree model by $(\mathcal{T}^{bas}, \psi^{bas}, \mathcal{F}^{bas})$. A nonbasic bisection is called **local** if it is not crossing with any member of F^{bas}, and **global** otherwise.

The local bisections are decomposed relatively to the nodes of $\mathcal{T}^{bas}$ by means of the following model. (see Fig. 5(a,b,c)). The **component** $\hat{V}_{\mathcal{N}}$ at a node $\mathcal{N}$ of $\mathcal{T}^{bas}$ is a quotient set $V_{\mathcal{N}}$ of V (or a quotient graph $\hat{G}_{\mathcal{N}}$ of G, in the case of a cut family F) as follows: for every branch $\mathcal{B}$ hanging at $\mathcal{N}$ in $\mathcal{T}^{bas}$, shrink the subset $(\psi^{bas})^{-1}(\mathcal{B})$ into a single **halo element** (resp., **halo node**); the corresponding quotient mapping is denoted by $\hat{\psi}_{\mathcal{N}}$. The following are some simple properties of components.

Lemma 3.1 ([10]) (i) *Any bisection compatible with a component is either local or basic. Moreover, every local bisection C is compatible with exactly one component, and every basic bisection is compatible with exactly two components (at the endnodes of the structural edge defining it in $\mathcal{T}^{bas}$) and are defined by single halo nodes in these components.*

(ii) *Any two crossing local bisections, as well as their corner bisections, are compatible with the same component.*

Assume now that we are looking for a cut model for a bisection family F of a set V. Let $F^{bas} \subseteq F$, and let F^{loc} and F^{gl} denote the corresponding subfamilies of F of local and global bisections, respectively. (Recall that here and everywhere in this section similar definitions are implicitly made for a cut family F, considering cuts as vertex bisections.) By Lemma 3.1(i), F^{loc} falls into parts $F^{loc}_{\mathcal{N}}$ corresponding to nodes $\mathcal{N}$ of $\mathcal{T}^{bas}$ (via compatibility with $\hat{V}_{\mathcal{N}}$). Following [10], our general approach is to represent the parts $F^{loc}_{\mathcal{N}}$ separately, and then to synthesize the entire representation for $F^{bas} \bigcup F^{loc}$. Let us define the appropriate type of such separate representations. Observe that, for any node $\mathcal{N}$, the neighbor group $F^{bas}_{\mathcal{N}}$ is exactly the set of basic bisections (cuts) compatible with $\hat{V}_{\mathcal{N}}$.

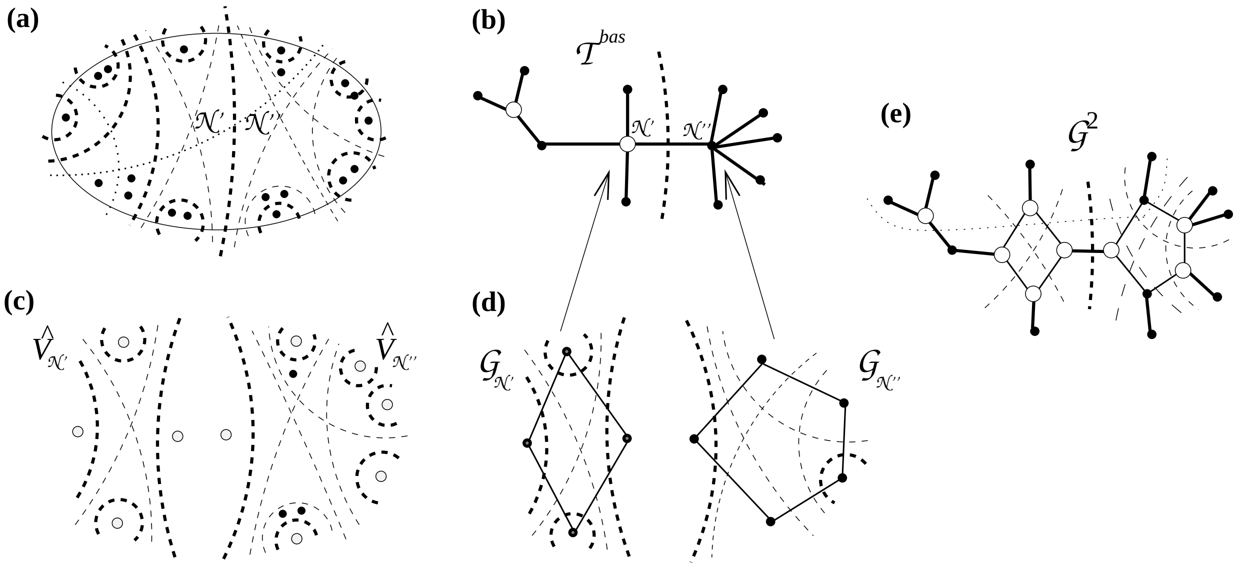

Figure 5: (a) decomposition of a bisection family w.r.t. a basic family (basic bisections are shown by thick dashed lines, local bisections by thin dashed lines, and global bisections by dotted lines); (b) the basic tree; (c) node sets of components at $\mathcal{N}'$ and at $\mathcal{N}''$ (halo nodes are shown gray), and the decomposition of F^{loc}; (d) local models; (e) implanting and the plant model $\mathcal{H}^2$.

Definition 3.2 *Let $\mathcal{N}$ be a node of $\mathcal{T}^{bas}$. A cut model $(\mathcal{G}_\mathcal{N}, \psi_\mathcal{N}, \mathcal{F}'_\mathcal{N})$ for $F'_\mathcal{N}$ is called a* **local model at $\mathcal{N}$** *if the following holds (see Fig. 5(d)):*

(i) $F^{loc}_\mathcal{N} \subseteq F'_\mathcal{N} \subseteq F^{loc}_\mathcal{N} \cup F^{bas}_\mathcal{N}$;

(ii) *for every branch $\mathcal{B}$ of $\mathcal{T}^{bas}$ hanging at $\mathcal{N}$, its preimage $(\psi^{bas})^{-1}(\mathcal{B})$ is mapped by $\psi_\mathcal{N}$ into a single node $\mathcal{N}_\mathcal{B}$ of $\mathcal{G}_\mathcal{N}$.*

Let $F'_\mathcal{N}$ be as in Definition 3.2(i). Let $(\mathcal{G}'_\mathcal{N}, \hat{\psi}'_\mathcal{N}, \mathcal{F}'_\mathcal{N})$ be a cut model for $\hat{\psi}_\mathcal{N}(F'_\mathcal{N})$. It is not hard to verify that the composition $(\mathcal{G}'_\mathcal{N}, \hat{\psi}_\mathcal{N} \circ \hat{\psi}'_\mathcal{N}, \mathcal{F}'_\mathcal{N})$ of $\hat{V}_\mathcal{N}$ and $\mathcal{G}'_\mathcal{N}$ is a local model at $\mathcal{N}$. [3] Hence, in applications to cut families, we can obtain a local model at $\mathcal{N}$ via the component $\hat{G}_\mathcal{N}$ by constructing a cut model for a family $\hat{\psi}_\mathcal{N}(F'_\mathcal{N})$ of cuts of $\hat{G}_\mathcal{N}$.

Assume now that there is given a local model $(\mathcal{G}_\mathcal{N}, \psi_\mathcal{N}, \mathcal{F}_\mathcal{N})$ for each node $\mathcal{N}$ of $\mathcal{T}^{bas}$ with $F^{loc}_\mathcal{N} \neq \emptyset$. Those local models can be naturally "implanted" into $\mathcal{T}^{bas}$ instead of the corresponding nodes to obtain a united cut model $(\mathcal{G}^2, \psi^2, \mathcal{F}^2)$ for $F^{loc} \cup F^{bas}$ as follows (for illustration see Fig. 5(e)).

[3]Moreover, it can be shown (see [10]), that a model $\mathcal{G}_\mathcal{N}$ is a local model at $\mathcal{N}$ if and only if it is a composition of $\hat{V}_\mathcal{N}$ and a cut mode $\mathcal{G}'_\mathcal{N}$ as above.

Let $\mathcal{N}$ be a node of $\mathcal{T}^{bas}$. For any structural edge ε of $\mathcal{T}^{bas}$ incident to $\mathcal{N}$, let $\mathcal{B}^{\varepsilon}$ denote the branch hanging at $\mathcal{N}$ on ε. The structural graph $\mathcal{G}_{\mathcal{N}}$ is implanted into $\mathcal{T}^{bas}$ by replacing the endpoint $\mathcal{N}$ of every structural edge ε incident to $\mathcal{N}$ by the node $\mathcal{N}_{\mathcal{B}^{\varepsilon}} = \psi_{\mathcal{N}}((\psi^{bas})^{-1}(\mathcal{B}^{\varepsilon}))$ of $\mathcal{G}_{\mathcal{N}}$, and then deleting $\mathcal{N}$. The structural graph $\mathcal{G}^2$ is obtained by simultaneously implanting into $\mathcal{T}^{bas}$ each structural graph $\mathcal{G}_{\mathcal{N}}$ instead of the corresponding node $\mathcal{N}$.

We will now define the model mapping φ^2 and the modeling family $\mathcal{F}^2$. Note that, by the construction of $\mathcal{G}^2$, the node set $\mathcal{V}^2$ of $\mathcal{G}^2$ is a disjoint union of the node sets of the local models implanted, and the nodes of $\mathcal{T}^{bas}$ that did not undergo implanting. Similarly, the edge set $\mathcal{E}^2$ of $\mathcal{G}^2$ is a disjoint union of the edge sets of the local models implanted, and the edges of $\mathcal{T}^{bas}$.

The model mapping φ^2 is defined as follows. For $v \in V$, let $\mathcal{N}_v = \psi^{bas}(v)$. Then

$$\varphi^2(v) \equiv \begin{cases} \psi_{\mathcal{N}_v}(v) & \text{if } \mathcal{N}_v \text{ undergo implanting} \\ \mathcal{N}_v & \text{otherwise.} \end{cases}$$

We now define the modeling family $\mathcal{F}^2$. Let us consider the modeling cuts in $\mathcal{F}_{\mathcal{N}}$ and in $\mathcal{F}^{bas}$ as edge sets. By [10], each such edge set is a cut of $\mathcal{G}^2$. The modeling family $\mathcal{F}^2$ is a union of the modeling cuts of the basic tree and of the implanted local models, i.e., $\mathcal{F}^2 = (\bigcup\{\mathcal{F}_{\mathcal{N}} : \text{there is a local model } \mathcal{G}_{\mathcal{N}}\}) \bigcup \mathcal{F}^{bas}$.

Theorem 3.3 ([10]) *$(\mathcal{G}^2, \psi^2, \mathcal{F}^2)$ is a cut model for $F^{bas} \bigcup F^{loc}$, and $(\psi^2)^{-1}$ takes $\mathcal{F}^{bas}$ onto F^{bas}.*

The model $(\mathcal{G}^2, \psi^2, \mathcal{F}^2)$ is called the **plant model** based on $\mathcal{T}^{bas}$ and the set of local models $\{\mathcal{G}_{\mathcal{N}}\}$. Note that a possibility to obtain a set of simple local models, needed for the construction of a simple plant model, depends much on an appropriate choice of a basic parallel family.

Remark: To be able to refer graphs serving as structural for plant models, let us introduce a way to define new types of graphs. We say that a graph is a **tree of graphs of type $\boldsymbol{\mathcal{M}}$** if every its block is a graph of the type $\mathcal{M}$. In this sense, cactus trees may be referred as trees of edges and cycles. Clearly, if all local models are 2-vertex connected and are of a certain type $\mathcal{M}$, then the structural graph of the corresponding plant model is a tree of edges and graphs of type $\mathcal{M}$.

The following Lemma shows several properties that are expanded from local models to their plant model. (Recall that a cut is called minimal if it is inclusion minimal, i.e., if no proper subset of its edges is a cut.)

Lemma 3.4 ([10]) (i) *If all modeling cuts of the local models are minimal, then all modeling cuts of the plant model are minimal; moreover, if all modeling families of the local models are the sets of their minimal cuts, then the modeling family of the plant model is the set of its minimal cuts.*

(ii) *If all the local models are condensed, then the plant model is also condensed.*

(iii) *If each local model at a node $\mathcal{N}$ is of size linear in the number of $(F^{loc}_{\mathcal{N}} \cup F^{bas}_{\mathcal{N}})$-atoms, then the plant model is linear in the number of $(F \cup F^{bas})$-atoms.*

(iv) *Any bisection in F^{loc} is represented in a plant model the same number of times as it was represented in the corresponding local model. Any bisection in F^{bas} is represented exactly once by an edge inherited from $\mathcal{T}^{bas}$, and, in addition, the same number of times as it is represented in the (at most two) local models at the nodes incident to this edge.*

Let us now discuss extension to modeling also F^{gl}. We suggest an approach that allows to keep the structural graph $\mathcal{G}^2$ and extend only the modeling family $\mathcal{F}^2$. For this purpose, let us restrict ourselves to choices of F^{bas} for which no global bisection divides any $(F^{loc} \cup F^{bas})$-atom (in Fig. 5(a), one of the shown global bisections is of this kind, while the other is not). Under this condition, any global bisection is compatible with any model for $(F^{loc} \cup F^{bas})$, and in particular, with the plant model $(\mathcal{G}^2, \psi^2, \mathcal{F}^2)$. Therefore, for any global bisection there exists a cut of $\mathcal{G}^2$ which ψ^2-induces it (notice that such a cut is always nonminimal, see Fig. 5(e)). In such a way, it is possible to model also the global cuts, by extending $\mathcal{F}^2$ with certain nonminimal cuts of $\mathcal{G}^2$, and thus obtain a model for the entire F.

4 Some properties of λ- and $(\lambda + 1)$-cuts

In several proofs, we use the following simple statements (the proof is omitted, for illustration see Fig. 6(a)).

Lemma 4.1 (i) *If the triangle of two parallel λ-cuts has a corner $(\lambda+1)$-cut, then λ is odd and the triangle has the two edges incident to this corner of the weight $\frac{\lambda+1}{2}$ and the third edge of the weight $\frac{\lambda-1}{2}$.*

(ii) *If the triangle of two parallel $(\lambda+1)$-cuts has a corner λ-cut, then λ is even and the triangle has the two edges incident to this corner of the weight $\frac{\lambda}{2}$ and the third edge of the weight $\frac{\lambda}{2}+1$.*

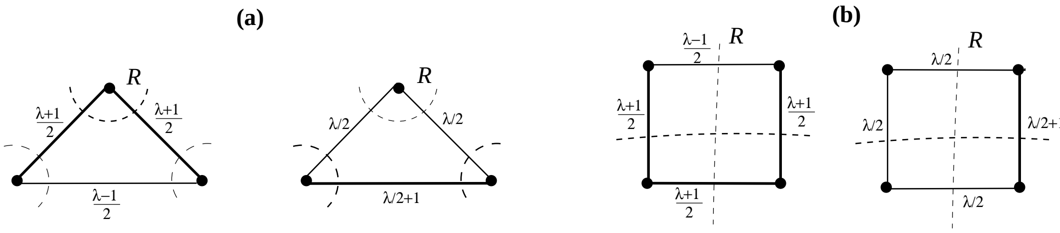

Figure 6: Possible triangles and squares of a λ-cut and a $(\lambda+1)$-cut: (a) triangles; (b) squares.

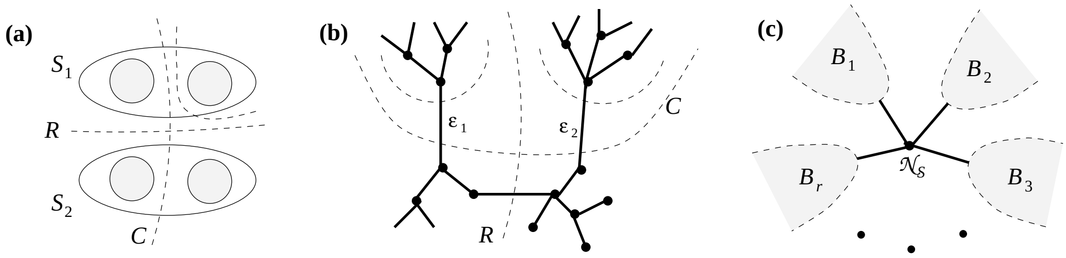

Figure 7: Illustration to the proof of Corollary 4.3.

The possible squares of crossing λ- and $(\lambda+1)$-cuts are given in the following Lemma (see for illustration Fig. 6(b)).

Lemma 4.2 *Let C be a $(\lambda+1)$-cut and R a λ-cut crossing with it. Then the $\{R, C\}$-square has no diagonal edges (so, it is a cycle) and:*

(i) *If λ is odd, then it has one edge of weight $\frac{\lambda-1}{2}$ and three edges of weight $\frac{\lambda+1}{2}$.*

(ii) *If λ is even, then it has one edge of weight $\frac{\lambda}{2}+1$ and three edges of weight $\frac{\lambda}{2}$.*

Proof: Since each corner cut of the $\{R, C\}$-square is of cardinality at least λ, the equations (2) and (3) (see Section 2) imply that $d_{23} = d_{14} = 0$ and that two adjacent corner cuts are λ-cuts, while the other two are $(\lambda+1)$-cuts. Assume, w.l.o.g., that $d_1 = \lambda$. This implies $d_{13} = \lambda - d_{12}$, $d_{24} = d_{12}$ (since $|R| = d_1$), $d_{34} = \lambda - d_{24} + 1$ (since $d_4 = \lambda + 1$), and thus $d_2 = 2d_{12}$. Now, if $d_2 = \lambda + 1$ then $d_{12} = \frac{\lambda+1}{2}$ which implies λ odd and part (i) of the Lemma, and if $d_2 = \lambda$ then $d_{12} = \frac{\lambda}{2}$ which implies λ even and part (ii) of the Lemma. □

Lemma 4.3 *Any $(\lambda+1)$-cut divides at most one $(\lambda+1)$-class.*[4] *More exactly:*

[4]This statement has the following generalization: *Any k-cut divides at most one k-class.*

(i) *If λ is odd, then any $(\lambda+1)$-cut that crosses a λ-cut does not divide any $(\lambda+1)$-class. Moreover, it is φ-induced by a (nonminimal) 2-cut of $\mathcal{T}^\lambda$; thus one of its sides can be partitioned into two sets of degree λ each.*

(ii) *If λ is even, then any $(\lambda+1)$-cut divides exactly one $(\lambda+1)$-class.*

Proof: Assume, in negation, that there is a $(\lambda+1)$-cut C that divides two distinct $(\lambda+1)$-classes, say S_1 and S_2 (see Fig. 7(a)). Let R be any λ-cut separating S_1 from S_2. Observe that R and C are crossing, and that each corner cut of their square divides either S_1 or S_2. By Lemma 4.2, the $\{R, C\}$-square has a corner cut of cardinality λ, contradicting that S_1 and S_2 are $(\lambda+1)$-classes.

We now prove (i) (see Fig. 7(b)). Assume that λ is odd. Let C be a $(\lambda+1)$-cut, and let R be a λ-cut crossing with C. By Lemma 4.2, C is defined by a set which is the union of two disjoint sets of degree λ each. Let ε_1 and ε_2 be the two structural edges of $\mathcal{T}^\lambda$ that correspond to the λ-cuts defined by these sets. It follows that C is φ-induced by the 2-cut $\{\varepsilon_1, \varepsilon_2\}$ of $\mathcal{T}^\lambda$, which finishes the proof of (i).

In order to prove (ii), it is sufficient to show that, for the case λ even, each $(\lambda+1)$-cut divides at least one $(\lambda+1)$-class. Let us first show that if λ is even, then the degree of any $(\lambda+1)$-class S of G is even. Let $\mathcal{B}_1, \mathcal{B}_2 \ldots, \mathcal{B}_r$ be the branches hanging on the corresponding node $\mathcal{N}_S = \varphi(S)$ in $\mathcal{T}^\lambda$ (see Fig. 7(c)). Observe that, by the definition of $\mathcal{T}^\lambda$, for each branch $\mathcal{B}$ hanging on a node $\mathcal{N}$ in $\mathcal{T}^\lambda$ holds $d(\varphi^{-1}(\mathcal{B})) = \lambda$. Let β be the number of edges with endpoints in preimages of distinct branches. Since

$$d(S) = \sum_{i=1}^{r} d(\varphi^{-1}(\mathcal{B}_i)) - 2\beta = r\lambda - 2\beta,$$

$d(S)$ is even.

Suppose now, in negation, that there is a $(\lambda+1)$-cut C which does not divide any $(\lambda+1)$-class. Then C is defined by $\bigcup_{i=1}^{t} S_i$, where $S_i, i = 1, \ldots, t$ are $(\lambda+1)$-classes of G. Let γ be the number of edges with one end in S_i and the other end in S_j, $i, j = 1, \ldots, t$, $i \neq j$. Now

$$|C| = d(\bigcup_{i=1}^{t} S_i) = \sum_{i=1}^{t} d(S_i) - 2\gamma$$

which implies that $|C| = \lambda + 1$ is even, a contradiction. □

5 The odd case

Recall that λ is the cardinality of a minimum cut of the graph G under consideration; thus, F^λ denotes the family of minimum cuts of G. Let the **cube graph** be the graph of the vertices and edges of the 3-dimensional cube. We prove, in a constructive way, the following Theorem (for an example see Fig. 13 at the end of this section):

Theorem 5.1 *In the case $\lambda \geq 3$ odd, for $F^\lambda \cup F^{\lambda+1}$ there exists a condensed cut model $(\mathcal{H}^2, \varphi^2, \mathcal{F}^2)$ of size $O(n)$ (moreover, linear in $n_{\lambda+2}$), with the following properties:*

(i) *The structural graph $\mathcal{H}^2$*

- *for $\lambda > 3$, is a cactus tree,*
- *for $\lambda = 3$, is a tree of edges, cycles, and cube graphs, such that each node of each cube graph is empty and is incident to exactly one bridge;*

(ii) *The modeling family $\mathcal{F}^2$ consists of:*

- *the 1-cuts (bridges),*
- *the minimal 2-cuts (which are all pairs of edges of any block of $\mathcal{H}^2$ that is a cycle),*
- *in the case $\lambda = 3$, for any block of $\mathcal{H}^2$ that is a cube graph, the three cuts consisting each of four its pairwise nonadjacent edges (see the cuts in Fig. 12(e));*
- *for a certain set Π of bridge paths such that any two of them have at most one edge in common, the (nonminimal) 2-cuts $\{\{\varepsilon', \varepsilon''\} : \varepsilon', \varepsilon'' \in \mathcal{P}, \mathcal{P} \in \Pi\}$.*

(iii) *The mapping $(\varphi^2)^{-1}$ takes the set of 1-cuts bijectively onto F^λ (thus the bridge-tree of $\mathcal{H}^2$ is a model isomorphic to $\mathcal{T}^\lambda$) and the set of other cuts in $\mathcal{F}^2$ onto $F^{\lambda+1}$.*

We call any model that satisfies properties (i-iii) in Theorem 5.1 **2-level cactus tree model** (assuming λ odd). Any model for a subfamily of $F^{\lambda+1}$, such that the bridge-tree of its structural graph is a model isomorphic to $\mathcal{T}^\lambda$ and whose modeling family consists of nonminimal 2-cuts as in Theorem 5.1, is called a **Π-model**.

As the basic family for modeling the family $F^\lambda \cup F^{\lambda+1}$, we choose its subfamily $(F^\lambda)^{sep}$ of all F^λ-separating cuts (i.e., the λ-cuts that do not cross any other λ-cut), and henceforth use the notation "local" and "global" cuts w.r.t. this family.

Throughout this section, it is assumed that $\lambda \geq 3$ is odd. Note that then, the family $(F^\lambda)^{sep}$ coincides with F^λ. Let $\mathcal{T}^\lambda$ denote the tree modeling it. Recall that $\varphi^{-1}(\mathcal{N})$, for any nonempty node $\mathcal{N}$ of $\mathcal{T}^\lambda$, is a $(\lambda+1)$-class of G.

A $(\lambda+1)$-cut which is φ-induced by a 2-cut of $\mathcal{T}^\lambda$ is called **degenerate**; by Lemma 4.3(i), the family of degenerate cuts contains the family of global $(\lambda+1)$-cuts. Note that a λ-cut φ-induced by ε crosses a degenerate cut φ-induced by $\{\varepsilon', \varepsilon''\}$ if and only if ε belongs to the path between (but not including) ε' and ε'' in $\mathcal{T}^\lambda$. Thus, a degenerate cut is local if and only if the two structural edges of $\mathcal{T}^\lambda$ defining it are adjacent.

In Sect. 5.1, we show existence of a Π-model, with the structural graph $\mathcal{T}^\lambda$, for the degenerate cuts; recall that those include all global $(\lambda+1)$-cuts. In Sect. 5.2, we show that the nondegenerate and a certain subset of local degenerate $(\lambda+1)$-cuts can be modeled by a cactus tree, for $\lambda > 3$, and by a tree of edges, cycles, and cube graphs, for $\lambda = 3$, in the way described in Theorem 5.1. Further, in Sect. 5.3, we merge the two above models into a 2-level cactus model and show that the constructed model is condensed and is linear in $n_{\lambda+2}$, thus finishing to prove Theorem 5.1. Finally, in Sect. 5.4, we show how, using a 2-level cactus tree model, to maintain efficiently the (λ_0+2)-classes under insertions of edges into G; for this purpose, we extend the algorithm for the case $\lambda_0 = 1$ [16], preserving its complexity.

5.1 Modeling of degenerate cuts

For a 2-cut $\mathcal{C} = \{\varepsilon', \varepsilon''\}$ of $\mathcal{T}^\lambda$, let $\mathcal{P}(\mathcal{C})$ denote the path consisting of $\varepsilon', \varepsilon''$, and the edges of the path between them in $\mathcal{T}^\lambda$.

Lemma 5.2 *Let $\mathcal{C} = \{\varepsilon', \varepsilon''\}$ be a 2-cut of $\mathcal{T}^\lambda$ which φ-induces a $(\lambda+1)$-cut C of G, and let $\mathcal{P}(\mathcal{C}) = \mathcal{N}_1, \mathcal{N}_2, \ldots, \mathcal{N}_r$. Let us remove the edges of $\mathcal{P}(\mathcal{C})$, and denote by $\mathcal{B}_i$ the resulting connected component containing $\mathcal{N}_i$ and by X_i its preimage $\varphi^{-1}(\mathcal{B}_i)$, $i = 1, \ldots, r$. Then the subsets X_i partition V, and the quotient graph $G_\mathcal{C}$ of this partition is a cycle with one edge of weight $\frac{\lambda-1}{2}$ and all the other edges of weight $\frac{\lambda+1}{2}$, where the union of the two parts connected by the edge of weight $\frac{\lambda-1}{2}$ in $G_\mathcal{C}$ defines C.* (For illustration see Fig. 8.)

Proof: In order to prove this Lemma, it is sufficient to show that $d(X_1, X_r) = \frac{\lambda-1}{2}$, $d(X_i, X_{i+1}) = \frac{\lambda+1}{2}$ for $i = 1, \ldots, r-1$, and $d(X_i, X_j) = 0$ otherwise.

Observe that the cut defined by each of X_1 and X_r is a λ-cut, and the cut defined by $X_1 \cup X_r$ is a $(\lambda+1)$-cut. Thus, by Lemma 4.1(i), $d(X_1, X_r) = \frac{\lambda-1}{2}$.

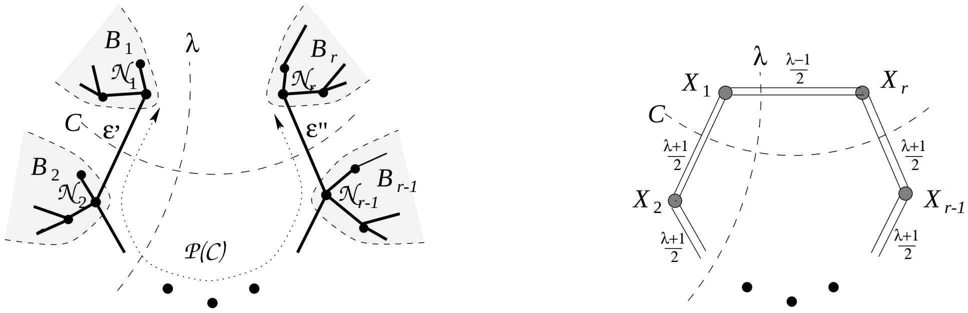

Figure 8: Illustration to the proof of Lemma 5.2.

Suppose in negation that there are i and j, $1 \leq i \leq r-2$, $i+2 \leq j \leq r$, such that $d(X_i, X_j) > 0$. Then we obtain a contradiction, since

$$\lambda < d(X_{i+1}) = d(\bigcup_{l=1}^{i} X_l) + d(\bigcup_{l=i+2}^{r} X_l) - 2d(\bigcup_{l=1}^{i} X_l, \bigcup_{l=i+2}^{r} X_l) \leq$$

$$2\lambda - 2(\frac{\lambda-1}{2} + d(X_i, X_j)) = \lambda + 1 - 2d(X_i, X_j) < \lambda.$$

Thus $d(X_i, X_{i+1}) = \lambda - d(X_1, X_r) = \lambda - \frac{\lambda-1}{2} = \frac{\lambda+1}{2}$, for $i = 1, \ldots, r-1$. $\square$

Corollary 5.3 *Let $\mathcal{C}$ be a 2-cut of $\mathcal{T}^\lambda$ which φ-induces a $(\lambda+1)$-cut of G. Then any two edges of $\mathcal{P}(\mathcal{C})$ form a cut of $\mathcal{T}^\lambda$, which φ-induces a $(\lambda+1)$-cut of G.*

Let us consider the set of paths $\mathcal{P}(\mathcal{C})$, for all 2-cuts $\mathcal{C}$ of $\mathcal{T}^\lambda$ that φ-induce $(\lambda+1)$-cuts of G. We call such a path **generating** if it is inclusion-maximal among those paths.

Lemma 5.4 *Two generating paths have at most one edge in common.*

Proof: Suppose, in negation, that two generating paths $\mathcal{P}'$ and $\mathcal{P}''$ have at least two structural edges in common. Then their intersection is a path of length at least two, say $\mathcal{P}' \cap \mathcal{P}'' = (\mathcal{N}_0, \ldots, \mathcal{N}_q)$, $q \geq 2$.

Let us first show that $\mathcal{P}'$ and $\mathcal{P}''$ are not contained both in a path of $\mathcal{T}^\lambda$. By maximality, $P' \not\subseteq P''$ and $P'' \not\subseteq P'$. Assume, in negation, that $\mathcal{P}' \cup \mathcal{P}''$ is a path and, w.l.o.g., that $\mathcal{N}_q$ is the endnode of $\mathcal{P}'$ and $\mathcal{N}_0$ is the endnode of $\mathcal{P}''$; let $\varepsilon_0 = (\mathcal{N}_0, \mathcal{N}_1)$, $\varepsilon_q = (\mathcal{N}_{q-1}, \mathcal{N}_q)$ (see Fig. 9(a)). We show that the two terminal structural edges $\varepsilon', \varepsilon''$ of $\mathcal{P}' \cup \mathcal{P}''$, where $\varepsilon' \in \mathcal{P}'$, $\varepsilon'' \in \mathcal{P}''$, form a 2-cut of $\mathcal{T}^\lambda$ that φ-induces a $(\lambda+1)$-cut of G, contradicting the maximality of $\mathcal{P}'$ and of $\mathcal{P}''$. Observe that the cuts $C' = \varphi^{-1}(\{\varepsilon', \varepsilon_q\})$ and $C'' = \varphi^{-1}(\{\varepsilon'', \varepsilon_0\})$ form a

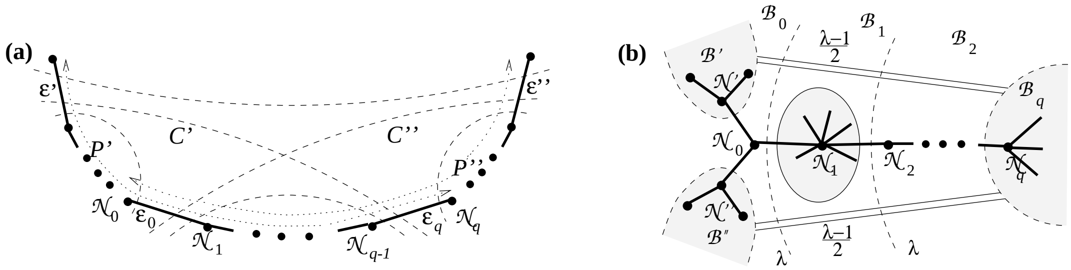

Figure 9: Impossibility of a long intersection of two global paths.

pair of crossing $(\lambda+1)$-cuts of G. One of the corner cuts of their square is φ-induced by the 2-cut $\{\varepsilon', \varepsilon''\}$. By Corollary 5.3, the other three corner cuts are all $(\lambda+1)$-cuts, since each of them is φ-induced by a 2-cut of $\mathcal{T}^\lambda$ that consists of two edges belonging both to $\mathcal{P}'$ or both to $\mathcal{P}''$. One can easily verify that this, together with the equations (2) and (3), imply that all the corner cuts in the C', C''-square are $(\lambda+1)$-cuts, which finishes the current proof.

Now, since $\mathcal{P}'$ and $\mathcal{P}''$ are not contained both in any path of $\mathcal{T}^\lambda$, at least one of $\mathcal{N}_0, \mathcal{N}_q$, say $\mathcal{N}_0$, has two adjacent nodes $\mathcal{N}', \mathcal{N}''$, where $\mathcal{N}' \in \mathcal{P}' \setminus \mathcal{P}''$ and $\mathcal{N}'' \in \mathcal{P}'' \setminus \mathcal{P}'$. Let $\mathcal{B}_q$ be the branch hanging on $\mathcal{N}_{q-1}$ and containing $\mathcal{N}_q$, and let $\mathcal{B}'$ (resp., $\mathcal{B}''$) be the branch hanging on $\mathcal{N}_0$ and containing $\mathcal{N}'$ (resp., $\mathcal{N}''$) (see Fig. 9(b)); these branches are disjoint. Then, by Lemma 4.1,

$$d(\varphi^{-1}(\mathcal{B}_q), \varphi^{-1}(\mathcal{B}')) = d(\varphi^{-1}(\mathcal{B}_q), \varphi^{-1}(\mathcal{B}'')) = \frac{\lambda - 1}{2}. \tag{4}$$

Let $X_i = \varphi^{-1}(\mathcal{B}_i)$, where $\mathcal{B}_i$ is the connected component containing $\mathcal{N}_i$ in $\mathcal{T}^\lambda \setminus \{(\mathcal{N}_0, \mathcal{N}_1), (\mathcal{N}_1, \mathcal{N}_2)\}$, $i = 0, 1, 2$. Observe that $\{X_0, X_1, X_2\}$ is a partition of V, and that $d(X_0) = d(X_2) = \lambda$. By (4), $d(X_0, X_2) \geq \lambda - 1$, and thus we have

$$d(X_1) = d(X_0) + d(X_2) - 2d(X_0, X_2) = 2\lambda - 2d(X_0, X_2) \leq 2,$$

a contradiction since $\lambda \geq 3$. □

By this Lemma and Corollary 5.3, the set of generating paths is a Π-model.

Remark: The suggested representation of the degenerate cuts is almost bijective. The only case of a double representation can occur for two disjoint pairs of edges incident to an empty node of degree 4 in $\mathcal{T}^\lambda$. Recall that for λ odd, no two λ-cuts cross. Thus, if $C = \delta(X, \bar{X})$ is a degenerate $(\lambda+1)$-cut, and X has a partition into two sets X_1, X_2 of degree λ each, then this partition is unique. Thus, a double representation can occur only if also $\bar{X}$ has also a (unique) partition into two sets X_1', X_2' of degree λ each. Simple computations show that the quotient graph of the partition $\{X_1, X_2, X_1', X_2'\}$ is a square with weights as follows:

each diagonal edge of weight $1 < \alpha \le \frac{\lambda-1}{2}$, two opposite side edges of weight $\frac{\lambda-1}{2}$, while the other two side edges of weight $\frac{\lambda+1}{2} - \alpha$. In particular, there is no λ-cut between any two of the four corners. Thus the corner cuts form a neighbor group which corresponds to an empty node of $\mathcal{T}^\lambda$ of degree 4.

We now establish linearity of this Π-model in $n_{\lambda+2}$. Recall that the size of $\mathcal{T}^\lambda$ is $O(n_{\lambda+1}) = O(n_{\lambda+2})$; therefore, it is sufficient to prove linearity in $|\mathcal{V}(\mathcal{T}^\lambda)|$.

Lemma 5.5 *A structural edge belongs to at most four paths of* Π *if* $\lambda > 3$*, and to at most three paths of* Π *if* $\lambda = 3$.

Proof: Let ε be a structural edge of $\mathcal{T}^\lambda$, and let R the λ-cut of G φ-induced by ε. Let $P = (\mathcal{N}_0, \mathcal{N}_1, \ldots, \mathcal{N}_q)$ be a path in Π containing ε, and let C_i be the cut φ-induced by $(\mathcal{N}_{i-1}, \mathcal{N}_i)$, $i = 1, \ldots, q$. By Lemma 4.1(i), $|C_1 \cap C_q| = \frac{\lambda-1}{2}$. Observe that, for any $i = 1, \ldots, q$, $C_1 \cap C_q \subseteq C_i$. In particular, the set $C_1 \cap C_q$ of $\frac{\lambda-1}{2}$ edges is contained in R. By the maximality of the paths in Π, for two distinct paths in Π such sets are distinct (but not necessarily disjoint). For $\lambda = 3$ holds $\frac{\lambda-1}{2} = 1$, which implies that those sets are disjoint. Hence, if $\lambda = 3$ then there can be at most three paths in Π containing ε.

Assume now that $\lambda > 3$, and let $\varepsilon = (\mathcal{N}', \mathcal{N}'')$. Since any two paths in Π containing ε have no other structural edge in common, it is sufficient to show that there are at most 4 structural edges incident to ε such that each of them belongs to some path of Π containing ε. For this purpose, we prove that each of $\mathcal{N}', \mathcal{N}''$ is incident to at most two such edges. We give a proof for $\mathcal{N}''$ (the proof for $\mathcal{N}'$ is similar).

Suppose, in negation, that this is not so. Then there are three structural edges, say, $\varepsilon_i = (\mathcal{N}'', \mathcal{N}_i)$, $i = 1, 2, 3$, such that each cut $\{\varepsilon, \varepsilon_i\}$, $i = 1, 2, 3$, φ-induces a $(\lambda+1)$-cut of G. Let $\mathcal{T}', \mathcal{T}'', \mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3$ be the connected components of $\mathcal{T}^\lambda \setminus \{\varepsilon, \varepsilon_1, \varepsilon_2, \varepsilon_3\}$, where $\mathcal{N}' \in \mathcal{T}'$, $\mathcal{N}'' \in \mathcal{T}''$, and $\mathcal{N}_i \in \mathcal{T}_i$, $i = 1, 2, 3$ (see Fig. 10). Since $d(\varphi^{-1}(\mathcal{T}')) = d(\varphi^{-1}(\mathcal{T}_i)) = \lambda$ and $d(\varphi^{-1}(\mathcal{T}' \cup \mathcal{T}_i)) = \lambda + 1$, then, by Lemma 4.1(i), $d(\varphi^{-1}(\mathcal{T}'), \varphi^{-1}(\mathcal{T}_i)) = \frac{\lambda-1}{2}$, $i = 1, 2, 3$. But then

$$\lambda = d(\varphi^{-1}(\mathcal{T}')) \ge \sum_{i=1}^{3} d(\varphi^{-1}(\mathcal{T}'), \varphi^{-1}(\mathcal{T}_i)) = 3\frac{\lambda - 1}{2},$$

a contradiction, since $\lambda > 3$. $\square$

The examples in Fig. 10(a,b) show that the bounds in Lemma 5.5 are tight.

Corollary 5.6 *The total length of the paths in* Π *is at most* $4(|\mathcal{V}(\mathcal{T}^\lambda)| - 1)$*; hence, their number is at most* $2(|\mathcal{V}(\mathcal{T}^\lambda)| - 1)$.

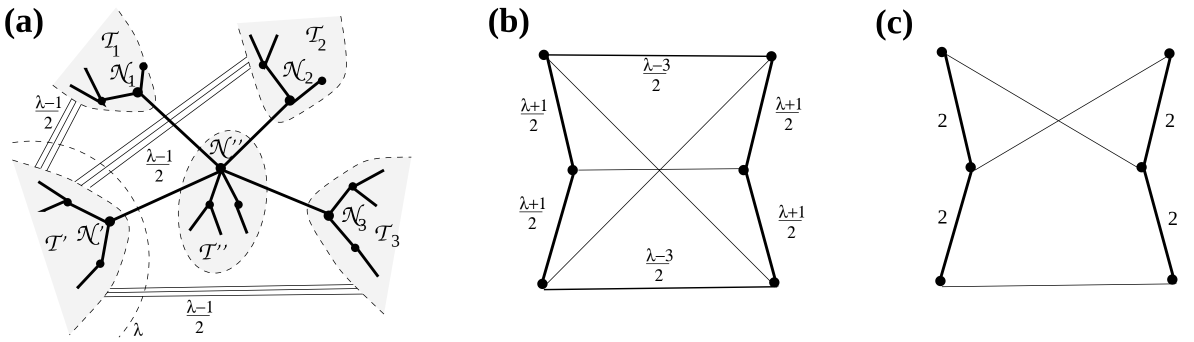

Figure 10: (a) Illustration to the proof of Lemma 5.5; (b,c) examples showing that the bounds in Lemma 5.5 are tight: (b) $\lambda > 3$, (c) $\lambda = 3$.

5.2 Modeling of nondegenerate cuts

We now turn to study nondegenerate $(\lambda+1)$-cuts. Let us call a $(\lambda+1)$-cut **essential** if it is nondegenerate or is a degenerate corner cut in the square of two crossing nondegenerate $(\lambda+1)$-cuts. Although the latter are already modeled in the Π-model, we add them to the nondegenerate cuts, since, as we show, the obtained family of essential cuts is modeled by a cactus tree if $\lambda > 3$, and by only a slightly more complicated model if $\lambda = 3$. The family of nondegenerate cuts only, does not have, in general, such models.

Note that any essential cut is local. Indeed, by Lemma 4.3(i), any nondegenerate cut is local; thus, by Lemma 3.1, any corner cut in a square of a pair of crossing nondegenerate cuts is also local. Clearly, any model for the essential cuts plus a Π-model for the degenerate cuts represent, together, all the $(\lambda+1)$-cuts. Let us denote by $\bar{F}$ the family of essential $(\lambda+1)$-cuts. Observe that the $(\bar{F} \cup F^\lambda)$-atoms are just the $(\lambda+2)$-classes of G, since no degenerate cut divides a $(\lambda+1)$-class.

In what follows, we construct a model for $\bar{F} \cup F^\lambda$ as a plant model based on $\mathcal{T}^\lambda$. The components w.r.t. $\mathcal{T}^\lambda$ are called $\boldsymbol{(\lambda+1)}$**-components** of G (this is a generalization of the concept of a 3-component suggested in [13] for $\lambda = 2$). The following important statement is straightforward.

Lemma 5.7 (i) *Any $(\lambda+1)$-component is a λ-connected graph.*

(ii) *The λ-cuts of a $(\lambda+1)$-component are cuts defined by a single halo node, and vice versa.*

Let $\mathcal{N}$ be an arbitrary node of $\mathcal{T}^\lambda$, and let $(\hat{G}_\mathcal{N}, \hat{\psi}_\mathcal{N})$ be the corresponding $(\lambda+1)$-component. Let $\bar{F}_\mathcal{N}$ be the family of essential cuts compatible with $\hat{G}_\mathcal{N}$. Recall that to

Figure 11: Conditions (C1), (C2) for a pair of crossing bisections.

obtain a local model at $\mathcal{N}$, it is sufficient to construct a cut model for the family $\hat{\psi}_{\mathcal{N}}(\bar{F}_{\mathcal{N}})$ of cuts of $\hat{G}_{\mathcal{N}}$. It is not hard to see, that this family coincides with the family of essential cuts of $\hat{G}_{\mathcal{N}}$. Henceforth, let us study this family, denoting it for simplicity by $\bar{F}_{\mathcal{N}}$, instead of $\hat{\psi}_{\mathcal{N}}(\bar{F}_{\mathcal{N}})$. In this section, we show that any such family can be modeled by a cactus tree (i.e., by the minimal cuts of a cactus tree), except for one specific case of a $(\lambda+1)$-component, which can occur for $\lambda = 3$ only. For this purpose, we use the following characterization (see Fig. 11).

Theorem 5.8 ([10]) *A bisection family F is modeled by a cactus tree if and only if for any pair of crossing bisections in F holds:*

(C1) *the four corner bisections are in F;*

(C2) *the diagonal bisection is not in F.*

Moreover, for any family satisfying conditions (C1), (C2), *there exists a cactus tree model of size linear in the number of F-atoms, which represents each bisection at most twice.*

Note that since any nondegenerate cut is local, then, by Lemma 3.1(ii), any two crossing nondegenerate cuts are compatible with the same component. In order to establish the cactus tree structure of the nondegenerate cuts, the following analog of Crossing Lemma [3, 7] is crucial.

Lemma 5.9 *Let C', C'' be a pair of crossing nondegenerate $(\lambda+1)$-cuts. Then all the corner cuts of their square are $(\lambda+1)$-cuts; moreover, the $\{C', C''\}$-square is a uniform $\frac{\lambda+1}{2}$-cycle.*

Proof: We first show that all corner cuts in the $\{C', C''\}$-square are $(\lambda+1)$-cuts. Assume, in negation, that this is not so. Then the equations (2) and (3) imply that at least one corner cut is a λ-cut; w.l.o.g. assume $d_1 = \lambda$. Since both C', C'' are nondegenerate, it follows that $d_2, d_3 \geq \lambda+1$, which together with (3) implies $d_2 = d_3 = \lambda+1$. But then the triangle of $\delta(A_2)$ and $\delta(A_2 \cup A_1)$ has two corner $(\lambda+1)$-cuts and one corner λ-cut, which by Lemma 4.1(ii) is impossible. Now, a simple computation using (2) and (3) shows that the $\{C', C''\}$-square has edge weights as required. $\square$

As a corollary, for any pair of crossing nondegenerate cuts in $\bar{F}$ holds both (C1) and (C2). Now, the question remains only for squares of crossing cuts in $\bar{F}$ such that at least one of them is degenerate. We show that such crossing pairs do not exist, excluding one specific case of a $(\lambda+1)$-component. Observe that, by Lemma 5.7(ii), a $(\lambda+1)$-cut of a $(\lambda+1)$-component is degenerate if and only if it is defined by two halo nodes.

Lemma 5.10 *Let $\bar{C}$ be an essential degenerate cut. Then $\bar{C}$ does not cross any essential cut, except for the case when $\lambda = 3$ and the $(\lambda+1)$-component that $\bar{C}$ is compatible with is a cube graph on halo nodes only.*

Proof: Let C', C'' be a pair of crossing nondegenerate $(\lambda+1)$-cuts such that $\bar{C}$ is a degenerate corner cut in their square. Consider the $\{C', C''\}$-square. W.l.o.g., assume that A_1, A_2 are contained in the same side of C' and that $\bar{C}$ is defined by A_1; hence, A_1 consists of two nodes of the degree λ each, say X_1, Y_1. Assume to the contrary that there is an essential $(\lambda+1)$-cut C crossing with $\bar{C}$ (see Fig. 12(a)). Clearly, C is crossing with at least one of C', C''.

Let us consider the case of nondegenerate C first. We show that, in this case, C divides each of A_2 and A_3 (see Fig. 12(b)(c)). Assume, in negation, that C does not divide A_2 (the proof for A_3 is similar). Then one of X_1, Y_1, say X_1, and A_2 are contained both in the same part of C. Now, if C is crossing with C' (resp., C''), then Y_1 (resp., X_1) is a corner cut in their square. In each case, a square of two nondegenerate $(\lambda+1)$-cuts has a corner λ-cut, a contradiction to Lemma 5.9.

So, C divides each of A_2, A_3 in addition to A_1 (see Fig. 12(d)). This implies that there are at least seven $\{C, C', C''\}$-atoms. From [24, Lemma 7.3], it follows that if F is a family consisting of three cuts and if there are more than six F-atoms, then F contains a cut of cardinality at least $\frac{4}{3}\lambda$. Applying this to $F = \{C, C', C''\}$ we obtain $\frac{\lambda+1}{\lambda} \geq \frac{4}{3}$, which is possible only if $\lambda = 3$.

Let us now fix $\lambda = 3$ and show that each of the cuts defined by A_2 and A_3 is degenerate. We prove this for A_2 (the proof for A_3 is similar). Observe that C crosses both C' and C''. Denote the parts into which C divides A_i, $i = 1, 2, 3$, by X_i, Y_i, where all X_i are in the same part of C (see Fig. 12(d)). By Lemma 5.9, each of the cuts $\delta(X_1 \cup X_2)$ and $\delta(Y_1 \cup Y_2)$ is a $(\lambda+1)$-cut, since it is a corner cut in the $\{C, C'\}$-square and since both C, C' are

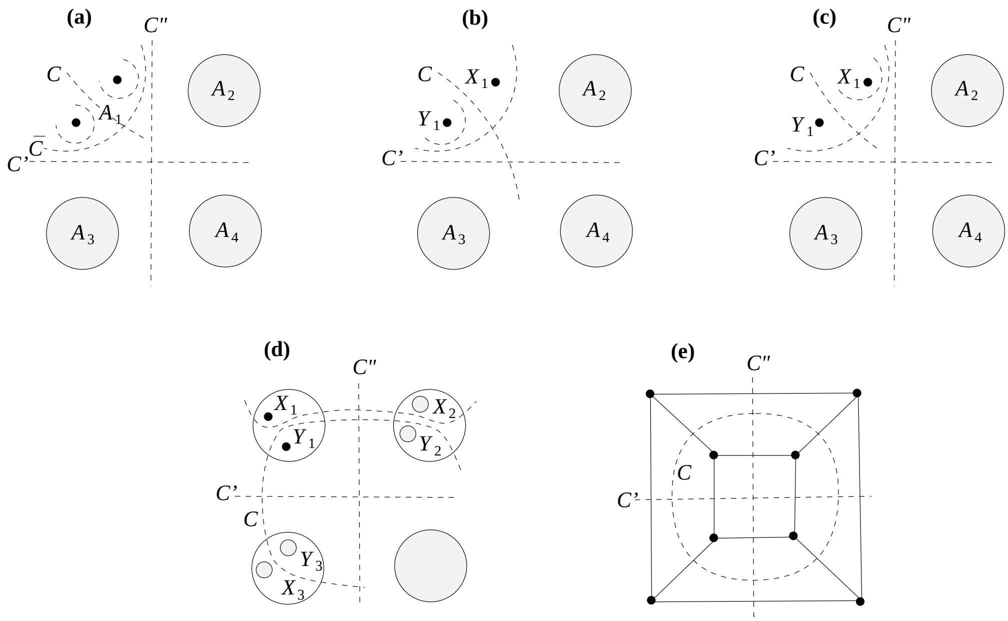

Figure 12: Illustration to the proof of Lemma 5.10.

nondegenerate. Now, $\delta(X_1)$ is a corner λ-cut in the $\{C'', \delta(X_1 \cup X_2)\}$-square, $\delta(Y_1)$ is a corner λ-cut in the $\{C', \delta(Y_1 \cup Y_2)\}$-square, and C' is nondegenerate. Hence, by Lemma 5.9, each of $\delta(X_1 \cup X_2)$, $\delta(Y_1 \cup Y_2)$ must be degenerate. Therefore, $d(X_2) = d(Y_2) = \lambda = 3$, which implies that $\delta(A_2)$ is degenerate, and thus X_2, Y_2 are singletons.

Now, using this and arguments as before applied to A_2 instead of A_1, we obtain that C divides A_4 into two sets, say X_4, Y_4, of the degree λ each. Summarizing, we have $d(X_i) = d(Y_i) = \lambda = 3$, $i = 1, \ldots, 4$; by Lemma 5.7(ii), all X_i, Y_i are single halo nodes. By Lemma 4.1(i), two nodes defining a degenerate 4-cut must be connected by exactly one edge. Using this, one can easily verify that the $(\lambda + 1)$-component is a cube graph (see Fig. 12(e)).

Now, let us consider the case of λ arbitrary and a degenerate essential cut C. Since C is crossing with C' or with C'', which are both nondegenerate cuts, we arrive at the same situation as before of a nondegenerate cut, which is now C, crossing a degenerate essential cut. Thus we get the statement of the theorem for C, i.e., $\lambda = 3$, and the $(\lambda+1)$-component that C is compatible with is a cube graph on halo nodes only. The same holds for $\bar{C}$, since, by Lemma 3.1(ii), C and $\bar{C}$, being crossing, are compatible with the same component. □

By this Lemma and Theorem 5.8, we obtain:

Corollary 5.11 *Let $\mathcal{N}$ be a node of $\mathcal{T}^\lambda$ with $\bar{F}_\mathcal{N} \neq \emptyset$. Then either*

(i) *for the family $\bar{F}_\mathcal{N}$ there exists a local model (w.r.t. $\mathcal{T}^\lambda$), which is a cactus tree with its minimal cuts, of size linear in the number of $\bar{F}_\mathcal{N}$-atoms, or;*

(ii) $\lambda = 3$, *and $\hat{G}_\mathcal{N}$ is a cube graph with halo nodes only.*

5.3 United model

Let us show how to construct a 2-level cactus tree model using the models described in the two previous sections. Note that a cube graph has only three nondegenerate 4-cuts; each such cut consists of four pairwise nonadjacent edges of the cube. Thus, if $\hat{G}_\mathcal{N}$ is a cube graph with halo nodes only, then $\hat{G}_\mathcal{N}$, with its three nondegenerate cuts, is a local model for the nondegenerate cuts in $F_\mathcal{N}$; the degenerate cuts in $F_\mathcal{N}$, although they are all essential, can be skipped, since they are represented in the Π-model. To accord conditions of Theorem 5.1, let us do the following. For every cactus tree type local model as in Corollary 5.11, replace every its bridge by two parallel edges with the same ends; clearly, such operation does not spoil the model and turns the structural graph be 2-connected. Let us now implant all the local models as in Corollary 5.11 instead of the corresponding nodes of $\mathcal{T}^\lambda$, as described in Section 3; we denote the resulting cut model for $\bar{F} \cup F^\lambda$ by $(\mathcal{H}^2, \varphi^2, \bar{\mathcal{F}}^2)$. Observe that this model bijectively represents the λ-cuts by its 1-cuts, which all are inherited from $\mathcal{T}^\lambda$. In other words, the bridge-tree of $\mathcal{H}^2$, obtained by shrinking every its 2-class, is isomorphic to $\mathcal{T}^\lambda$. By Corollaries 5.6 and 5.11 and Lemma 3.4(iii), the size of the model is linear in the number of the $\bar{F}$-atoms, i.e., the $(\lambda+2)$-classes of G. The structural graph $\mathcal{H}^2$ is a cactus tree in the case $\lambda > 3$, or a tree of edges, cycles, and cube graphs, in the case $\lambda = 3$.

Identifying each edge of $\mathcal{T}^\lambda$ with the corresponding bridge of $\mathcal{H}^2$, the set of generating paths forms a set of bridge-paths of $\mathcal{H}^2$, called **generating bridge-paths**. Let us add to $\bar{\mathcal{F}}^2$ the cuts corresponding to the generating bridge-paths, and denote by $\mathcal{F}^2$ the obtained modeling family. By Corollary 5.6, this operation retains the linearity of the size in the number of the $(\lambda+2)$-classes of G. It is easy to see that the model $\mathcal{H}^2 = (\mathcal{H}^2, \varphi^2, \mathcal{F}^2)$, together with the set of the generating bridge-paths (which is a Π-model), satisfies the conditions (i-iii) of Theorem 5.1, i.e., is a 2-level cactus tree model for $F^\lambda \cup F^{\lambda+1}$.

Note that every local model implanted is condensed, since its modeling cuts partition the node set of the structural graph into singletons; this is clear for a cactus tree, and can be easily verified for a cube. Thus, by Lemma 3.4(ii), the model $\mathcal{H}^2$ is also condensed. This finishes the proof of Theorem 5.1.

Observe that any cut in $F^{\lambda+1}$ is represented by $\mathcal{H}^2$ at most four times, since anyone of

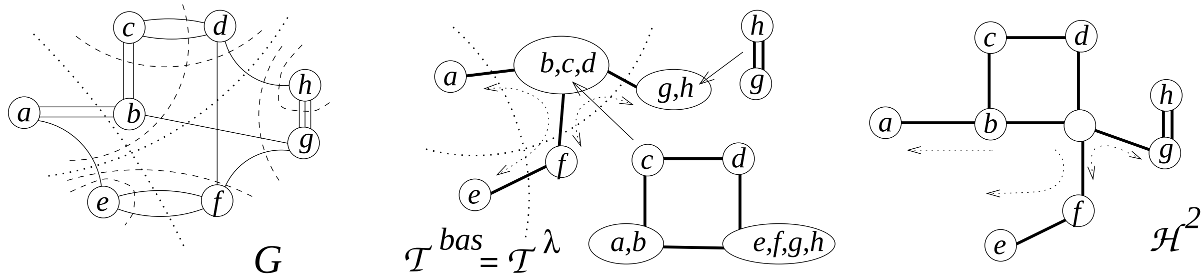

Figure 13: Example of the construction of a 2-level cactus tree model, the case $\lambda = 3$. (Generating paths and some of generating cuts are shown by dotted lines.)

the Π-model and the cactus tree model represent cuts at most twice, while cuts modeled by cubes and cuts in F^{λ} are represented bijectively.

Remark. Lemmas 5.2 and 5.9 imply the following interesting property of paths in a Π-model (we do not use it, and thus omit the proof): *Let $\varepsilon', \varepsilon''$ be any two edges in a path in Π. Then the shortest path between them in $\mathcal{H}^2$ has at most one structural edge in common with any cycle of $\mathcal{H}^2$.*

5.4 Incremental maintenance

Herein we show how, using a 2-level cactus tree model, to maintain efficiently the $(\lambda + 2)$-classes under insertions of edges into G.

Theorem 5.12 *For λ_0 odd, the $(\lambda_0 + 2)$-classes of G can be maintained under a sequence of u updates* Insert-Edge *and q queries* Same-$(\lambda_0 + 2)$-Class(x, y)? *in $O((u+q+n)\alpha(u+q, n))$ total time. The initialization time is polynomial in n, and the space required is $O(n)$.*

Our main idea is to extend the algorithm for the case $\lambda = 1$ [16] to the case of arbitrary λ odd, preserving its complexity. For $\lambda > 3$, our model is structurally similar to that of [16], except for the Π-model, and this goal is achieved by replacing the insertion of any edge by equivalent, in a sense, sequence of edge insertions; each of those insertions is processed separately by means of the algorithm [16]. Moreover, we show how our algorithm can be extended to the case $\lambda = 3$, when also cube graphs appear in the structural graph.

Notice that we can easily reduce the complexity of maintenance to $O((u+q+n_{\lambda_0+2})\alpha(u+q, n_{\lambda_0+2}))$ in the following way. At the preprocessing stage, we can build the quotient graph G' by shrinking each of the n_{λ_0+2} $(\lambda_0 + 2)$-classes of G into a single supervertex and apply our algorithm to G', with n_{λ_0+2} supervertices, instead of G. In this version, the current $(\lambda_0 + 2)$-class of a vertex v of G is found as the current $(\lambda_0 + 2)$-class of the supervertex

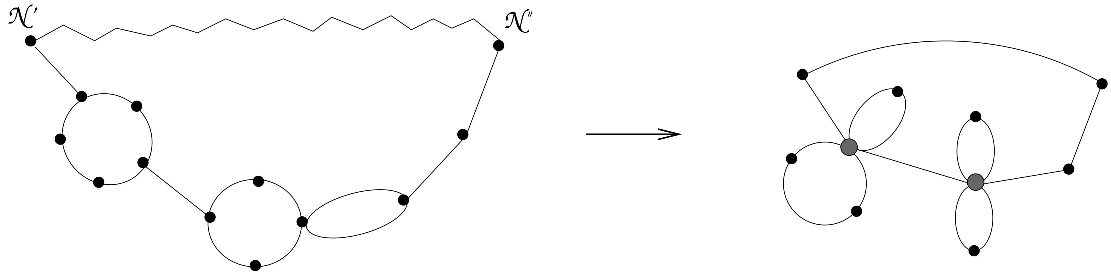

Figure 14: *Squeeze*($\mathcal{N}'$,$\mathcal{N}''$) operation. (New nodes arising from shrinkings are shown by bigger circles.)

of G' corresponding to the initial $(\lambda_0 + 2)$-class of v. This is done via two queries, where finding the supervertex can be supported by a static data structure in $O(1)$ time.

5.4.1 Incremental transformations, the case $\lambda > 3$

Let us consider the case $\lambda > 3$, when the structural graph $\mathcal{H}^2$ is a cactus tree. We first discuss the transformations of a 2-level cactus tree model caused by insertion of an edge (x, y) into G. Similarly to [16, 13, 10], when an edge (x, y) is inserted, we define the **squeezing path** $\mathcal{P}_{xy}$ to be the unique path-of-edges-and-cycles in $\mathcal{H}^2$ between the nodes $\varphi^2(x)$ and $\varphi^2(y)$. The **attachment nodes** of $\mathcal{P}_{xy}$ are defined as $\varphi^2(x)$, $\varphi^2(y)$, and all the articulation nodes of $\mathcal{H}^2$ separating them. Each cycle in $\mathcal{P}_{xy}$ contains exactly two attachment nodes. To **squeeze a cycle** means to shrink its two attachment nodes into a new node (recall that shrinking includes deletion of arising loops). As usually, we assume that shrinking a set of nodes of a structural graph implies the corresponding change of the model mapping: the new mapping is the composition of the original mapping and the quotient one.

Given two nodes ($\mathcal{N}'$,$\mathcal{N}''$) of a model whose structural graph is a cactus tree, we define the following operation (see Fig. 14):

Squeeze($\mathcal{N}'$,$\mathcal{N}''$)**:**

- add an edge with endnodes $\mathcal{N}'$,$\mathcal{N}''$;
- squeeze every cycle on the path between $\mathcal{N}'$ and $\mathcal{N}''$.

It is easy to show that such an operation keeps the property of the structural graph to be a cactus tree; this operation is the same as used in [16].

Let us denote by $\hat{\mathcal{P}}_{xy}$ the bridge-path formed by the bridges of $\mathcal{P}_{xy}$. The set of intersections $\{\hat{\mathcal{P}}_{xy} \cap \mathcal{P} : \mathcal{P} \in \Pi\}$ which have at least two edges and the single edges of $\hat{\mathcal{P}}_{xy}$ not

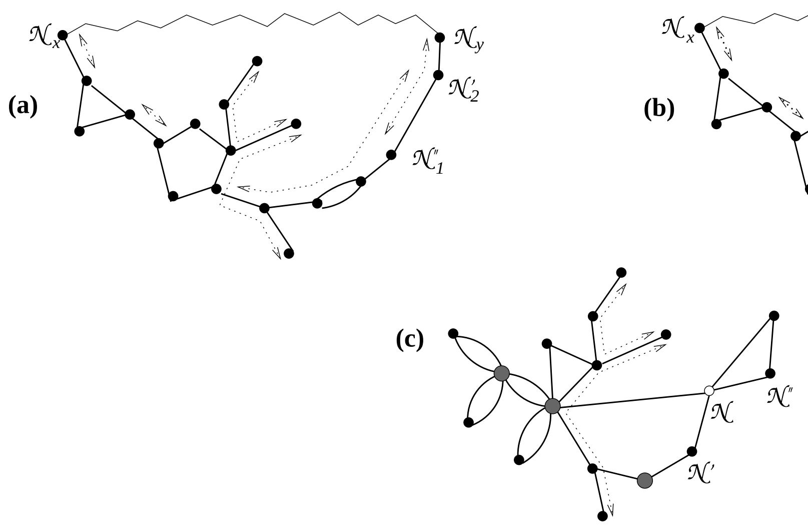

Figure 15: Transformations of a 2-level cactus tree model: (a) initial model; (b) transformation (A); (c) transformations (B),(C). (Dotted lines show generating paths and cycle-generating subpaths.)

belonging to any such intersection are called the **cycle-generating subpaths** of $\mathcal{P}_{xy}$ (or of $\hat{\mathcal{P}}_{xy}$). The transformations of the 2-level cactus tree model under insertion of an edge $e = (x, y)$ are as follows (see Fig. 15):

(A) For any two cycle-generating subpaths $\mathcal{P}' = \mathcal{N}'_1, \ldots, \mathcal{N}'_2$ and $\mathcal{P}'' = \mathcal{N}''_1, \ldots, \mathcal{N}''_2$, that have a structural edge $(\mathcal{N}''_1, \mathcal{N}'_2)$ in common:

(i) an empty node $\mathcal{N}$ is inserted into this common edge (i.e., this edge is replaced by $\mathcal{N}$ and the two structural edges $(\mathcal{N}''_1, \mathcal{N})$ and $(\mathcal{N}, \mathcal{N}'_2)$);

(ii) $\mathcal{P}'$ and $\mathcal{P}''$ are replaced by the two new cycle-generating subpaths $\mathcal{N}'_1, \ldots, \mathcal{N}$ and $\mathcal{N}, \ldots, \mathcal{N}''_2$, respectively.

(Observe that after this modification, no two cycle-generating subpaths have a structural edge in common.)

(B) The endnodes of the cycle-generating subpaths partition $\mathcal{P}_{xy}$ into a sequence of paths-of-edges-and-cycles. For each such path with endnodes $\mathcal{N}', \mathcal{N}''$, the operation $\mathit{Squeeze}(\mathcal{N}', \mathcal{N}'')$ is executed.

(C) The new Π-model is defined to be the set of paths $\{\mathcal{P} \setminus \hat{\mathcal{P}}_{xy} : \mathcal{P} \in \Pi, |\mathcal{P} \setminus \hat{\mathcal{P}}_{xy}| \geq 2\}$, i.e., for every path in Π, its remaining bridges form a path of the new Π-model if their number is at least two. (Note that each generating path in Π gives at most one generating path in the new Π-model, since every path in the new Π-model consists of all remaining bridges from some path in Π.)

Our aim now is to show that the result is a 2-level cactus tree model for the new graph $G_{xy} = G \cup (x, y)$, provided $\lambda(G_{xy}) = \lambda(G)$. Let us denote the objects modified as above by $\mathcal{H}^2_{xy}, \varphi^2_{xy}$, and Π_{xy}. The modeling family $\mathcal{F}^2_{xy}$ is defined to consist of the 1-cuts and minimal 2-cuts of $\mathcal{H}^2_{xy}$ and of the nonminimal 2-cuts of the form $\{\{\varepsilon', \varepsilon''\} : \varepsilon', \varepsilon'' \in \mathcal{P}, \mathcal{P} \in \Pi_{xy}\}$.

Theorem 5.13 *If $\lambda(G_{xy}) = \lambda(G)$, then the model $(\mathcal{H}^2_{xy}, \varphi^2_{xy}, \mathcal{F}^2_{xy})$ is a condensed 2-level cactus tree model of size $O(n)$ for G_{xy}.*

Proof: In order to prove our statement, it is sufficient to show that the cut model $(\mathcal{H}^2_{xy}, \varphi^2_{xy}, \mathcal{F}^2_{xy})$, together with Π_{xy}, satisfies for G_{xy} properties (i-iii) of Theorem 5.1 (condensity is obtained automatically, as above). For convenience, let us modify the order of transformations as follows.

(1) Every cycle of $\mathcal{P}_{xy}$ is squeezed (as a result, $\mathcal{P}_{xy}$ turns into $\hat{\mathcal{P}}_{xy}$).
Execute (C).

(2) Execute (A).

(3) For each cycle-generating subpath with ends $\mathcal{N}_1, \mathcal{N}_2$, the edge $(\mathcal{N}_1, \mathcal{N}_2)$ is added.

It is easy to see that these transformations lead to the same model $(\mathcal{H}^2_{xy}, \varphi^2_{xy}, \mathcal{F}^2_{xy})$ and to the same set of paths Π_{xy}.

Let us prove property (i). Clearly, contraction of an edge or squeezing of a cycle in a cactus tree keeps it a cactus tree. Thus, the structural graph after transformation (1) is a cactus tree. Also, any insertion of a node into an edge at transformation (2) retains the model graph a cactus tree. Let us, at transformation (3), add the edges sequentially. Then, at each stage, there is a unique path in the current cactus tree between the endnodes of the added edge, and this path does not contain any edge of a cycle. Therefore, after adding such an edge, no two cycles have an edge in common, i.e., we arrive at another cactus tree. Hence, $\mathcal{H}^2_{xy}$ is a cactus tree.

Property (ii) is satisfied by the definition of $\mathcal{F}^2_{xy}$. Moreover, the fact that any two paths in Π have at most one edge in common implies immediately the same property for Π_{xy}.

Let us prove property (iii). First, let us show that $(\varphi^2_{xy})^{-1}$ takes bijectively the family of 1-cuts of $\mathcal{H}^2_{xy}$ onto the family of λ-cuts of G_{xy}. Observe that a bisection of V defines a λ-cut of G_{xy} if and only if it defines a λ-cut of G that does not get the inserted edge. It is easy to see that those bisections are represented bijectively in $\mathcal{H}^2$ by the bridges not belonging to $\mathcal{P}_{xy}$. By construction: (a) any such bridge remains in $\mathcal{H}^2_{xy}$, and it φ^2_{xy}-induces in $\mathcal{H}^2_{xy}$ the same bisection of V that it φ^2-induces in $\mathcal{H}^2$, (b) the bridges of $\mathcal{H}^2$ belonging to $\mathcal{P}_{xy}$ after transformation (3) stop being bridges, and (c) no new bridges are made, which suffices.

To finish proving property (iii), it remains to show that $(\varphi^2_{xy})^{-1}$ takes the 2-cuts in $\mathcal{F}^2_{xy}$ onto the family of $(\lambda+1)$-cuts of G_{xy}. Let us consider the $(\lambda+1)$-cuts of G_{xy}: such one is either a λ-cut of G that gets the inserted edge, or a $(\lambda+1)$-cut of G that does not get the inserted edge. Observe that the former are represented in $\mathcal{H}^2$ by the bridges belonging to $\mathcal{P}_{xy}$, and the latter by the minimal 2-cuts not dividing $\mathcal{P}_{xy}$ and by the cuts $\{\varepsilon', \varepsilon''\}$ of the Π-model for which either $\varepsilon', \varepsilon'' \in \mathcal{P}_{xy}$ or $\varepsilon', \varepsilon'' \notin \mathcal{P}_{xy}$.

Let us trace transformations (1-3). It is easy to see that after transformation (1) only minimal 2-cuts of $\mathcal{H}^2$ are affected, and exactly cuts dividing $\mathcal{P}_{xy}$ disappear. According to transformation (2), let us insert the corresponding empty nodes, denoting the path arising from $\hat{\mathcal{P}}_{xy}$ by $\hat{\mathcal{P}}_{xy}$ as well. In the obtained model:

- The λ-cuts that get the inserted edge are modeled by the single edges of $\hat{\mathcal{P}}_{xy}$;
- The $(\lambda+1)$-cuts that do not get the inserted edge are modeled by:
  - the minimal 2-cuts not dividing $\hat{\mathcal{P}}_{xy}$;
  - the nonminimal 2-cuts $\{\varepsilon', \varepsilon''\}$, where $\varepsilon', \varepsilon''$ are two bridges either both belonging to the same cycle-generating subpath, or $\varepsilon', \varepsilon'' \in \mathcal{P} \setminus \hat{\mathcal{P}}_{xy}, \mathcal{P} \in \Pi$ (the latter are the modeling cuts of the Π_{xy}-model).

Observe that at transformation (3) we only add certain edges, and thus the node set and model mapping remain the same. Let us denote the intermediate model obtained after transformation (2) together with the set of modeling cuts listed above by $(\hat{\mathcal{H}}^2, \varphi^2_{xy}, \hat{\mathcal{F}}^2)$. Observe that, by the above discussion, this is a model for the $(\lambda+1)$-cuts of G_{xy}.

Let us show that the family of node bisections defining the 2-cuts in $\mathcal{F}^2$ coincides with the family of node bisections defining the cuts in $\hat{\mathcal{F}}^2$. Recall that any 2-cut in $\mathcal{F}^2_{xy}$ consists of two edges belonging to the same cycle of $\hat{\mathcal{H}}^2_{xy}$, or of two bridges of the form $\{\varepsilon', \varepsilon''\} : \varepsilon', \varepsilon'' \in \mathcal{P}, \mathcal{P} \in \Pi_{xy}$. For any bisection $\mathcal{B}$ of the node set of $\mathcal{H}^2_{xy}$ (and thus of that of $\hat{\mathcal{H}}^2$), it can be easily verified that:

- $\mathcal{B}$ defines a cut of $\mathcal{H}^2_{xy}$ that consists of two edges belonging to the same cycle if and only if it defines in $\hat{\mathcal{H}}^2$ a cut that consists either of two edges belonging to the same cycle, or of a single bridge forming a cycle-generating subpath, or of two bridges belonging to the same cycle-generating subpath.

- $\mathcal{B}$ defines a cut of $\mathcal{H}^2_{xy}$ that consists of two bridges $\varepsilon', \varepsilon''$, such that $\varepsilon', \varepsilon'' \in \mathcal{P}, \mathcal{P} \in \Pi_{xy}$, if and only if it defines in $\hat{\mathcal{H}}^2$ a cut that consists of the same two bridges $\varepsilon', \varepsilon''$, with $\varepsilon', \varepsilon'' \in \mathcal{P} \setminus \hat{\mathcal{P}}_{xy}, \mathcal{P} \in \Pi$.

Thus, $(\varphi^2_{xy})^{-1}$ takes the 2-cuts in $\mathcal{F}^2_{xy}$ onto the λ-cuts of G that get the inserted edge and the $(\lambda+1)$-cuts that do not get the inserted edge. Thus $(\varphi^2_{xy})^{-1}$ takes the family of 2-cuts in $\mathcal{F}^2_{xy}$ onto the family of $(\lambda+1)$-cuts of G_{xy}. This finishes the proof of property (iii) and thus the proof of the Theorem. $\square$

Observe that the transformations of our model, as described, is insensible to increasing of the connectivity of the current graph from λ_0 to λ_0+1 (since nothing depends on *existence* of λ_0-cuts). Therefore, in the next section we discuss implementation for the general case of maintenance of the (λ_0+2)-classes of a graph undergoing any sequence of edge insertions.

5.4.2 Implementation

Let us discuss now implementation of incremental dynamics of (λ_0+2)-classes by means of the 2-level cactus tree model. We use as a subroutine the algorithm [16] that solves our problem for the case $\lambda_0 = 1$. To be free to use it for our purposes, let us give for it formal specifications independent of its concrete semantics. The algorithm [16], given a cut model for a set V whose structural graph is a cactus tree, maintains the classes defined by the nonempty preimages of its nodes under any sequence of update operations

Compress(x,y): do *Squeeze*$(\mathcal{N}_x, \mathcal{N}_y)$, where x, y are two elements of V and $\mathcal{N}_x, \mathcal{N}_y$ are the nodes they are mapped to, respectively,

and at any time is able to answer for any two elements $x, y \in V$ the query

Same-Class$\{x,y\}$*?*: Return "true" if x and y are mapped to the same node of the current model, and "false" otherwise.

In the case the size of the model is $O(n)$, $n = |V|$, the total time complexity for u updates and q queries is $O((u+q+n)\alpha(u+q,n))$ and the space required is $O(n)$.

The algorithm [16] uses a certain tree-like data structure, whose nodes represent the nodes of the model and its cycles. To support dynamics, to each "node" node is assigned a certain union-find type data structure, while to each "cycle" node is assigned a certain data structure of set-splitting type; any primitive operation on those data structures is executed in $O(\alpha(u+q,n))$ amortized time.

Each *Compress*(x,y) operation is done in three phases. First, *Find*(x) and *Find*(y) return $\mathcal{N}_x, \mathcal{N}_y$, respectively. Then, the squeezed path is found, in time linear in its length. Finally, *Squeeze*$(\mathcal{N}_x, \mathcal{N}_y)$ is executed; the number of primitive data structure operations during it is linear in the length of the squeezed path. Each query *Same-Class*$\{x,y\}$*?* is answered by checking $Find(x)$ and $Find(y)$ for equality.

It is shown in [16], that the total length of all paths undergoing *Squeeze* operations during the algorithm, for any sequence of updates, is $O(n')$, where n' is the size of the initial cactus tree. Hence, the total number of primitive operations for finding all squeezed paths and for reorganizing the data structure is $O(n')$, while the total number of *Find* operations is $2(u+q)$. The space required is $O(n+n')$. For the case $n' = O(n)$, this leads to the time and space complexities given above. [5]

In order to maintain the (λ_0+2)-classes of G, we apply the incremental algorithm and the data structure of [16] to the cut model $(\mathcal{H}^2, \varphi^2)$. The additional information on the 2-level cactus tree model is implemented and maintained as follows.

The data structure used in [16] contains the representation of all bridges of the model. Using it, we keep for each bridge of $\mathcal{H}^2$ the list of names of at most four paths in Π containing it. Such name lists provide forming the cycle-generating paths in time linear in the length of $\hat{\mathcal{P}}_{xy}$, by a single pass along that path. Paths in Π are maintained as double-linked edge lists; therefore, excluding a cycle-generating path from a path in Π can be done in $O(1)$ time. Observe that there is no need to update name lists during the algorithm; the space remains linear.

Our algorithm is the following extension of the algorithm [16]. First, given the endpoints x and y of an inserted edge, we call the subroutine of algorithm [16] that returns the squeezed path $\mathcal{P}_{xy}$. We scan that path to find all the cycle-generating subpaths (using the lists defined above) and insert an empty node into each bridge that is contained in two cycle generating subpaths; this can be done in time linear in $|\mathcal{P}_{xy}|$. Finally, for each cycle-generating subpath with endnodes $\mathcal{N}', \mathcal{N}''$, we call subroutine $Squeeze(\mathcal{N}', \mathcal{N}'')$.[6]

To adjust the algorithm and the data structure of [16] to insertions of empty nodes, preserving the complexity, let us use the following trick. Instead of inserting empty nodes

[5] In [16] was considered the case of a model that has no empty nodes. However, the same algorithm works also if there are empty nodes in the model. One of the ways to see this, is by means of the following simple reduction: for every empty node $\mathcal{N}$ of the model, add a "dummy" preimage element $v_{\mathcal{N}}$ to the ground set.

[6] In the case $\lambda_0 = 1$ considered in [16], always the whole $\hat{\mathcal{P}}_{xy}$ plays the role of the single cycle-generating subpath, since in this case *any* pair of edges of $\mathcal{T}^{\lambda_0}$ models a 2-cut of G.

during the algorithm, all possible insertions are done at the initialization by inserting an empty node into *every* bridge of $\mathcal{H}^2$ which belongs to at least two paths of the Π-model (the inserted empty nodes will have degree 2); the total size of the representation remains $O(n)$. The name lists for the two bridges incident to an inserted node are copies of the list for the bridge that this node was inserted into. Thus, no insertions will be needed. However, if all the insertions are done at the initialization, we must modify the transformations to adjust them to the case when the two bridges incident to an inserted node belong to only one cycle-generating subpath, when such an insertion should not have taken place. To reverse the wheel backwards, in such a case, we can simply contract one of the "twin" bridges incident to the inserted node, say $(\mathcal{N}, \mathcal{N}')$, which can be executed by calling twice *Squeeze*$(\mathcal{N}, \mathcal{N}')$; the number of such calls is again linear in $|\mathcal{P}_{xy}|$.

Let us show that the time complexity of our algorithm is the same $O((u+q+n)\alpha(u+q, n))$ as for the algorithm [16]. Indeed, the initial size of the cactus tree model is, as in [16], $O(n)$. In fact, the proof in [16] of the linear bound for the total sum of the lengths of all squeezed paths during the algorithm can be easily generalized to any formally possible sequence of such paths, independently of their semantics. Hence, the total sum of the lengths of the squeezed paths used in our algorithm is, as in [16], $O(n)$. Recall that the additional time spent for finding cycle-generating subpaths and for modifying the Π-model is $O(|P_{xy}|)$, and, hence, does not lead to an increase of complexity.

Let us now show that the space complexity is $O(n)$. Except for the Π-model and the corresponding name lists, all the parts of our model are of the same type as in [16]; their updates are executed by means of the algorithm [16]. Therefore, the space required for those parts is, as in [16], $O(n + n') = O(n)$. Now, the initial space related to the Π-model is $O(n)$, and during the algorithm it can only decrease. Thus, the total space requirements are $O(n)$. This finishes the proof of Theorem 5.12 for the case $\lambda_0 > 3$.

5.4.3 Generalization to the case $\lambda = 3$

Let us now discuss the case $\lambda_0 = 3$ when there are local models which are cubes. As we show, this case can be reduced to the previous one. Let us, for each cube of $\mathcal{H}^2$, choose arbitrarily any its cut formed by four nonadjacent edges; recall that such a cut models a $(\lambda_0 + 1)$-cut, henceforth called "hidden". Then, let us contract everyone of those structural edges, transforming the cube into a 2-uniform cycle of length 4, and then delete one edge from every pair of its parallel edges. Clearly, any 2-cut which consists of two edges of such a cycle φ^2-induces a $(\lambda_0 + 1)$-cut of G: it induces a nondegenerate $(\lambda_0 + 1)$-cut if those edges are nonadjacent and a degenerate $(\lambda_0 + 1)$-cut otherwise. Execution of this operation for every cube of $\mathcal{H}^2$ results in a cut model, whose structural graph is a cactus tree and which satisfies all properties of a 2-level cactus tree model, except for the representation of the hidden $(\lambda_0 + 1)$-cuts. Therefore, maintaining the atoms w.r.t. the λ_0- and nonhidden $(\lambda_0 + 1)$-cuts under insertions of edges into G can be processed by using the algorithm suggested above for the case $\lambda_0 > 3$.

Clearly, the subfamily of hidden $(\lambda_0 + 1)$-cuts is parallel, and thus is modeled by a tree. We maintain the corresponding atoms by treating this tree model separately; obviously, the time complexity cannot be greater than that of the algorithm of [16] (e.g., the algorithm [13, Section 3.2], with the time complexity $O(u + q + n)$ and required space $O(n)$, can be applied). Moreover, since the union of the cut sets represented by our two models at any stage is the entire set of λ_0- and $(\lambda_0 + 1)$-cuts of the current graph, then the answer to a query *Same-4-Class*(x, y)*?* is positive if and only if both x and y belong to the same atom w.r.t. *each* of these two systems. In this way, the case $\lambda_0 = 3$ is reduced to the previous one, which finishes the proof of Theorem 5.12.

6 Concluding remarks

1. Observe that the properties mentioned in Theorem 5.1 are similar to those of the cactus tree model for the minimum cuts, though more complicated. Since the structure of the modeling cuts is explicit and, in a sense, simple, and since the representation is compact, our model seems to be convenient to represent the minimum and minimum+1 cuts of graphs in various applications.

2. Our difficulties to include the case $\lambda = 3$ into the general scheme seem to be explainable. According to [24], there are special properties of cut families when cardinalities of the cuts are strictly within $\frac{4}{3}\lambda$; observe that if $\lambda = 3$ then $\frac{\lambda+1}{\lambda} = \frac{4}{3}$, while for $\lambda > \frac{4}{3}$ holds $\frac{\lambda+1}{\lambda} < \frac{4}{3}$.

3. In [5], the k-cuts of an arbitrary graph G dividing only one its 3-class S are described as generated by the k-cuts of the corresponding *3-component* $\bar{S}$ (which is defined slightly differently than in this paper); the 3-components are shown to be 3-connected graphs. Using this description and the 2-level cactus tree models for all such 3-components of G, a complete description of k-cuts, $k \leq 4$, of G can be given. The structure thus obtained is used in [11] to maintain effectively the 5-classes in an arbitrary incremental graph.

4. Let us call a cut of a weighted graph **subminimum** if its weight is the second minimum. Let us suggest as an open problem generalizing the results of this paper to modeling the minimum and subminimum λ'-cuts, in the case $\frac{\lambda'}{\lambda} \leq \frac{4}{3}$, for an arbitrary weighted graph. To achieve this generalization, the techniques used in this paper separately for odd and even cases have to be combined; therefore, the new construction is expected to have the difficulties of both these cases simultaneously, and maybe even more.

Acknowledgments. The authors are very grateful to Jeffery Westbrook for fruitful discussions and to Alek Vainshtein for useful comments and help in improving the presentation of this paper. The authors also thank an anonymous referee whose comments helped improving the presentation of this paper.

References

[1] A. A. Benczur, "Augmenting undirected connectivity in $\tilde{O}(n^3)$ time", *Proc. 26th Annual ACM Symp. on Theory of Computing*, ACM Press, 1994, 658–667.

[2] A. A. Benczur, "The structure of near-minimum edge cuts", In *Proc. 36st Annual Symp. on Foundations of Computer Science*, 1995, 92–102.

[3] R. E. Bixby, "The minimum number of edges and vertices in a graph with edge connectivity n and m n-bonds", *Networks* **5**, 1975, 253–298.

[4] T. Cormen, C. Leiserson, and R. Rivest. *Introduction to Algorithms*, McGraw-Hill, New York, NY, 1990.

[5] Ye. Dinitz, "The 3-edge components and the structural description of all 3-edge cuts in a graph", *Proc. 18th International Workshop on Graph-Theoretic Concepts in Computer Science (WG92), Lecture Notes in Computer Science*, v.657, Springer-Verlag, 1993, 145–157.

[6] Ye. Dinitz, "Maintaining the 4-edge connected components of a graph on line", *Proc. 2-nd Israel Symposium on Theory of Computing and Systems (ISTCS'93)*, IEEE Computer Society Press, 1993, 88–97.

[7] E. A. Dinic, A. V. Karzanov and M. V. Lomonosov, "On the structure of the system of minimum edge cuts in a graph", *Studies in Discrete Optimization*, A. A. Fridman (Ed.), Nauka, Moscow, 1976, 290–306 (in Russian).

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

[9] Ye. Dinitz and Z. Nutov, "A 2-level cactus tree model for the minimum and minimum+1 edge cuts in a graph and its incremental maintenance. Part II: the even case", a manuscript.

[10] Ye. Dinitz and Z. Nutov, "Cactus-tree type models for families of bisections of a set", a manuscript.

[11] Ye. Dinitz and Ronit

[12] Ye. Dinitz and A. Vainshtein, "The connectivity carcass of a vertex subset in a graph and its incremental maintenance", *Proc. 26th Annual ACM Symp. on Theory of Computing*, ACM Press, 1994, 716–725 (see also TR-CS0804, Technion, Haifa, Israel, June 1994).

[13] Ye. Dinitz and J. Westbrook, "Maintaining the Classes of 4-Edge-Connectivity in a Graph On-Line", Technical Report #871, Dept. of Comp. Sci., Technion, Haifa, israel, 1995, 47p. (To appear in *Algorithmica*).

[14] H. N. Gabow, "Applications of a poset representation to edge connectivity and graph rigidity", *Proc. 32nd Symp. on Foundations of Computer Science*, 1991, 812–821.

[15] H. N. Gabow, "A representation for crossing set families with application to submodular flow problems", *Proc. 4th Annual ACM-SIAM Symp. on Discrete Algorithms*, 1993, 202–211.

[16] Z. Galil and G. F. Italiano, "Maintaining the 3-edge-connected components of a graph on line", *SIAM J. Computing* **22** (1), 1993, 11–28.

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

[18] D. R. Karger, "Global min-cuts in RNC, and other ramifications of a simple min-cut algorithm", In *Proc. 4th ACM-SIAM Symposium on Discrete Algorithms*, 1993, 21–30.

[19] A. V. Karzanov and E. A. Timofeev, "Efficient algorithm for finding all minimal edge-cuts of a non-oriented graph", *Cybernetics*, 156–162. (Translated from Kibernetica, Vol. 22, No. 2, 1986, 8–12).

[20] S. Khuller and R. Thurimella, "Approximation algorithms for graph augmentation", *Journal of Algorithms* **14**, 1993, 214–225.

[21] J. A. La Poutré, J. van Leeuwen, and M. H. Overmars. “Maintenance of 2-and 3-edge-connected components of graphs”, *Discrete Mathematics* **114**, 1993, 329–359.

[22] D. Naor, D. Guisfield and C. Martel, “A fast algorithm for optimally increasing the edge connectivity”, In *Proc. 31st Annual Symp. on Foundations of Computer Science*, 1990, 698–707.

[23] H. Nagamochi and T. Kameda, “Canonical cactus tree representation for minimum cuts”, *Japan J. of Ind. and Appl. Math.*, **11**, 1994, no.3, 343–361.

[24] H. Nagamochi, K. Nishimara and T. Ibaraki, “Computing all small cuts in an undirected network”, *SIAM J. on Discrete Mathematics*, **10**, No. 3, 1997, 469–481.

[25] J.C. Picard and M. Queranne, “On the structure of all minimum cuts in a network and applications”, *Math. Programming Study*, **13**, 1980, 8–16.

[26] J. Westbrook, “Incremental algorithms for four-edge connectivity”, Extended Abstract of the lecture held at the CS seminar of Bell Labs on March 10, 1993.