A 2-level cactus tree model for the system of minimum and minimum+1 edge cuts of a graph and its incremental maintenance. Part II: the even 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@tx.technion.ac.il

November 4, 1998

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 even $\lambda \geq 4$. The case of odd connectivity is considered in the companion paper (Part I).

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.

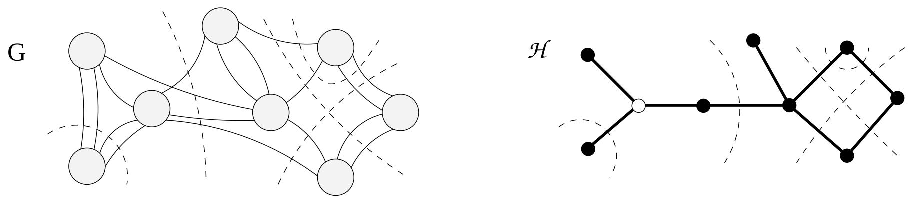

Figure 1: Cactus-tree model of a 4-connected graph. (The gray blobs represent vertex classes of 5-connectivity. Some cuts of the graph and their representing cuts of the cactus tree are shown by dashed lines.)

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 (see [4, 10, 11, 3, 17, 9, 8, 1] and others); 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, and let λ denote the cardinality of a minimum cut of G. When analyzing the connectivity of G, the natural first "stratum" is the set of its minimum cuts 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 minimal cuts of the cactus tree model $(\mathcal{H}, \varphi)$ [4] :[1] $\mathcal{H}$ 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 minimal cuts of $\mathcal{H}$ onto 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 $O(n)$.

The cactus tree $\mathcal{H}$ 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 [15]). In this paper, we consider the unique version of $\mathcal{H}$ in which all empty nodes of degree 3 are replaced by cycles.

[1] An equivalent model was suggested independently in [10].

Since the structure of the cactus tree $\mathcal{H}$ and its connection to G are simple, and since this representation is compact, it is a good model for connectivity studies. 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 [10].

In this paper for the case λ even and in the companion paper [6] for the case λ odd, we suggest an extension of the cactus tree model, called the 2-level cactus tree model. Our models represents, in a clear and compact way, the system of the λ- and $(\lambda+1)$-cuts and of the vertex classes of $(\lambda+2)$-connectivity of a graph; their sizes are $O(n)$. Previous results are as follows: Galil and Italiano [11] and La Poutré et al. [13] suggested a structure for the case $\lambda = 1$, Dinitz [3] and Westbrook [17] suggested another structure for the case $\lambda = 2$. Also 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 $\lambda = 2$.

Benczur in [1] gives a geometric representation of the cuts of weight less than $\frac{6}{5}\lambda$. His model is less compact in comparison with our models: its size is $O(n^2)$. For $\lambda \leq 5$ our model is stronger, since $\lambda + 1 > \frac{6}{5}\lambda$; 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 [1] is stronger than our models. It is worth to mention that practical networks usually require rather small connectivity.

A natural object to represent a cut of a graph is the corresponding bisection (partition into two nonempty parts) of the vertex set; in a connected graph, there is a bijective correspondence between the cuts of the graph and the bisections of its vertex set [6]. 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.

We define a cut model for a bisection family F of a set V as a triple $(\mathcal{G}, \psi, \mathcal{F})$, where $\mathcal{G}$ is a ("structural") graph, ψ is a mapping from V to the node set of $\mathcal{G}$, and $\mathcal{F}$ is a family of cuts of $\mathcal{G}$ such that $\psi^{-1}(\mathcal{F}) = F$ ("modeling family"). The constructive generic "2-level" approach for modeling bisection families of a set by cut models [4, 5, 7] is as follows. Two bisections are called parallel if they collectively partition V into 3 parts. We choose a certain family F^{bas} consisting of mutually parallel bisections (henceforth, we call them "basic"); it is modeled by a tree $\mathcal{T}^{bas}$. The bisections in $F \setminus F^{bas}$ are classified w.r.t. F^{bas} and modeled based on $\mathcal{T}^{bas}$. Bisections parallel to all basic ones are called local; each of them is naturally assigned to a node of $\mathcal{T}^{bas}$. For every subfamily of local bisections assigned to the same node, we construct a cut model (so called "local model"). We obtain a united model for all basic and local bisections by a natural "implanting" of each local model instead of the related node of $\mathcal{T}^{bas}$. The modeling family can be extended to represent also the nonlocal bisections in $F \setminus F^p$ (called "global") if we choose F^{bas} such that the partition of V by the basic and local bisections coincides with the partition of V by the whole family F. In this paper, the natural "implanting" of local models is modified in order to simplify the representation of global bisections.

Another type of cut models used in this paper is the "r-skeleton" [6] that represents the cuts of G up to cardinality $\lambda + r - 1$ and the set of classes of $(\lambda + r)$-connectivity that V is cut into by these cuts. When passing from cuts of G to the modeling cuts in the skeleton, the cardinality decreases by a constant value. This provides an immediate reduction of the incremental maintenance problem for the original graph to the corresponding problem for its skeleton model.

As in [11, 13] ($\lambda = 1$) and in [9] ($\lambda = 2$), we use our models for the incremental maintenance of the classes of $(\lambda_0 + 2)$-connectivity, $\lambda_0 \geq 3$, where λ_0 is the connectivity of the initial graph. This means that we support our structures 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-$(\lambda_0 + 2)$-*Class*(x, y)*?*: Return "true" if two given vertices x and y belong to the same $(\lambda_0 + 2)$-class of G, and "false" otherwise.

The 2-level cactus tree model suggested in this paper is a skeleton of connectivity 2. Thus we have a reduction of the maintenance problem to the case $\lambda_0 = 2$, preserving the complexity. For an arbitrary sequence of u updates *Insert-Edge* and q queries *Same*-$(\lambda_0 + 2)$-*Class*(x, y)*?*, total time required is $O(u + q + n \log n)$, using the algorithm [9], or $O((u+q+n)\alpha(u+q, n))$, using that of [17], where α is the inverse of the Ackerman function (which grows extremely slow, see, e.g., [2]).

A 2-level cactus tree model for an arbitrary graph can be constructed in polynomial time by using the deterministic algorithm of [16] or the randomized algorithm of [12] for listing of near minimum cuts.

This paper is organized as follows. Section 2 brings basic definitions and notations. Section 3 introduces our tools: 2-level cut modeling and skeleton models, and presents some properties of λ- and $(\lambda + 1)$-cuts; these two sections being abridged versions of the corresponding part of [6], except for the new Lemma 3.8. Section 4 deals with both statics and dynamics for the case of even λ. Section 5 contains concluding remarks.

For a more detailed introduction see [6]. The preliminary version of this and the companion papers is Extended Abstract [5].

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 $|E| = m$. 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). By following statement, it is legal to study cuts as vertex bisections.

Proposition 2.1 ([6]) *For every cut of a connected graph, there is a unique bisection of the vertex set defining it.*

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 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

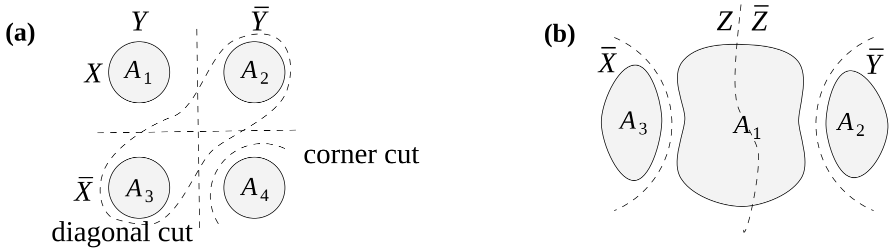

Figure 2: Relations between bisections: (a) crossing bisections; (b) parallel bisections and a bisection between them.

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. For any graph H, let $V(H)$ and $E(H)$ denote the vertex and edge sets of H, respectively.

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 parallel bisections $\{X, \bar{X}\}$ and $\{Y, \bar{Y}\}$, where $\bar{X} \cap \bar{Y} = \emptyset$, 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.

When V is the vertex set of a graph G, similar definitions are used for cuts, considering them as bisections of V. The quotient graph defined by the four corner sets of two cuts C, C' is called the $\{C, C'\}$**-square**. An edge of the square belonging to both C and C' is called a **diagonal edge**; 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 $i \neq j = 1, \ldots, 4$.

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 $\boldsymbol{l}$**-uniform** if the weight of every edge in the latter is l.

3 Modeling tools and auxiliary statements

In this Section, we introduce our modeling techniques. Section 3.1 introduces the hierarchic 2-level approach of [5, 7] to the construction of cut models for families of bisections. In Sect. 3.2, we introduce skeleton models of [6]. Sect. 3.3 presents some properties of cuts proved in [6].

3.1 Cut models

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 [4]. Following [7], 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 $\psi^{-1}(\mathcal{F}) = F$; there $\mathcal{F}$ 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.

[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 $\mathcal{C}, \mathcal{F}, \mathcal{G}, \mathcal{V}(\mathcal{G})$.

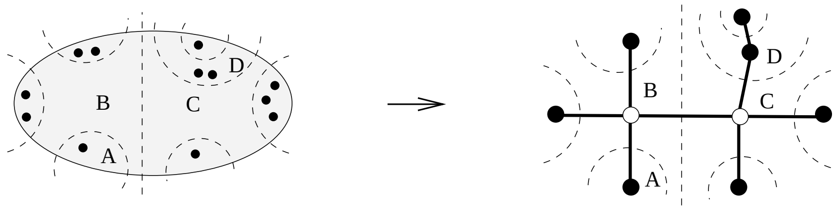

Figure 3: A parallel family and its tree model. (The nodes whose preimages are empty are shown white. Bisections and cuts are shown by dashed lines.)

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 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 the modeling family partitions the node set of the structural graph into singletons.

Generalizing [4] (and following [7]), we allow *indirect* descriptions of the modeling family. Modeling cuts are groupped into so called **bunches**. Each bunch is presented in a clear and compact way; its definition consists of its type and a constant number of parameters, e.g., references to certain edges sets. Observe that such a *structural description* can be preferable in comparison with a trivial listing of all the cuts (for example, see in Introduction on the cactus tree model).

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 $\mathcal{F}$ (when it is indirect of the above type, its size is the total length of bunch definitions and auxiliary sets). Observe that thus 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.

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 [14], $|F^p| = O(|V|)$ (in fact, $|F^p| \leq 2|V| - 3$). Following [4], we represent such a family by the naturally defined **tree model** $(\mathcal{T}^p, \psi^p, 1\text{-cuts of } \mathcal{T}^p)$, where $\mathcal{T}^p$ is a tree (see Fig. 3, for a formal definition see [9, 7])). This model is condensed and is bijective, i.e., every bisection in F^p is ψ^p-induced by a unique 1-cut of $\mathcal{T}^p$. For a node $\mathcal{N}$ of $\mathcal{T}^p$, the family of bisections $F^p_{\mathcal{N}} = \{(\psi^p)^{-1}(\mathcal{C}) : \mathcal{C} = (\mathcal{N}, \mathcal{N}') \text{ is a 1-cut of } \mathcal{T}^p\}$ is called the **neighbor group at $\boldsymbol{\mathcal{N}}$**. Two bisections belong to the same neighbor group if and only if there is no other bisection in F^p between them.

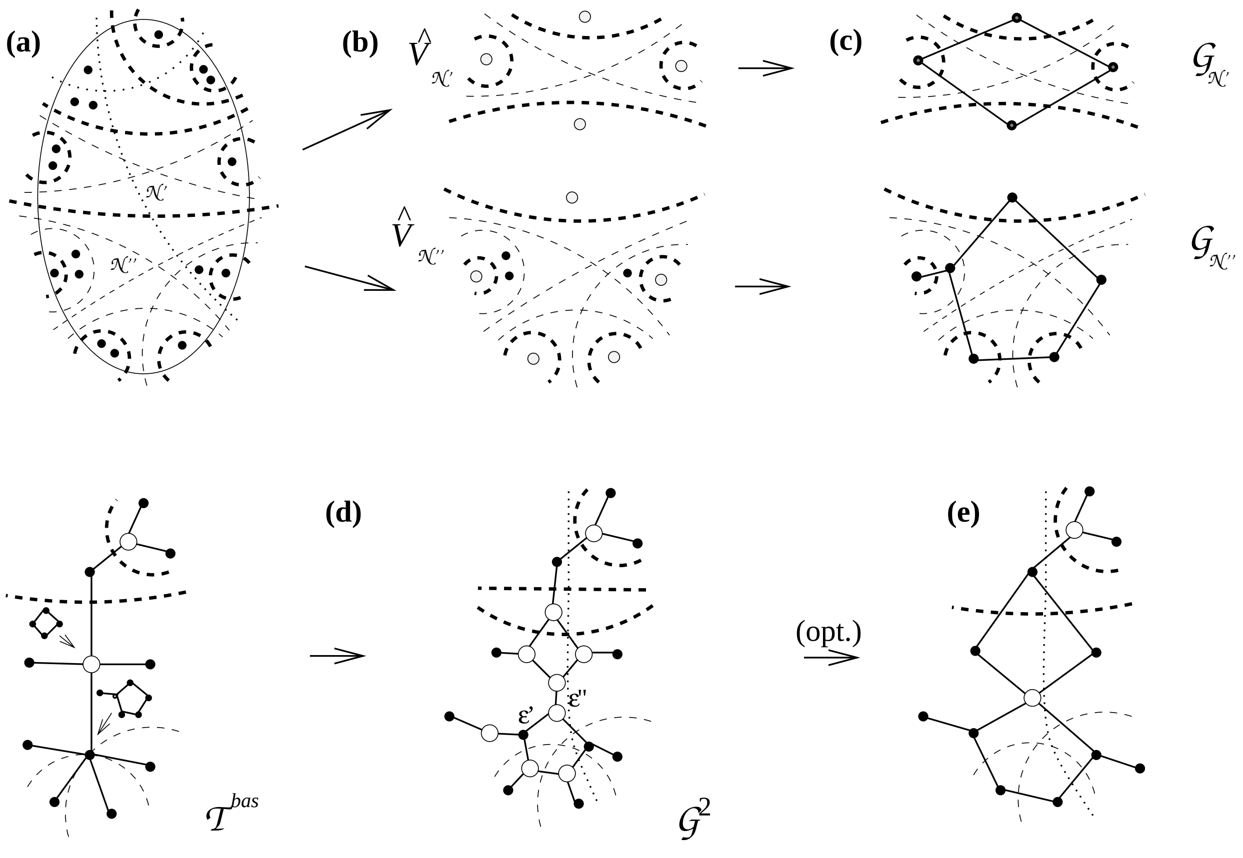

Figure 4: (a) decomposition of a bisection family w.r.t. a basic family (basic bisections are shown by thick dashed lines, local bisections by thin dashed lines, and global bisections by dotted lines); (b) node sets of components at $\mathcal{N}'$ and at $\mathcal{N}''$ (halo nodes are shown gray), and the decomposition of F^{loc}; (c) local models; (d) implanting; (e) the result of optional contractions (in (d) and (e) modeling cuts are shown partly; one of the global bisections is not modeled).

Assume we are looking for a cut model for a bisection family F of a set V. Given a parallel bisection family, we use the following classification of bisections in F w.r.t. that family (see Fig. 4(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 in F is called **local** if it is not crossing with any member of F^{bas}, and **global** otherwise; F^{loc} and F^{gl} denote the corresponding subfamilies of F, respectively. (Recall that here and everywhere in this section similar definitions are implicitly made for a cut family F, considering cuts as vertex bisections.)

We decompose the local bisections relatively to the nodes of $\mathcal{T}^{bas}$ by means of the following model (see Fig. 4(b)). The **component** $\hat{V}_{\mathcal{N}}$ at a node $\mathcal{N}$ of $\mathcal{T}^{bas}$ is defined to be the quotient set of V (or the quotient graph $\hat{G}_{\mathcal{N}}$ of G, in the case of a cut family F) which is obtained from V by shrinking, for every branch $\mathcal{B}$ hanging at $\mathcal{N}$ in $\mathcal{T}^{bas}$, the subset $(\psi^{bas})^{-1}(\mathcal{B})$ into a single **halo element** (resp., **halo node**); the corresponding quotient mapping is denoted by $\hat{\psi}_{\mathcal{N}}$.

Lemma 3.1 ([7]) *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.*

By Lemma 3.1, F^{loc} falls into parts $F^{loc}_{\mathcal{N}}$ corresponding to nodes $\mathcal{N}$ of $\mathcal{T}^{bas}$ (via compatibility with $\hat{V}_{\mathcal{N}}$). Our general approach is to represent the parts $F^{loc}_{\mathcal{N}}$ separately, and then to sinthesize the entire representation for $F^{bas} \bigcup F^{loc}$ (everywhere in this paper holds $F^{bas} \subseteq F$, which implies $F^{bas} \bigcup F^{loc} \subseteq F$). 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 cuts compatible with $\hat{V}_{\mathcal{N}}$. A **local model** at $\mathcal{N}$ is a cut model $(\mathcal{G}_{\mathcal{N}}, \psi_{\mathcal{N}}, \mathcal{F}'_{\mathcal{N}})$ with the following properties:

- the modeled family $F'_{\mathcal{N}}$ satisfies $F^{loc}_{\mathcal{N}} \subseteq F'_{\mathcal{N}} \subseteq F^{loc}_{\mathcal{N}} \cup F^{bas}_{\mathcal{N}}$;
- 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 $\mathcal{G}_{\mathcal{N}}$.[3]

It can be easily shown that $\mathcal{G}_{\mathcal{N}}$ is a local model at $\mathcal{N}$ if and only if it is a composition of $(\hat{V}_{\mathcal{N}}, \hat{\psi}_{\mathcal{N}}, \hat{\psi}_{\mathcal{N}}(F'_{\mathcal{N}}))$ and a cut model for $\hat{\psi}_{\mathcal{N}}(F'_{\mathcal{N}})$, with $F'_{\mathcal{N}}$ as above (see Fig. 4(c)). Hence, in applications to cut families, we usually obtain a local model at $\mathcal{N}$ via the component $\hat{G}_{\mathcal{N}}$ by constructing a cut model for the family $\hat{\psi}_{\mathcal{N}}(F'_{\mathcal{N}})$ of cuts of $\hat{G}_{\mathcal{N}}$, for an appropriate family $F'_{\mathcal{N}}$.

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 model for $F^{loc} \cup F^{bas}$ as follows (for illustration see Fig. 4(d)). 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

[3] This requirement is naturally fulfilled in known models. Moreover, it is not very restrictive in general, since any model for $F'_{\mathcal{N}}$, $F^{loc}_{\mathcal{N}} \subseteq F'_{\mathcal{N}} \subseteq F^{loc}_{\mathcal{N}} \cup F^{bas}_{\mathcal{N}}$, can be easily modified to satisfy it, by appropriate shrinkings.

$\mathcal{N}_{\mathcal{B}^\varepsilon} = \psi_{\mathcal{N}}((\psi^{bas})^{-1}(\mathcal{B}^\varepsilon))$ of $\mathcal{G}_{\mathcal{N}}$, and then deleting $\mathcal{N}$. Let us denote the resulting graph by $\mathcal{G}^2$.

Since V is the union of nonempty preimages of nodes of $\mathcal{T}^{bas}$, it is sufficient to define the mapping $\psi^2 : V \to \mathcal{V}(\mathcal{G}^2)$ on everyone of such preimages. For a nonempty node $\mathcal{N}$ of $\mathcal{T}^{bas}$, ψ^2 takes elements in $(\psi^{bas})^{-1}(\mathcal{N})$ as the mapping $\psi_{\mathcal{N}}$ does, in the case $\mathcal{N}$ has undergone implanting, and takes all of them to $\mathcal{N}$, otherwise.

Let us define the modeling family. By the construction, there is a natural bijective correspondence between the structural edges of $\mathcal{G}^2$ and the structural edges of the local models and of the basic tree. Let us consider the modeling cuts in $\mathcal{F}_{\mathcal{N}}$ and in $\mathcal{F}_{bas}$ as edge sets. By [7], the above correspondence applied to them defines edge sets of cuts of $\mathcal{G}^2$, and those cuts represent, via ψ^2-inducing, the same bisections of V; in what follows we identify such corresponding edges and such corresponding cuts. In this sense, we define

$$\mathcal{F}^2 = (\bigcup\{\mathcal{F}'_{\mathcal{N}} : \text{there is a local model } \mathcal{G}_{\mathcal{N}}\}) \bigcup \mathcal{F}^{bas}.$$

Theorem 3.2 ([7]) *$(\mathcal{G}^2, \psi^2, \mathcal{F}^2)$ 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 $\{\mathcal{G}_{\mathcal{N}}\}$.

The following Lemma shows several properties that are expanded from local models to their plant model.

Lemma 3.3 ([7]) (i) *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.*

(ii) *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.*

Observe that if a basic bisection $\mathcal{B}$ is represented in at least one local model, then the structural edge inherited from $\mathcal{T}^{bas}$ that represents $\mathcal{B}$ can be contracted without losing representation for $\mathcal{B}$ (see, for example, edges ε' and ε'' in Fig. 4(d,e)). Executing such contractions is optional, and each of them can be done independently from the others; henceforth, we refer to them as "optional contractions".

An important particular choice of F^{bas} (used in [4]) is as follows. We say that a bisection in F is **separating** if it is parallel to any other member of F. We denote the subfamily of all separating bisections in F by F^{sep}; clearly, it is parallel. Choosing $F^{bas} = F^{sep}$ implies $F^{bas} \subseteq F$ and $F^{gl} = \emptyset$.

3.2 Skeleton models

A particular type of cut models used in this paper is skeleton models.

Definition 3.4 *A cut model $(\mathcal{G}, \psi)$ for a graph G is called an* **r-skeleton** *if it is condensed, and ψ^{-1} takes the set of all $(\lambda(\mathcal{G}) + i)$-cuts of $\mathcal{G}$ onto the set of all $(\lambda(G) + i)$-cuts of G, for any $i = 0, \ldots, r-1$.*

A simple example of an r-skeleton is the quotient graph of G w.r.t. the partition of V into $(\lambda(G) + r)$-classes.

It is easy to see, by the definition of the skeleton, that for any $i \leq r$, the ψ-preimages of $(\lambda(\mathcal{G}) + i)$-classes of $\mathcal{G}$ are exactly the $(\lambda(G) + i)$-classes of G. This implies the following statement.

Lemma 3.5 *Let $(\mathcal{G}, \psi)$ be an r-skeleton for G and let i be an integer not exceeding r. Then the answer to a query* Same-$(\lambda(G) + i)$-Class(x, y)? *for G is the same as the answer to the query* Same-$(\lambda(\mathcal{G}) + i)$-Class$(\psi(x), \psi(y))$? *for $\mathcal{G}$.*

An important property of a skeleton is that it is, in a sense, stable under insertions of edges into G (note that such an insertion can increase the connectivity of a graph at most by 1).

Theorem 3.6 ([6]) *Let $(\mathcal{G}, \psi)$ be an r-skeleton for G. Let G' be the graph obtained from G by inserting an edge (x, y); let $(\mathcal{G}', \psi')$ be the model obtained by inserting the edge $(\psi(x), \psi(y))$ into $\mathcal{G}$ and, then, shrinking each of the arising nonsingleton $(\lambda(\mathcal{G}) + r)$-classes. If $\lambda(G') = \lambda(G)$, then $(\mathcal{G}', \psi')$ is an r-skeleton for G'; if $\lambda(G') = \lambda(G) + 1$ and $r > 1$, then $(\mathcal{G}', \psi')$ is an $(r-1)$-skeleton for G'.*

Corollary 3.7 ([6]) *Let $(\mathcal{G}, \psi)$ be an r-skeleton for G. Let G undergo edge insertions and let λ_G^0 and $\lambda_{\mathcal{G}}^0$ denote the initial connectivities of G and $\mathcal{G}$, respectively. Then the time complexity of maintaining the $(\lambda_G^0 + i)$-classes of G under any sequence of updates* Insert-Edge *and queries* Same-$(\lambda_G^0 + i)$-Class?, *$i = 1, \ldots, r$, is the same as the time complexity of maintaining the $(\lambda_{\mathcal{G}}^0 + i)$-classes of $\mathcal{G}$ undergoing the corresponding updates and queries.*

Observe that the cactus model $(\mathcal{H}, \varphi)$ with weights (i.e., multiplicities) 2 for its tree-edges and 1 for edges belonging to a cycle (henceforth called "cycle-edges") is a 1-skeleton for G, with $\lambda(\mathcal{H}) = 2$. We use this version of a weighted cactus model in this paper.

3.3 Some properties of $\boldsymbol{\lambda}$- and $(\boldsymbol{\lambda}+\mathbf{1})$-cuts

Recall that F^λ denotes the family of minimum cuts of G. As recommended at the end of Sect. 3.1, let us consider as the basic family for modeling $F = F^\lambda \cup F^{\lambda+1}$ its subfamily $(F^\lambda)^{sep}$; let $\mathcal{T}^\lambda$ denote the tree modeling it. Observe that the cactus tree model $(\mathcal{H}, \varphi)$ is a plant model for F^λ based on $\mathcal{T}^\lambda$ and built by implanting into it cycles (i.e., local models which are cycles), in the way of Theorem 3.2. By [4], for an arbitrary node of $\mathcal{T}^\lambda$ that a local model is implanted instead of it, the corresponding component has a special uniform cycle structure and coincides with that local model. We elaborate this property as follows (recall that the version of the cactus tree model that we use is with cycles of length 3 implanted instead of all empty nodes of degree 3).

Lemma 3.8 *When constructing the cactus tree model, a local model $\mathcal{G}_\mathcal{N}$ is implanted instead of a node $\mathcal{N}$ if and only if the component $\hat{G}_\mathcal{N}$ is the $\frac{\lambda}{2}$-uniform cycle on the set of its halo nodes.*

In this case, $\mathcal{N}$ is an empty node, the length of the cycle $\hat{G}_\mathcal{N}$ is equal to the degree of $\mathcal{N}$ in $\mathcal{T}^\lambda$, and the local model coincides with $\hat{G}_\mathcal{N}$ together with its minimal cuts, i.e., 2-cuts if edge weights ignored.

Proof: The *if*-part and the last group of statements are straightforward, by definitions of a component and of a local model. Indeed, if $\hat{G}_\mathcal{N}$ is a $\frac{\lambda}{2}$-uniform cycle, then this cycle, with edge weights ignored, together wits its 2-cuts is a cactus tree model for the λ-cuts of $\hat{G}_\mathcal{N}$; if the length of this cycle is at least 4, then the subfamily $(F^\lambda)^{loc}_\mathcal{N}$ is nonempty, and the local model must be implanted; if this length is 3, we implant a triangle by our choice.

The *only-if*-part for the case $(F^\lambda)^{loc}_\mathcal{N}$ nonempty is given in [4]. The only remaining case is an empty node of degree 3. The corresponding component has exactly three nodes, all of them halo, and all three its cuts are of cardinality λ. An easy computation shows that each of the three edges must have the weight $\frac{\lambda}{2}$. $\square$

By the construction of a plant model, since only empty nodes of $\mathcal{T}^\lambda$ are affected by implanting, φ coincides with the model mapping for $\mathcal{T}^\lambda$. According to [4], all the optional contractions are executed. Recall that $\varphi^{-1}(\mathcal{N})$, for any nonempty node $\mathcal{N}$, is a $(\lambda+1)$-class of G. The components w.r.t. $\mathcal{T}^\lambda$ are called $(\boldsymbol{\lambda}+\mathbf{1})$**-components** of G (observe that this is a generalization of the concept of a 3-component given in [3]). The following important statement is straightforward.

Lemma 3.9 (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.*

For establishing the structure of the global $(\lambda+1)$-cuts, we use the following statements.

Lemma 3.10 ([6]) *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 has one edge of weight $\frac{\lambda}{2}+1$ and three edges of weight $\frac{\lambda}{2}$.*

Corollary 3.11 ([6]) *If λ is even, then any $(\lambda+1)$-cut divides exactly one $(\lambda+1)$-class.*

Remark: By Corollary 3.11(ii), if λ is even, then the totality of Connectivity Carcasses [8] of all $(\lambda+1)$-classes of G represents all the $(\lambda+1)$-cuts of G. However, Section 4 presents a more compact data structure, which, in addition, allows more efficient incremental maintaining.

4 The even case

In this section we prove in a constructive way the following reduction to the case $\lambda = 2$ (for examples see Figures 17 and 18 at the end of this section).

Theorem 4.1 *For any connected graph with even connectivity, there exists a 2-skeleton of connectivity 2 and size $O(n_{\lambda+2}) = O(n)$.*

This provides, in particular, by Section 3.2, a straightforward reduction of the (λ_0+2)-classes maintenance problem, for an arbitrary graph with $\lambda_0 \geq 4$ even, to the 4-classes maintenance problem for its 2-skeleton; for the latter problem there exists an efficient algorithm [9].

In this section the modeled family $F^\lambda \cup F^{\lambda+1}$ is denoted by F and the basic family is $(F^\lambda)^{sep}$. We begin with showing, in Section 4.1, a method to generate the global cuts from the local cuts and the cactus tree model $\mathcal{H}$. Further, in Section 4.2, a special modification of a $(\lambda+1)$-component is suggested that enables to establish a structure of its local cuts. Finally, unwrapping the reductions, we build in Section 4.3 a 2-skeleton for G, whose connectivity is 2; its 2- and 3-cuts have a simple explicitly given structure.

4.1 Structure of global cuts

Let us recall some general information, according to Section 3.1. The family $(\mathcal{F}^\lambda)^{sep}$ is modeled by the tree $\mathcal{T}^\lambda$. To each node $\mathcal{N}$ of $\mathcal{T}^\lambda$ corresponds the neighbor group $F^\lambda_\mathcal{N}$ formed by λ-cuts that are φ-induced by the cuts of $\mathcal{T}^\lambda$ defined by single edges incident to $\mathcal{N}$. Each separating λ-cut belongs to exactly two neighbor groups and is induced by a structural edge of $\mathcal{T}^\lambda$ with the endnodes corresponding to those neighbor groups. All λ-cuts of G are modeled by the cactus $\mathcal{H}$ that is obtained by implanting cycles instead of certain empty nodes of $\mathcal{T}^\lambda$ and executing all the optional contractions (i.e., contracting each bridge that is incident to an implanted cycle).

Let us consider a cycle $\mathcal{L}$ of $\mathcal{H}$ implanted instead of a node $\mathcal{N}_\mathcal{L}$ of $\mathcal{T}^\lambda$. By Lemma 3.8, the component of $\mathcal{N}_\mathcal{L}$ is a $\frac{\lambda}{2}$-uniform cycle on the halo nodes. Hence, to the structural edges ε of $\mathcal{L}$ correspond pairwise disjoint sets E_ε of edges of G, each of cardinality $\frac{\lambda}{2}$ (edges in such a set connect vertices of G, which are mapped into branches of $\mathcal{T}^\lambda$ hanging on $\mathcal{N}_\mathcal{L}$ and neighboring according to the order given by $\mathcal{L}$). By Lemma 3.8, all nonempty nodes of $\mathcal{T}^\lambda$ remain in $\mathcal{H}$ (i.e., are not replaced by a local model).

By Corollary 3.11, any global $(\lambda+1)$-cut of G divides exactly one $(\lambda+1)$-class. Let us fix such a class S and analyze the global $(\lambda+1)$-cuts dividing it. For simplicity, we denote the (nonempty) node $\varphi(S)$ by $\mathcal{N}_S$, the component $\hat{G}_{\mathcal{N}_S}$ by $\hat{G}_S$, and the neighbor group $F^\lambda_{\mathcal{N}_S}$ by F^λ_S. The cuts in the neighbor group F^λ_S, together, separate S from $V \setminus S$. For any cut $R \in F^\lambda_S$, let us denote the corresponding bisection of V by $\{X_R, \bar{X}_R\}$, where $S \subseteq X_R$; notice that, for two distinct cuts $R_1, R_2 \in F^\lambda_S$, the sets $\bar{X}_{R_1}, \bar{X}_{R_2}$ are disjoint. Shrinking of every such set $\bar{X}_R$ into a halo node results in $\hat{G}_S$.

Let us consider a global $(\lambda+1)$-cut C dividing S. Since it is not local, it crosses at least one cut R in F^λ_S. Let us show, what a structure corresponds to such a crossing cut pair (for example see Fig. 5).

Lemma 4.2 *Let $F^\lambda_{\mathcal{N}(R)}$ be the neighbor group that contains R and is distinct from F^λ_S. Then:*

(i) *the node $\mathcal{N}(R)$ of $\mathcal{T}^\lambda$ is replaced by a cycle $\mathcal{L}_{\mathcal{N}(R)}$ in $\mathcal{H}$, and $\mathcal{N}_S$ is a node of this cycle in $\mathcal{H}$;*

(ii) *the quotient graph of the partition into the $(F^\lambda_{\mathcal{N}(R)} \cup C)$-atoms is a cycle $\mathcal{L}_{\mathcal{N}(R),C}$ with one edge $\hat{\varepsilon}$ of the weight $\frac{\lambda}{2}+1$, such that the contraction of $\hat{\varepsilon}$ results in $\mathcal{L}_{\mathcal{N}(R)}$ (hence, all the other edges of $\mathcal{L}_{\mathcal{N}(R),C}$ are of the weight $\frac{\lambda}{2}$). Moreover, R is defined by the two edges of $\mathcal{L}_{\mathcal{N}(R),C}$ adjacent to $\hat{\varepsilon}$, and C is defined by $\hat{\varepsilon}$ and some other edge ε of this*

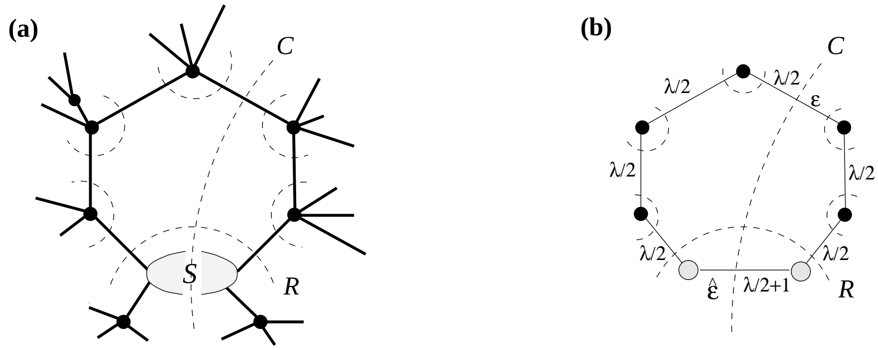

Figure 5: A structure that corresponds to a global cut and a basic cut crossing it: (a) the cycle $\mathcal{L}_{\mathcal{N}(R)}$ in $\mathcal{H}$ (shown by heavy lines); (b) the cycle $\mathcal{L}_{\mathcal{N}(R),C}$.

cycle not adjacent to $\hat{\varepsilon}$ (i.e., C is φ-induced by the cut $\{\hat{\varepsilon}, \varepsilon\}$ of $\mathcal{H}$).

Proof: Observe that, since R and C are crossing, they divide V into four parts, and the structure of their square is given by Lemma 3.10(ii). Assume, first, that this is the final partition into the $(F^{\lambda}_{\mathcal{N}(R)} \cup C)$-atoms. Then the resulting quotient graph $\mathcal{L}_{\mathcal{N}(R),C}$ is given by this Lemma and is as required. Clearly, the component at $\mathcal{N}(R)$ is the result of shrinking, in this graph, the two nodes that S is mapped into them; so, it is a cycle of length 3 with all edges of weight $\frac{\lambda}{2}$. By Lemma 3.8, the local model $\mathcal{L}_{\mathcal{N}(R)}$ is as required as well. The last two observations of the Lemma are straightforward.

Assume now that R and C do not partition V into the $(F^{\lambda}_{\mathcal{N}(R)} \cup C)$-atoms. Then, one of the two corner λ-cuts of the $\{R, C\}$-square, say R_1, is not separating and, so, is crossing with a λ-cut R' (see Fig. 6(a)). As any crossing pair of λ-cuts, R_1 and R' are cuts compatible with the same component $\hat{G}_{\mathcal{N}}$ of $\mathcal{T}^{\lambda}$, which is a uniform $\frac{\lambda}{2}$-cycle of length at least 4. Let us prove that the corresponding neighbor group $F^{\lambda}_{\mathcal{N}}$ coincides with $F^{\lambda}_{\mathcal{N}(R)}$, by showing that R belongs to $F^{\lambda}_{\mathcal{N}}$. If R, to the contrary, does not belong to the neighbor group $F^{\lambda}_{\mathcal{N}}$, then there exists a separating (w.r.t. F^{λ}) cut $R'' \in F^{\lambda}_{\mathcal{N}}$ between R and R_1 (see Fig. 6(b)). It is easy to see that any cut between R_1 and R must cross the second corner λ-cut R_2 of the $\{R, C\}$-square. But then R'' is not separating, a contradiction. Hence $F^{\lambda}_{\mathcal{N}} = F_{\mathcal{N}(R)}$.

Now, the above uniform cycle is the component $\hat{G}_{\mathcal{N}(R)}$, and the cuts R, R_1, and R_2 are compatible with it; the corresponding cycle $\mathcal{L}_{\mathcal{N}(R)}$ (with weights ignored) is implanted into $\mathcal{T}^{\lambda}$ instead of the node $\mathcal{N}(R)$. The cut R, as a separating one, is defined by two adjacent edges of $\hat{G}_{\mathcal{N}(R)}$, say, ε_1 and ε_2, and there is one more its edge, ε, such that the pairs $\varepsilon_1, \varepsilon$ and $\varepsilon_2, \varepsilon$ define the cuts R_1, and R_2, respectively. Now, the combination of the cycle $G_{\mathcal{N}(R)}$ and the $\{R, C\}$-square is the extended cycle $\mathcal{L}_{\mathcal{N}(R),C}$ as required. The last two observations

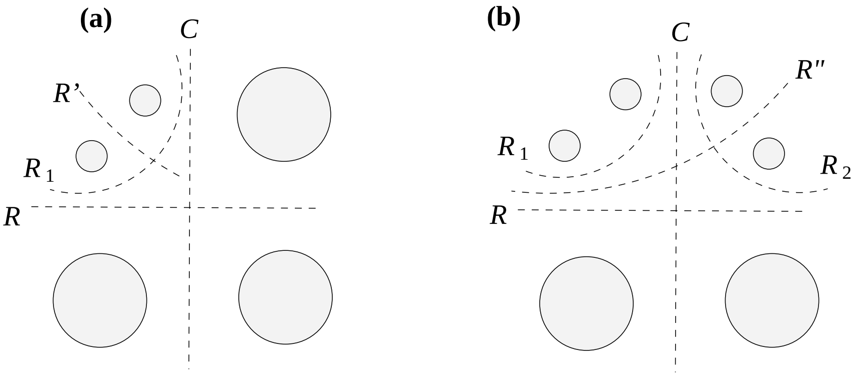

Figure 6: Illustration to the proof of Lemma 4.2.

of the Lemma follow immediately. □

Corollary 4.3 *Let a global cut C dividing a $(\lambda+1)$-class S, a basic cut $R \in F_S^\lambda$, and a cycle $\mathcal{L}_{\mathcal{N}(R)}$ be as in Lemma 4.2. Then each edge of C either has both endpoints in X_R or has both of them in $\bar{X}_R$. Moreover, there exists a structural edge $\varepsilon \in \mathcal{L}_{\mathcal{N}(R)}$ not incident to $\mathcal{N}_S$ such that the set of edges of C that have both endpoints in $\bar{X}_R$ coincides with E_ε.*

Let us estimate the number of cuts R as in Corollary 4.3 and establish the general structure generated by a global cut.

Lemma 4.4 *Let C be a global cut that divides a $(\lambda+1)$-class S. Then C is crossing with at most two cuts of F_S^λ.*

Proof: By Corollary 4.3, each cut R of F_S^λ which is crossing with C contributes to C a set of exactly $\frac{\lambda}{2}$ edges with both endpoints in $\bar{X}_R$; they are not edges of $\hat{G}_S$, since the entire $\bar{X}_R$ is mapped to a single halo node of $\hat{G}_S$. For any two distinct such cuts R' and R'', those edge sets are pairwise disjoint, since the edge endpoints lie in disjoint vertex sets $\bar{X}_{R'}$ and $\bar{X}_{R''}$ (see Fig. 7(a)). Let t denote the number of such cuts; then $t\frac{\lambda}{2} \leq \lambda+1$, which implies that if $\lambda \geq 4$ then $t \leq 2$. □

The following Corollary is implied by Corollary 4.3 and Lemma 4.4 (see for illustration Fig. 7).

Corollary 4.5 *Let C be a global cut that divides a $(\lambda+1)$-class S. Then there exists a cycle $\mathcal{L}$ of $\mathcal{H}$ incident to $\mathcal{N}_S$ (type 1) or exist two such cycles $\mathcal{L}_1$ and $\mathcal{L}_2$ (type 2) as follows.*

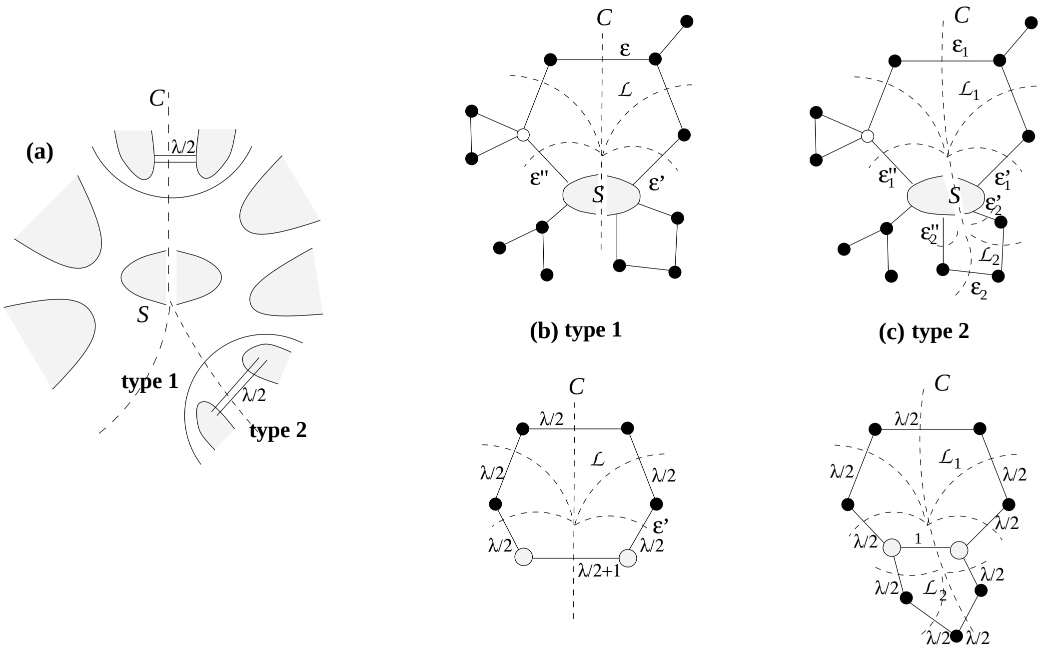

Figure 7: Bunch generated by a global cut C: (a) edges in $C \setminus E(\hat{G}_S)$; (b) type 1; (c) type 2.

The set $C \setminus E(\hat{G}_S)$ coincides with the set E_ε, where $\varepsilon \in \mathcal{L}$, or with the disjoint union of E_{ε_1} and E_{ε_2}, where $\varepsilon_1 \in \mathcal{L}_1$ and $\varepsilon_2 \in \mathcal{L}_2$, respectively, and ε, or each of $\varepsilon_1, \varepsilon_2$, respectively, is not incident to $\mathcal{N}_S$.

Let us gain more information about $(\lambda+1)$-cuts from the structures obtained. We consider a global cut C of the type 1 first (see Fig. 7(b)). Relying on the structure of quotient graph $\mathcal{L}_{\mathcal{N}(R),C}$ given by Lemma 4.2(ii), it is easy to see that the replacement of E_ε in C by $E_{\tilde{\varepsilon}}$, for any $\tilde{\varepsilon} \in \mathcal{L}_{\mathcal{N}(R)}$, results in a $(\lambda+1)$-cut of G as well. Moreover, such a cut is global if $\tilde{\varepsilon}$ is not incident to $\mathcal{N}_S$ in $\mathcal{H}$, and local otherwise. We call the group of all such $(\lambda+1)$-cuts a **bunch** of type 1. Clearly, any bunch can be easily reconstructed as described above given anyone of its cuts and the cycle $\mathcal{L}_{\mathcal{N}(R)}$. The bunch defined by a cut C and a cycle $\mathcal{L}$ is said to be **generated by C and $\mathcal{L}$**, or, briefly, generated by C (or by $\mathcal{L}$) if the second object is understood.

In particular, observe that we can use for such a reconstruction of a bunch anyone of its two local cuts. Such a pair—a local cut C and the corresponding cycle $\mathcal{L}$—can be identified in terms *local for the component $\hat{G}_S$* as follows. For any cycle $\mathcal{L}$ of $\mathcal{H}$ containing a node $\mathcal{N}_S$, let $b_\mathcal{L}$ denote the halo node of $\hat{G}_S$ corresponding to the branch of $\mathcal{T}^\lambda$ that contains the nodes of $\mathcal{L}$, except for $\mathcal{N}_S$. By Lemma 4.2, the set of edges incident to $b_\mathcal{L}$ is halved by $E_{\varepsilon'}, E_{\varepsilon''}$, where $\varepsilon', \varepsilon''$ are the structural edges of $\mathcal{L}$ incident to $\mathcal{N}_S$ in $\mathcal{H}$. A local cut C dividing S and a cycle $\mathcal{L}$ generate a bunch of type 1 if and only if $C \cap \delta(b_{\mathcal{L}})$ coincides with $E_{\varepsilon'}$ or with $E_{\varepsilon''}$.

When a global cut C is of type 2, then the modifications of C similar to that described for type 1 can be done independently for cycles $\mathcal{L}_1$ and $\mathcal{L}_2$ (see Fig. 7(c)). Indeed, a modification done for any one of the cycles affects, for the other one, only the corner $(\lambda+1)$-cuts of the square considered in the proof of Lemma 4.2; so, nothing in the structure given by Lemma 4.2 for the latter cycle is affected. In this case, the bunch of type 2 generated by C, $\mathcal{L}_1$, and $\mathcal{L}_2$, consists of all modifications of C resulting from replacement of E_{ε_1} and E_{ε_2} by any other $E_{\tilde{\varepsilon}_1}$ and $E_{\tilde{\varepsilon}_2}$, where $\tilde{\varepsilon}_1 \in \mathcal{L}_1$, $\tilde{\varepsilon}_2 \in \mathcal{L}_2$.

There are four local cuts in such a bunch; each of them, together with the cycles $\mathcal{L}_1$ and $\mathcal{L}_2$, can be used to generate the entire bunch. Such a local cut C and the cycles $\mathcal{L}_1$ and $\mathcal{L}_2$ are recognized, similarly to the case of type 1, as a cut C that halves properly each of $\delta(b_{\mathcal{L}_1})$ and $\delta(b_{\mathcal{L}_2})$ in $\hat{G}_S$, where $b_{\mathcal{L}_1}$ and $b_{\mathcal{L}_2}$ are the two halo nodes in $\hat{G}_S$ corresponding to $\mathcal{L}_1$ and $\mathcal{L}_2$, respectively.

Clearly, all bunches of the two described types cover the set of global $(\lambda+1)$-cuts.

For any two sets A and B, let $A \triangle B = (A \cup B) \setminus (A \cap B)$ denote the symmetric difference of A and B. We say that a partition of a set is **balanced** if its parts have equal cardinality.

We say that a halo node b of $\hat{G}_S$ is **distinguished** if there exists a $(\lambda+1)$-cut C' of $\hat{G}_S$ such that $|C' \cap \delta(b)| = \frac{\lambda}{2}$; such C' is said to be a **distinguishing cut**. Since $d(b) = \lambda$, the cut $C'' = C' \triangle \delta(b)$ has the same property, and we say that $\{C', C''\}$ is a **distinguishing pair**. Note that, in this case, the partition $\{\delta(b) \cap C', \delta(b) \cap C''\}$ of $\delta(b)$ is balanced.

The previous discussion implies the following property:

Lemma 4.6 *Let C be a local cut dividing a $(\lambda+1)$-class S and $\mathcal{L}$, $\mathcal{L}_1$, $\mathcal{L}_2$ be cycles of $\mathcal{H}$ containing $\mathcal{N}_S$. Then:*

(i) *C and $\mathcal{L}$ generate a bunch of type 1 if and only if C distinguishes in $\hat{G}_S$ the halo node $b_{\mathcal{L}}$ so that the arising balanced partition of $\delta(b_{\mathcal{L}})$ coincides with the partition induced by $\mathcal{L}$.*

(ii) *C and $\mathcal{L}_1$, $\mathcal{L}_2$ generate a bunch of type 2 if and only if C distinguishes in $\hat{G}_S$ each of the halo nodes $b_{\mathcal{L}_1}$ and $b_{\mathcal{L}_2}$ so that the arising balanced partitions of $\delta(b_{\mathcal{L}_1})$ and $\delta(b_{\mathcal{L}_2})$ coincide with the partitions induced by $\mathcal{L}_1$ and $\mathcal{L}_2$, respectively.*

As a consequence, any model for the local $(\lambda+1)$-cuts of a component $\hat{G}_S$, with additional information on all the cases when a local $(\lambda+1)$-cut halves prperly the set of edges incident

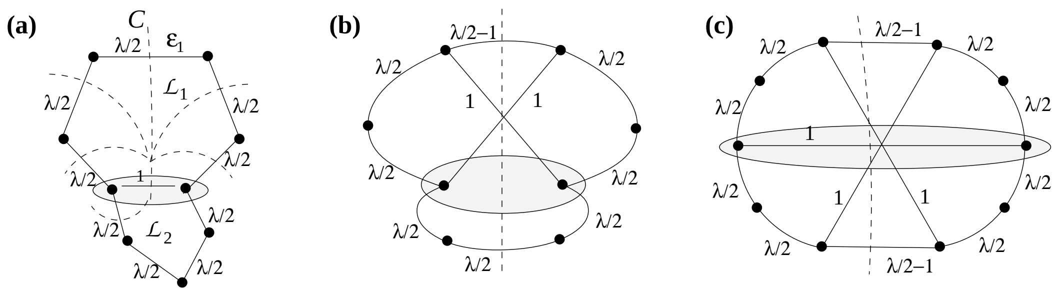

Figure 8: (a) a bunch of type 1 contained in a bunch of type 2; (b) a global cut of type 1 contained in 2 bunches; (c) a local cut C contained in 4 bunches.

to a halo node, can serve as an implicit model for *all* the $(\lambda+1)$-cuts dividing S. We use this reduction in what follows.

Let us conclude with some observations. Notice that each bunch of type 2 contains exactly four bunches of type 1: each one of them is obtained by fixing one among the four structural edges $\varepsilon_1', \varepsilon_1'' \in \mathcal{L}_1$ and $\varepsilon_2', \varepsilon_2'' \in \mathcal{L}_2$ incident to $\mathcal{N}_S$ in $\mathcal{H}$ (see Fig. 8(a), where ε_2'' is chosen). In what follows, we consider only the maximal-inclusion bunches (i.e., do not consider bunches of type 1 contained in a bunch of type 2), if not said the contrary.

Observe that, by the previous discussion, no two bunches share a global cut of type 2. However, distinct bunches can share a global cut of type 1 or a local cut (see examples in Fig. 8(b,c)).

Lemma 4.7 *A global $(\lambda+1)$-cut is contained in at most 2 bunches, and a local one in at most 4.*

Proof: Let us, first, prove an auxiliary statement: *any local $(\lambda+1)$-cut C distinguishes at most 4 halo nodes in its component.* Indeed, the contrary implies that some side of C contains at least 3 of such nodes; since the sets of edges incident to them in C are disjoint, $3\frac{\lambda}{2} \leq |C| = \lambda+1$ must hold, a contradiction if $\lambda \geq 4$.

By similar reasons, any global $(\lambda+1)$-cut halves the set of edges incident to a halo node for at most 2 halo nodes of $\hat{G}_S$. This property immediately implies the first statement of the Lemma.

For a local $(\lambda+1)$-cut C, if it generates only bunches of type 2, then one can imagine at most six pairs from the 4 nodes distinguished by it, each defining (together with C) a bunch of type 2. To reduce the bound six to four, let us fix an arbitrary halo node $b_{\mathcal{L}}$ distinguished by C. Let $\bar{C}$ be an arbitrary global cut generated by C and $\mathcal{L}$. Observe that any bunch of type 2 generated by C, $b_{\mathcal{L}}$, and any other halo node b induces a bunch of type 1 generated

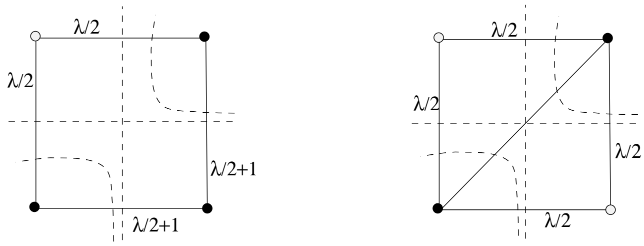

Figure 9: Two possible cases of crossing $(\lambda+1)$-cuts for λ even.

by $\bar{C}$ and b. By proved above, there can be at most two latter bunches, and thus of the former bunches as well. The total bound four for the current case follows.

If C generates a bunch of type 1 with exactly one of the nodes, then the three remaining nodes can form at most three pairs to generates bunches of type 2, totally four; the remaining cases can be proved similarly. □

4.2 Structure of local cuts

We study the local $(\lambda+1)$-cuts as the $(\lambda+1)$-cuts of the component $\hat{G}_S$; by Corollary 3.11, each of them divides S.

It turns out that the family F_S^{loc} of those cuts can be restored from a certain parallel cut family as follows. Crossings of such cuts have a very specific structure: flipping of a certain halo node from one side of any cut in a crossing pair to the other its side turns that crossing pair into a parallel cut pair (as in Fig 9). Moreover, for all crossing pairs that are "parallelized" using the same halo node, the changes in the edge set of a cut by such a flipping are the same. We modify the graph $\hat{G}_S$, by splitting-into-two each distinguished halo node, in a way, and replace each pair of cuts in F_S^{loc} that distinguish a halo node b by a single cut dividing the two new nodes that b is split into; the original cuts can be easily restored from the new cut. The new cut family of the new graph turns to be parallel, i.e., modeled by a tree. In such a way, the cuts in F_S^{loc} have a simple and compact but implicit representation; moreover, in the next Section, we proceed to an (explicit) cut model for the entire F based on those tree models.

We begin with establishing a specific structure of squares of crossing $(\lambda+1)$-cuts in $\hat{G}_S$.

Lemma 4.8 *Let C', C'' be a pair of crossing $(\lambda+1)$-cuts in a graph with even connectivity λ. Then, at least one of the corner cuts of their square is a λ-cut. Moreover, if C is such a corner cut, then both side edges in it are of weight $\frac{\lambda}{2}$, while its diagonal edge is of weight 0.*

Proof: Assume, in negation, that all the four corner cuts of C', C'' have cardinality strictly greater than λ, i.e., at least $\lambda+1$. Recall from Section 2 that in the $\{C', C''\}$-square holds $d_1+d_4 = |C'|+|C''|-2d_{23} = 2(\lambda+1)-2d_{23}$, and $d_2+d_3 = |C'|+|C''|-2d_{14} = 2(\lambda+1)-2d_{14}$. Thus, all the corner cuts are $(\lambda+1)$-cuts, and $d_{23} = d_{14} = 0$. One can easily verify that in this case each side edge of the $\{C', C''\}$-square must have the non-integral weight $\frac{\lambda+1}{2}$, a contradiction.

Let now C be a corner λ-cut in the $\{C', C''\}$-square, say $C = \delta(A_1)$. We show that $d_{12} = d_{13} = \frac{\lambda}{2}$. If this is not so, then assume, w.l.o.g., that $d_{12} < \frac{\lambda}{2}$. But then $d(A_2) = d(A_1 \cup A_2) - (d_{13} + d_{14}) + d_{12} = (\lambda+1) - (\lambda - d_{12}) + d_{12} = 2d_{12} + 1 \leq 2(\frac{\lambda}{2} - 1) + 1 < \lambda$, a contradiction. $\Box$

In fact, there exist only two possible cases of the square of cuts as in Lemma 4.8, which are presented in Fig. 9 (this fact can be easily established, and this leads to an alternative proof of Lemma 4.8).

Recall that $\hat{G}_S$ is λ-connected and that the only its λ-cuts are the cuts defined by single halo nodes. Therefore, any corner cut as in Lemma 4.8 is defined by a halo node; such a halo node b is distinguished by any of the crossing cuts, and the arising balanced partition of $\delta(b)$ is defined by the side edges of the square incident to b. Observe that flipping of b to the other side of anyone of the crossing cuts turns it into a corner cut, i.e., turns the crossing pair into a parallel one.

The following fact is crucial.

Lemma 4.9 *If a halo node b is distinguished by several pairs of $(\lambda+1)$-cuts, then the arising balanced partitions of $\delta(b)$ coincide.*

Proof: Suppose, in negation, that there is a halo node b which is distinguished by two pairs of $(\lambda+1)$-cuts, say $\{C_1, C_3\}$ and $\{C_2, C_4\}$, such that the arising partitions of $\delta(b)$ do not coincide (see Fig. 10(a)). The negation assumption implies that there are exactly five $\{C_1, C_2, C_3, C_4\}$-atoms, and one of them is $\{b\} = X_0$. Denote the other atoms by X_1, X_2, X_3, X_4, where $\delta(X_i \cup X_{i+1}) = C_i$, $i = 1, 2, 3$, and $\delta(X_4 \cup X_1) = C_4$. Then holds:

$$5\lambda \leq \sum_{i=0}^{4} d(X_i) = 2 \sum_{0 \leq i < j \leq 4} d(X_i, X_j) = \sum_{i=1}^{4} |C_i| - 2(d(X_1, X_3) + d(X_2, X_4)) \leq 4(\lambda+1),$$

which is possible only if $\lambda = 4$ and $d(X_1, X_3) = d(X_2, X_4) = 0$. Clearly, in this case $d(X_0, X_i) = 1$, $i = 1, 2, 3, 4$, must hold. Further computations lead to $d(X_i, X_j) = \frac{3}{2}$ for any $i \neq j = 1, 2, 3, 4$ (see Fig. 10(b)), which is impossible for an unweighted graph. $\Box$

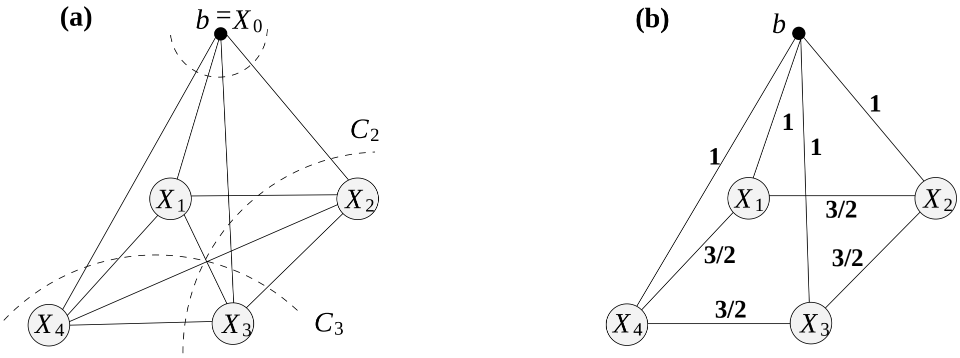

Figure 10: Canonicity of the partition of edges incident to a distinguished node b ((a) $\lambda > 4$, (b) $\lambda = 4$).

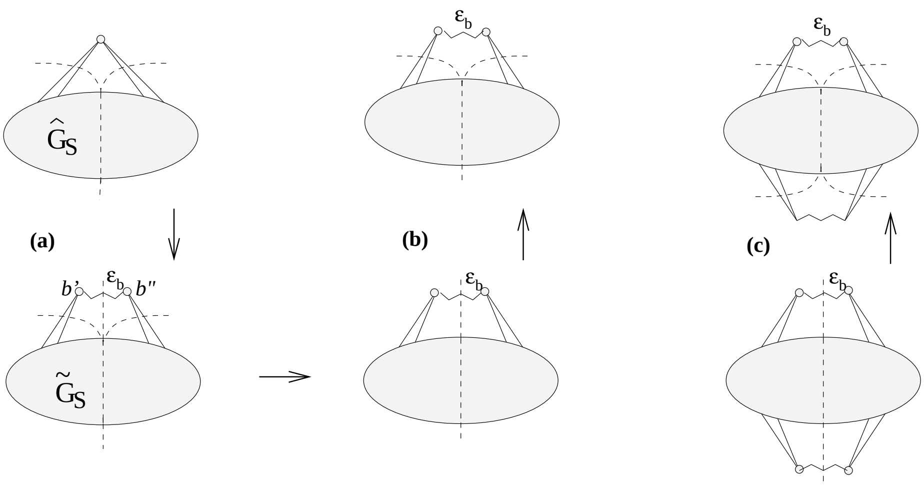

Figure 11: (a) a distinguished node and its split pair; (b) local cuts restored from a $(\lambda+1)$-cut that contains a single split edge; (c) local cuts restored from a $(\lambda+1)$-cut that contains two split edges.

For any distinguished halo node b, let us denote the canonical balanced partition of $\delta(b)$ by $\{\delta'(b), \delta''(b)\}$. The **split graph** $\tilde{G}_S$ of $\hat{G}_S$ is defined as follows (see Fig. 11(a)). Each distinguished node b is replaced by the two **split nodes** b', b'' (the **split pair** of b) and the **split edge** $\varepsilon_b = (b', b'')$ of the weight/multiplicity $\frac{\lambda}{2}$; each edge $e = (b, v)$ is replaced by (b', v) if $e \in \delta'(b)$ or by (b'', v) if $e \in \delta''(b)$.

Clearly, $\hat{G}_S$ is a quotient graph of $\tilde{G}_S$, obtained by shrinking every split pair of $\tilde{G}_S$ into a single node. Each non-split edge (resp., non-split node) of $\tilde{G}_S$ naturally corresponds to an edge (resp., a nondistinguished node) of $\hat{G}_S$, and vice versa; in what follows we identify such corresponding objects in $\tilde{G}_S$ and $\hat{G}_S$. Similarly, we identify each cut of $\tilde{G}_S$ that does not contain any split edge with the corresponding cut of $\hat{G}_S$. In this way, the $(\lambda+1)$-cuts of $\tilde{G}_S$ that do not contain any split edge are the cuts in F_S^{loc}, which we intend to model. Any other $(\lambda+1)$-cut of $\tilde{G}_S$ contains one or two split edges and is called a **split cut**.

The following statement, concerning general graphs, can be easily verified.

Fact 4.10 *Let C be a cut and v be a vertex such that $C \neq \delta(v)$, in a graph. Then, $C' = C \triangle \delta(v)$ is a cut, and $|C'| = |C| + d(v) - 2|C \cap \delta(v)|$. Hence, if $d(v)$ is even and $|C \cap \delta(v)| > \frac{d(v)}{2}$, then $|C'| \leq |C| - 2$.*

Let us now establish some important properties of the split graph $\tilde{G}_S$.

Lemma 4.11 (i) *$\tilde{G}_S$ is λ-connected;*

(ii) *R is a λ-cut of $\tilde{G}_S$ if and only if it is defined either by a halo node, or by a split node, or by a split pair.*

Proof: (i) Among minimum cuts of $\tilde{G}_S$, let R be one that contains a minimum number of split edges. Assume that R contains at least one split edge, say ε_b. If $R = \delta(b')$ then $|R| = \lambda$ and we are done. Otherwise, by Fact 4.10, the cut $R' = R \triangle \delta(b')$ is also a minimum cut of $\tilde{G}_S$ and contains less split edges than R, a contradiction. Therefore, R is a cut of $\hat{G}_S$ and, hence, $|R| \geq \lambda$ as required.

(ii) Clearly, any cut defined by a halo node, or by a split node, or by a split pair is a λ-cut. Assume, in negation, that there exists a λ-cut of $\tilde{G}_S$ of another type. Among such cuts, let R be one that contains the minimum number of split edges. Observe that R contains at least one split edge, say ε_b. Indeed, otherwise R is a λ-cut of $\hat{G}_S$, hence, it must be defined by a halo node in $\hat{G}_S$, i.e., it must be defined by a halo node or by a split pair in $\tilde{G}_S$.

Since $\tilde{G}_S$ is λ-connected, then by Fact 4.10, each of $R' = R \triangle \delta(b')$, $R'' = R \triangle \delta(b'')$ is a λ-cut of $\tilde{G}_S$ that contains less split edges than R. We obtain a contradiction to the definition of R by showing that at least one of R', R'' is defined neither by a halo node, nor by a split node, nor by a split pair. By Corollary 3.11(ii), S is nonempty; let $s \in S$. By construction, one of R', R'' contains s and both b', b'' on the same its side, while the other its side is nonempty. Then, the other one of R', R'' contains s on its one side, and at least three nodes on its other side, i.e., is as required, which finishes the proof. □

Corollary 4.12 (i) *Let C be a $(\lambda+1)$-cut of $\tilde{G}_S$ containing a split edge $\varepsilon_b = (b', b'')$. Then, $C \cap \delta(b') = \varepsilon_b$. Hence, $C' = C \triangle \delta(b')$ is also a $(\lambda+1)$-cut.*

Moreover, for any node $v \neq b''$ of $\tilde{G}_S$ of degree λ, $|C \cap \delta(v)| = \frac{\lambda}{2}$ implies $|C' \cap \delta(v)| = \frac{\lambda}{2}$.

(ii) *For any* $(\lambda+1)$*-cut* C *and any* halo *node* $\tilde{b}$ *of* $\tilde{G}_S$, $|C \cap \delta(\tilde{b})| < \frac{\lambda}{2}$.

Proof: Since $\tilde{G}_S$ is λ-connected, by Fact 4.10, no $(\lambda+1)$-cut of $\tilde{G}_S$ contains more than $\frac{\lambda}{2}$ edges incident to a node of degree λ.

(i) The above property implies straightforwardly the first part of (i), and, observing that $C' \cap \delta(v)$ contains $C \cap \delta(v)$, also the second its part.

(ii) Assume, in negation, that C contains exactly $\frac{\lambda}{2}$ edges from $\delta(\tilde{b})$. W.l.o.g., let C contain the minimal number of split edges among all such cuts. Observe that C contains at least one split edge, say $e_b = (b', b'')$. Indeed, otherwise C is a cut of $\hat{G}_S$, therefore, $\tilde{b}$ is a distinguished node in $\hat{G}_S$, and, hence, $\tilde{b}$ cannot be a halo node in $\tilde{G}_S$.

Now, by (i), $C' = C \triangle \delta(b')$ is a $(\lambda+1)$-cut that contains $\frac{\lambda}{2}$ edges from $\delta(b)$ and contains less split edges than C, a contradiction. □

Let $\tilde{F}_S^{\lambda+1}$ denote the family of all $(\lambda+1)$-cuts of $\tilde{G}_S$ that do not contain all the $\frac{\lambda}{2}$ non-split edges incident to any split node. In other words, a $(\lambda+1)$-cut $\tilde{C}$ belongs to $\tilde{F}_S^{\lambda+1}$ if and only if for every split node b' such that $|\tilde{C} \cap \delta(b')| = \frac{\lambda}{2}$ holds $\tilde{C} \cap \delta(b') = \varepsilon_b$, where ε_b is the split edge incident to b'. Let us see what is the relation between $\tilde{F}_S^{\lambda+1}$ and F_S^{loc}.

Let us consider a split cut $\tilde{C}$ in $\tilde{F}_S^{\lambda+1}$; recall that any split cut contains either one or two split edges. If $\tilde{C}$ contains a single split edge ε_b, then, by Corollary 4.12(i), each of the two cuts $\tilde{C}\triangle\delta(b')$ and $\tilde{C}\triangle\delta(b'')$ is a $(\lambda+1)$-cut cut of $\hat{G}_S$ distinguishing b (see Fig. 11(b)). Similarly, by Corollary 4.12(i), if it contains two split edges ε_{b_1} and ε_{b_2}, then each of the four cuts $\tilde{C}\triangle\delta(b_1')\triangle\delta(b_2')$, $\tilde{C}\triangle\delta(b_1'')\triangle\delta(b_2')$, $\tilde{C}\triangle\delta(b_1')\triangle\delta(b_2'')$, $\tilde{C}\triangle\delta(b_1'')\triangle\delta(b_2'')$ is a $(\lambda+1)$-cut distinguishing both b_1 and b_2 in $\hat{G}_S$ (see Fig. 11(c)). All those cuts are said to be **restored** from $\tilde{C}$.

It is easy to see, using Corollary 4.12(ii), that a $(\lambda+1)$-cut of $\tilde{G}_S$ belongs to both F_S^{loc} and $\tilde{F}_S^{\lambda+1}$ if and only if it is a nondistinguishing cut of $\hat{G}_S$. Such a cut is said to be restored from itself.

Lemma 4.13 *The family of cuts restored from the cuts in* $\tilde{F}_S^{\lambda+1}$ *coincides with* F_S^{loc}.

Proof: The fact that all restored cuts are in F_S^{loc} was already indicated above. It remains to show that for any cut $C \in F_S^{\lambda+1}$ there exists a cut in $\tilde{F}_S^{\lambda+1}$ such that C is restored from it.

If C is nondistinguishing in $\hat{G}_S$, then it is restored from itself. Now, let C distinguish in $\hat{G}_S$ a halo node, say b_1. Let b_1' be the split node of b_1 for which in $\tilde{G}_S$ holds $|C \cap \delta(b_1')| = \frac{\lambda}{2}$. Now, $C_1 = C\triangle\delta(b_1')$ is a $(\lambda+1)$-cut of $\tilde{G}_S$, and $C = C_1\triangle\delta(b_1')$. Therefore, if $C_1 \in \tilde{F}_S^{\lambda+1}$, then C is restored from C_1 and we are done. Otherwise, the definition of $\tilde{F}_S^{\lambda+1}$ (**??**) implies that there exists a halo node $b_2 \neq b_1$ in $\hat{G}_S$ such that in $\tilde{G}_S$, for one of b'_2, b''_2, say b'_2, holds: $|C_1 \cap \delta(b'_2)| = \frac{\lambda}{2}$ and $C_1 \cap \delta(b'_2) \neq e_{b_2}$. The cut $C_2 = C_1 \triangle \delta(b'_2)$ is a $(\lambda+1)$-cut which contains two split edges; evidently, it belongs to $\tilde{F}_S^{\lambda+1}$. Now, $C = C_2 \triangle \delta(b'_1) \triangle \delta(b'_2)$, which finishes the proof. □

Lemma 4.6 and Lemma 4.13 imply the following 2-stage reduction:

Reduction:

(i) *Given a description of the $(\lambda+1)$-cuts of $\hat{G}_S$, the global $(\lambda+1)$-cuts dividing S are* generated *by certain distinguishing cuts among them.*

(ii) *Given a description of the cuts in $\tilde{F}_S^{\lambda+1}$, the $(\lambda+1)$-cuts of $\hat{G}_S$ are* restored *from them.*

In such a way, the family $\tilde{F}_S^{\lambda+1}$ represents the family of all $(\lambda+1)$-cuts dividing S. (Observe that, though the graph $\tilde{G}_S$ serves to represent those cuts, it is not a model in the usual sense of this paper.)

Let us show that the family $\tilde{F}_S^{\lambda+1}$ is parallel. Lemmas 4.8 and 4.11(ii) and Corollary 4.12(ii) imply that at least one corner cut of a crossing pair of $(\lambda+1)$-cuts of $\tilde{G}_S$ is defined by a split node or by a split pair. It is easy to see that in both cases at least one of the crossing cuts contains exactly $\frac{\lambda}{2}$ non-split edges incident to some split node, and, hence, does not belong to $\tilde{F}_S^{\lambda+1}$. This implies the desired parallel property of $\tilde{F}_S^{\lambda+1}$.

Hence, for the family $\tilde{F}_S^{\lambda+1}$ exists a tree model $(\tilde{\mathcal{T}}_S, \tilde{\varphi}_S)$, where $\tilde{\varphi}_S$ maps $V(\tilde{G}_S)$ to $\mathcal{V}(\tilde{\mathcal{T}}_S)$. The model $(\tilde{\mathcal{T}}_S, \tilde{\varphi}_S)$ represents, via the above Reduction, all the $(\lambda+1)$-cuts of $\hat{G}_S$; moreover, by Corollary 3.11, the set of models $\{(\tilde{\mathcal{T}}_S, \tilde{\varphi}_S) : S \text{ is a } (\lambda+1)\text{-class of } G\}$ represents all the $(\lambda+1)$-cuts of G. In the next Section we show how this representation can be "unwrapped", united over all $(\lambda+1)$-classes of G, and implemented in a compact way.

4.3 2-skeleton model and the incremental maintenance

Now we unwrap Reduction(ii), arriving at a 2-skeleton for $\hat{G}_S$ (for illustration follow Fig. 12). For convenience of explanation, let us assign colors to structural edges of our model. First, we assign to structural edges of $\tilde{\mathcal{T}}_S$ the **blue** color. Second, for each halo node b of $\hat{G}_S$, we add to $\tilde{\mathcal{T}}_S$ the (halo) node denoted also by b and two **red** edges of weight 1 each: $(b, \tilde{\varphi}_S(b'))$

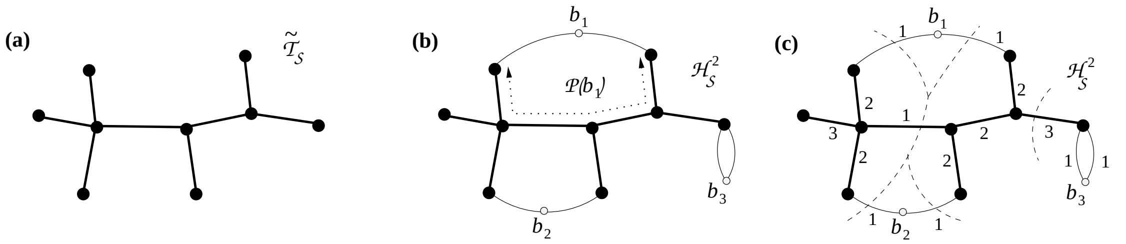

Figure 12: Construction of a 2-skeleton for $\hat{G}_S$: (a) the tree $\tilde{\mathcal{T}}_S$; (b) $\mathcal{H}_S^2$; (c) restoration of the distinguishing cuts. (Blue edges are shown thicker than red ones.)

and $(b, \tilde{\varphi}_S(b''))$, if b is a distinguished node, and two parallel edges $(b, \tilde{\varphi}_S(b))$, otherwise. For the graph $\mathcal{H}_S^2$ obtained, each its cut defined by a single halo node is a 2-cut. The mapping $\varphi_S^2 : V(\hat{G}_S) \to \mathcal{V}(\mathcal{H}_S^2)$ takes the vertices of S as $\tilde{\varphi}_S$ does and each halo node b of $\hat{G}_S$ to the node b of $\mathcal{H}_S^2$ (in what follows we identify them). Clearly, the pair of red edges incident to a halo node b φ_S^2-induces the λ-cut corresponding to b.

Let us call the unique blue path $\mathcal{P}(b)$ that connects the ends of the red edges incident to a halo node b the **projection path** of b. Observe that a structural edge belongs to $\mathcal{P}(b)$ if and only if it $\tilde{\varphi}_S$-induces a cut of $\tilde{G}_S$ dividing $\{b', b''\}$, i.e., containing e_b.

Lemma 4.14 (i) *A structural edge of $\mathcal{H}_S^2$ is contained in at most two projection paths.*

(ii) *Two projection paths have at most one structural edge of $\mathcal{H}_S^2$ in common.*

Proof: For proving (i) observe that if ε is a blue structural edge of $\mathcal{H}_S^2$, then for each projection path $P(b)$ containing ε, the split edge e_b is contained in the $(\lambda+1)$-cut of $\tilde{G}_S$ corresponding to ε. Those split edges are distinct, and each of them has the weight $\frac{\lambda}{2}$. Since $\lambda \geq 4$, the number of such edges is at most two, as required.

We now prove part (ii). Suppose, in negation, that two projection paths $\mathcal{P}_1$ and $\mathcal{P}_2$ have at least two structural edges in common. Since $\mathcal{P}_1$ and $\mathcal{P}_2$ are paths in a tree, their intersection is a path of length at least two. Let $(\mathcal{N}_1, \mathcal{N}), (\mathcal{N}, \mathcal{N}_2)$ be two adjacent edges of this intersection path. Remove $(\mathcal{N}_1, \mathcal{N})$ and $(\mathcal{N}, \mathcal{N}_2)$, and let $\mathcal{B}$ be the connected part of $\tilde{\mathcal{T}}_S$ containing $\mathcal{N}$ (see Fig. 13(b)). Since the $(\lambda+1)$-cuts of $\tilde{G}_S$ that are $\tilde{\varphi}_S$-induced by $(\mathcal{N}_1, \mathcal{N})$ and $(\mathcal{N}, \mathcal{N}_2)$ are distinct, the set $\tilde{\varphi}_S^{-1}(\mathcal{B})$ is nonempty (see Fig. 13(a)). Since those cuts has two split edges of weight $\frac{\lambda}{2}$ in common, $\tilde{\varphi}_S^{-1}(\mathcal{B})$ defines a cut of $\tilde{G}_S$ of cardinality at most 2, a contradiction.

□

We assign to any blue structural edge ε of $\mathcal{H}_S^2$ the weight $w(\varepsilon)$ equal to 3 minus the

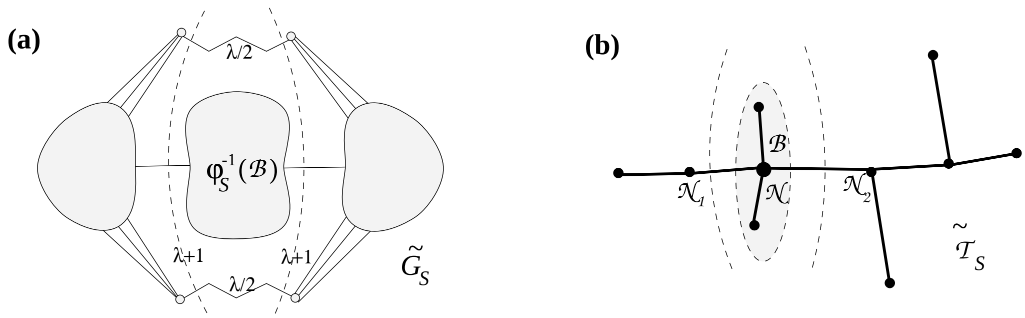

Figure 13: Impossibility of a long intersection of two projection paths.

number of projection paths containing ε (Fig. 12(c)). Taking into account these weights, observe that any blue edge ε, together with either zero, or one, or two red edges: one edge from each red pair that defines a projection path containing ε, is a 3-cut of $\mathcal{H}_S^2$; thus, ε is contained in either one, or two, or four such cuts, respectively. It is easy to see that those cuts φ_S^2-induce exactly the $(\lambda+1)$-cuts of $\hat{G}_S$ restored from the cut of $\tilde{G}_S$ defined by ε.

Lemma 4.15 (i) *$\mathcal{H}_S^2$ is 2-connected; it has $O(|V(\hat{G}_S)|)$ structural edges.*

(ii) *The model $(\mathcal{H}_S^2, \varphi_S^2)$ is a 2-skeleton for $\hat{G}_S$.*

Proof: (i) By construction, there are no bridges in $\mathcal{H}_S^2$ when weights ignored, except for the blue edges that do not belong to any projection path. Since each such edge has weight 3, there are no 1-cuts in $\mathcal{H}_S^2$.

By Section 3.1, the size of the tree $\mathcal{T}_S$, i.e., the number of blue edges, is linear in $|V(\hat{G}_S)|$. The number of red edges is at most twice the number of projection paths, and the latter is at most twice the number of blue edges, by Lemma 4.14.

(ii) The condensity property of $\mathcal{H}_S^2$ follows from the simple observation that the 3-cuts of $\mathcal{H}_S^2$ as in the discussion before the Lemma, together with 2-cuts defined by single halo nodes, partition the node set of $\mathcal{H}_S^2$ into singletons.

Now, the only thing remained is to show that $(\varphi_S^2)^{-1}$ takes the sets of 2- and 3-cuts of $\mathcal{H}_S^2$ onto the sets of λ- and $(\lambda+1)$-cuts of $\hat{G}_S$, respectively. From the previous discussion follows that for any λ- or $(\lambda+1)$-cut of $\hat{G}_S$ there exists a 2- or 3-cut of $\mathcal{H}_S^2$, respectively, φ_S^2-inducing it. It remains to prove that any 2- or 3-cut of $\mathcal{H}_S^2$ φ_S^2-induces λ- or $(\lambda+1)$-cut of $\hat{G}_S$, respectively.

First, let us analyze a cut C consisting of red edges only. Observe that the entire tree formed by the blue edges is on the same side of C. The other side consists only from halo

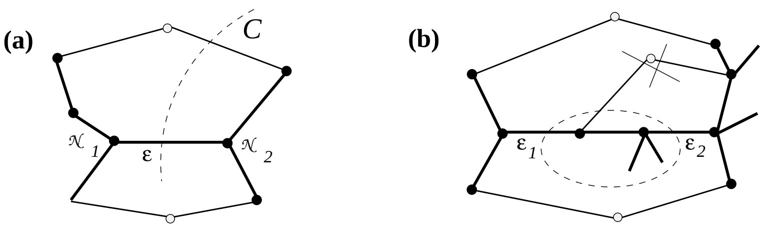

Figure 14: No 2-cut of $\mathcal{H}_S^2$ contains a blue edge.

nodes b, so that both the two edges incident to b are in C. Therefore, if $|C| = 2$ then C is defined by a single halo node; otherwise, $|C| \geq 4$.

Now we prove that no 2-cut C of $\mathcal{H}_S^2$ contains a blue edge. Let us consider two cases. Assume, first, that C consists of a blue edge $\varepsilon = (\mathcal{N}_1, \mathcal{N}_2)$ of weight 1 and a single red edge Since $w(\varepsilon) = 1$, there are two projection paths containing ε. There exists at least one of those paths such that the both red edges corresponding to it do not belong to C. Then, $\mathcal{N}_1$ and $\mathcal{N}_2$ are connected by the edges of that path, excluding ε, and those red edges (see Fig. 14(a)), a contradiction.

Second, assume that C consists of two blue edges ε_1, ε_2 of weight one each. Deletion of them breaks $\mathcal{T}_S$ into three parts: a "middle" one and two "terminal" ones (see Fig. 14(b)). No projection path containing anyone of ε_1, ε_2 connects the middle and a terminal parts; indeed, otherwise, the two red edges corresponding to that path connect the two sides of C. Since both $w(\varepsilon_1)$, $w(\varepsilon_2$ are 1, there must be at least two projection paths containing ε_1 or ε_2, and, hence, connecting the two terminal parts. Both those paths contain both of ε_1 and ε_2, a contradiction to Lemma 4.14.

We proceed to 3-cuts. Let us show that any 3-cut of $\mathcal{H}_S^2$ contains exactly one blue edge. It was already shown that three red edges cannot form a 3-cut. Let $\mathcal{C}$ be a 3-cut of $\mathcal{H}_S^2$, and let $\varepsilon_1, \ldots, \varepsilon_r$, $1 \leq r \leq 3$, be the blue edges in C. Assume, in negation, that r equals 2 or 3. Deletion of $\varepsilon_1, \ldots, \varepsilon_r$ partitions the tree $\tilde{\mathcal{T}}_S$ into $r + 1$ subtrees $\mathcal{T}_0, \ldots, \mathcal{T}_r$ (see Fig. 15). It can be easily shown that the edges ε_j and the subtrees $\mathcal{T}_i$ can be renumbered so that there exist nodes $\mathcal{N}_i$, $i = 0, \ldots, r$, such that $\mathcal{N}_i$ belongs to $\mathcal{T}_i$ and, for $i \geq 1$, is incident to ε_i; w.l.o.g., assume that our numbering is such a one. Let $\mathcal{Q}_S^2$ be the quotient graph obtained from $\mathcal{H}_S^2$ by shrinking every subtree $\mathcal{T}_i$ into the node $\mathcal{N}_i$, $i = 0, \ldots, r$; let $\mathcal{X}, \bar{\mathcal{X}}$ be the sides of C in $\mathcal{Q}_S^2$, where $\mathcal{N}_0 \in \mathcal{X}$. Observe that

(a) C is compatible with $\mathcal{Q}_S^2$ and contains all its r blue edges;

(b) since C is of odd cardinality 3, each of $\mathcal{X}, \bar{\mathcal{X}}$ contains at least one node of odd degree;

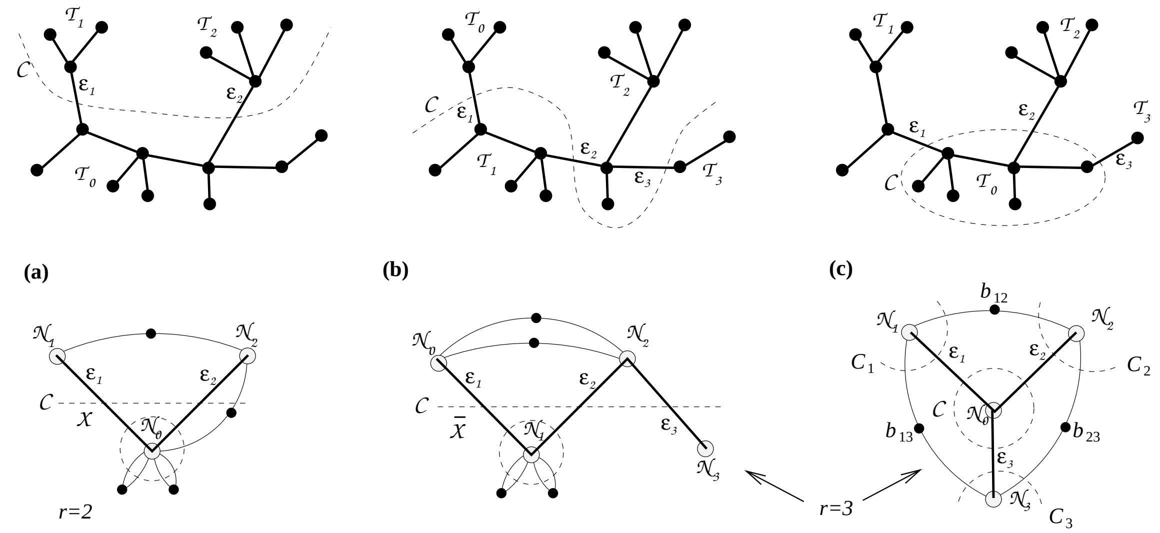

Figure 15: Illustration to the proof of Lemma 4.15 (upper parts show divisions of $\mathcal{T}_S$ and lower parts show graphs $\mathcal{Q}_S^2$).

since all nodes in $\mathcal{Q}_S^2$, except of $\mathcal{N}_0, \ldots, \mathcal{N}_r$, are halo and, thus, have degree 2 each, such a node must be one of $\mathcal{N}_0, \ldots, \mathcal{N}_r$;

(c) for any projection path with one endnode in $\mathcal{X}$ and the other in $\bar{\mathcal{X}}$, the cut C contains exactly one red edge of the pair of red edges corresponding to this path.

Let us consider two cases:

$r = 2$: W.l.o.g., let $\mathcal{T}_0$ be the "middle" subtree (see Fig. 15(a)). Let $q \in \{0, 1\}$ be the number of projection paths containing both ε_1 and ε_2 (those are the projection paths having one endnode in $\mathcal{T}_1$ and the other in $\mathcal{T}_2$). It can be easily verified (see Fig. 15(a)) that

$$d(\mathcal{N}_0) = 2(|\mathcal{X}| - 1) + 6 - 2q,$$

a contradiction, since $d(\mathcal{N}_0)$ must be odd.

$r = 3$: In this case, C contains no red edges, while each of the edges $\varepsilon_1, \varepsilon_2, \varepsilon_3$ has the weight 1 and thus is contained in exactly two projection paths. Clearly, one side of C contains at least two nodes from $\mathcal{N}_0, \mathcal{N}_1, \mathcal{N}_2, \mathcal{N}_3$; assume, w.l.o.g., that $\mathcal{N}_1, \mathcal{N}_3 \in \bar{\mathcal{X}}$. Let us consider two possibilities:

$\mathcal{N}_2 \in \mathcal{X}$: W.l.o.g., $\varepsilon_i = (\mathcal{N}_{i-1}, \mathcal{N}_i)$, $i = 1, 2, 3$ (see Fig. 15(b)). Since C does not contain any red edge, each of the two projection paths containing ε_1 must have one endnode in $\mathcal{T}_0$ and the other in $\mathcal{T}_2$. However, then both $\varepsilon_1, \varepsilon_2$ are contained in each of those two projection paths, a contradiction to Lemma 4.14(ii).

$\mathcal{N}_2 \in \bar{\mathcal{X}}$: In this case $\varepsilon_i = (\mathcal{N}_0, \mathcal{N}_i)$, $i = 1,2,3$ (see Fig. 15(c)). Let us consider the projection paths containing at least one of $\varepsilon_1, \varepsilon_2, \varepsilon_3$. Using arguments as in the case $\mathcal{N}_2 \in \mathcal{X}$, it can be easily verified that there are exactly three such paths: one contains $\varepsilon_1, \varepsilon_2$, another $\varepsilon_2, \varepsilon_3$, and the last one $\varepsilon_1, \varepsilon_3$. Let us denote by b_{ij}, $i < j = 1,2,3$, the halo node corresponding to the projection path containing $\varepsilon_i, \varepsilon_j$. Let $\mathcal{C}_i$ be the 3-cut of $\mathcal{H}_S^2$ which consists of the edge ε_i and the two red edges $(\mathcal{N}_i, b_{ij})$, and let $\mathcal{X}_i$ be the part of $\mathcal{C}_i$ containing $\mathcal{T}_i$, $i = 1,2,3$. Recall that each $\mathcal{C}_i$, $i = 1,2,3$, φ_S^2-induces a $(\lambda+1)$-cut of $\hat{G}_S$. Let us denote $X_i = (\varphi_S^2)^{-1}(\bar{\mathcal{X}}_i)$, $i = 1,2,3$, and $X_0 = (\varphi_S^2)^{-1}(\mathcal{X})$. By the construction, for $j > i = 1,2,3$ holds: $d(X_i) = \lambda + 1$, $d(b_{ij}) = \lambda$, and $d(b_{ij}, X_i) = d(b_{ij}, X_j) = \frac{\lambda}{2}$ **(??)**. Hence, $d(b_{ij}, X_0) = 0$, $d(X_i, X_0) \leq 1$. Now, $d(X_0) = 3 - 2 \cdot \Sigma_{i<j} d(X_i, X_j)$. This implies that $d(X_0)$ is equal to 1 or 3, hence, X_0 is nonempty; therefore, X_0 defines a 1- or 3-cut in a graph with connectivity $\lambda \geq 4$, a contradiction.

So, each 3-cut of $\mathcal{H}_S^2$ contains exactly one blue edge. It is easy to see that the set of red edges belonging to a 3-cut of $\mathcal{H}_S^2$ containing a single blue edge ε is formed by taking exactly one edge from each red pair that defines a projection path containing ε. By the previous discussion, any such cut φ_S^2-induces a $(\lambda+1)$-cut, which finishes the proof. □

Let us analyze the obtained representation. Observe that the structure of the 2- and 3-cuts of each skeleton model $(\mathcal{H}_S^2, \varphi_S^2)$ is simple: those are, respectively, the cuts defined by the single halo nodes and the cuts defined by the single blue edges as described before Lemma 4.15. The totality of those models, together with $(\mathcal{H}, \varphi)$, represents all the λ- and $(\lambda+1)$-cuts of G, via Reduction(i) established for the global cuts in Section 4.1. This representation can be implemented in a compact way if we keep, (i) for each $(\lambda+1)$-class S, a pointer from the node $\varphi(S)$ of $\mathcal{H}$ to the model $(\mathcal{H}_S^2, \varphi_S^2)$, (ii) for each halo node b of $\mathcal{H}_S^2$, a pointer from it to the corresponding cycle $\mathcal{L}(b)$ (possibly of the length 2) of $\mathcal{H}$, and, (iii) for some of the halo nodes b, a flag indicating that b is distinguished and the partition $\{\delta'(b), \delta''(b)\}$ of $\delta(b)$ coincides with its partition induced by $\mathcal{L}(b)$.

In this implementation, the bunches are represented as follows. A bunch of type 2 is defined by a blue edge ε of weight 1, such that each of the two corresponding halo nodes (whose projection paths contain ε) has the flag as above. A bunch of type 1 is defined either by a blue edge of weight 1, such that exactly one of those halo nodes has this flag, or by a blue edge of weight 2, such that the unique corresponding halo node has this flag.

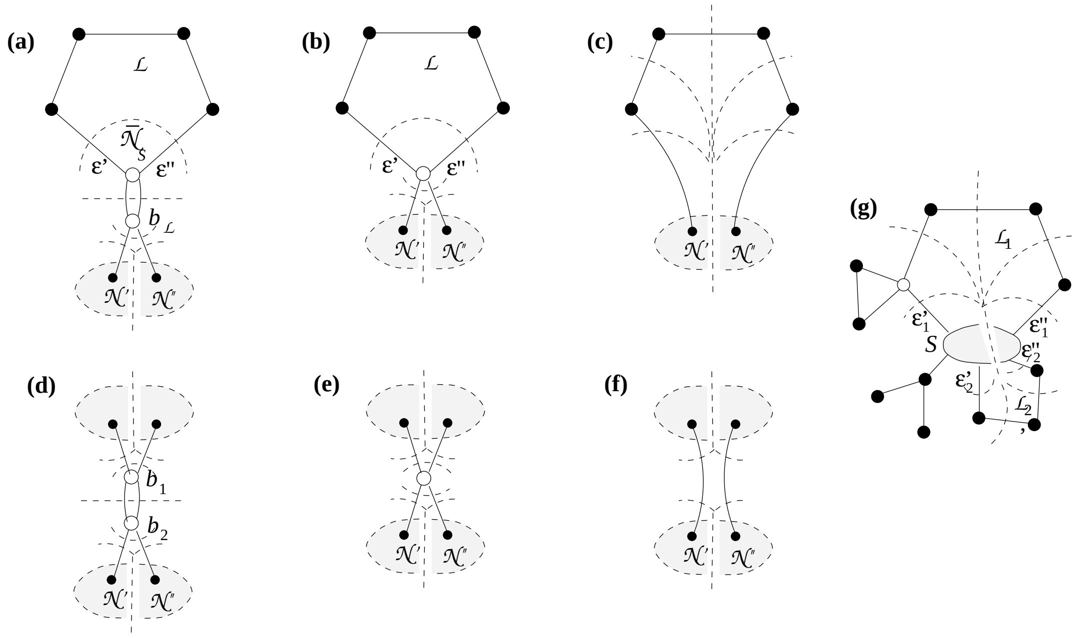

Figure 16: Illustration to the opening procedure.

By Lemma 4.15(i), the size of the considered representation is $O(n_{\lambda+2}) = O(n)$, since the sum of the numbers of nodes in all $(\lambda+1)$-components is $O(n_{\lambda+2})$ (indeed, as can be easily verified, there are, totally, at most $n_{\lambda+2}$ nonhalo nodes and $2 \cdot |\mathcal{V}(\mathcal{T}^\lambda)| = O(n_{\lambda+1}) = O(n_{\lambda+2})$ halo nodes). This representation generalizes the representation for the case $\lambda = 3$ given in [3], and we think that the algorithm and the data structure for the incremental maintenance of the 4-classes of G given in [9] can be adjusted for this generalization. However, the total skeleton model that we suggest in what follows provides, using Corollary 3.7, a *straightforward* reduction to the case $\lambda = 3$.

Now, we build a 2-skeleton $(\mathcal{H}^2, \varphi^2)$ of size $O(n_{\lambda+2}) = O(n)$ for G, with $\lambda(\mathcal{H}^2) = 2$. Let us use the version of $\mathcal{H}$ with weights 2 for tree-edges and 1 for cycle-edges, as in Section 3.2. Recall that with these weights, for λ even, $(\mathcal{H}, \varphi)$ is a 1-skeleton for G, with $\lambda(\mathcal{H}) = 2$. Observe that, in the sense of Section 3.1, any 2-skeleton $(\mathcal{H}^2_S, \varphi^2_S)$ is a local model at the node $\varphi(S)$ of $\mathcal{T}^\lambda$, of size linear in the number of the $(\lambda+2)$-classes in S and halo nodes of $\hat{G}_S$; moreover, any cycle $\mathcal{L}$ of $\mathcal{H}$ is a local model at the corresponding node $\mathcal{N}_\mathcal{L}$ of $\mathcal{T}^\lambda$. By Theorem 3.2 and Lemma 3.3, implanting all such local models into $\mathcal{T}^\lambda$ results in a model $(\bar{\mathcal{H}}^2, \varphi^2)$, of size linear in $n_{\lambda+2}$, that represents the λ- and local $(\lambda+1)$-cuts of G by its 2- and 3-cuts, respectively. We now adjust it, using Reduction(i), to represent also the global $(\lambda+1)$-cuts (see Fig. 16).

Let us paint red each cycle $\mathcal{L}$ of $\mathcal{H}$ implanted into $\mathcal{T}^\lambda$ (in what follows we identify it in $\mathcal{H}$ and in $\bar{\mathcal{H}}^2$). Let $\varepsilon = (\mathcal{N}_\mathcal{L}, \mathcal{N}_S)$ be a structural edge of $\mathcal{T}^\lambda$ such that a cycle $\mathcal{L}$ is implanted instead of $\mathcal{N}_\mathcal{L}$ and the local model $\mathcal{H}^2_S$ is implanted instead of $\mathcal{N}_S$ into $\mathcal{T}^\lambda$ (notice that in $\bar{\mathcal{H}}^2$ both $\mathcal{N}_\mathcal{L}, \mathcal{N}_S$ are empty nodes) and let $(\bar{\mathcal{N}}_S, b_\mathcal{L})$ be the "remainder" of ε in $\bar{\mathcal{H}}^2$, where $\bar{\mathcal{N}}_S \in \mathcal{L}, b_\mathcal{L} \in \mathcal{H}^2_S$ (see Fig. 16(a)). In the case $b_\mathcal{L}$ is distinguished and $\{\delta'(b_\mathcal{L}), \delta''(b_\mathcal{L})\}$ coincides with the partition of $\delta(b_\mathcal{L})$ induced by $\mathcal{L}$, we perform the following operation, in order to give an explicit representation to the bunch corresponding to $b_\mathcal{L}$ and $\mathcal{L}$. W.l.o.g., let us assume that $\delta'(b_\mathcal{L}) = E_{\varepsilon'}$ and $\delta''(b_\mathcal{L}) = E_{\varepsilon''}$, where $\varepsilon', \varepsilon''$ are the structural edges incident to $\bar{\mathcal{N}}_S$ in $\mathcal{L}$. We contract to a single new node the "reminder" $(\bar{\mathcal{N}}_S, b_\mathcal{L})$ in $\bar{\mathcal{H}}^2$ (according to the optional contraction described in Sect. 3.1) (see Fig. 16(b)). Then, we delete this node and identify each red edge incident to $b_\mathcal{L}$ in $\mathcal{H}^2_S$ with the one of $\varepsilon', \varepsilon''$ corresponding to it via the assumed coincidence (see Fig. 16(c)). We call this operation the **opening** of $\mathcal{H}^2_S$ into $\mathcal{L}$. (*Remark*: this operation can be executed in one phase as follows: if ε' corresponds to $(b_\mathcal{L}, \mathcal{N}')$ and ε'' to $(b_\mathcal{L}, \mathcal{N}'')$, then we delete $\bar{\mathcal{N}}_S$ and $b_\mathcal{L}$, together with $(\bar{\mathcal{N}}_S, b_\mathcal{L})$, $(b_\mathcal{L}, \mathcal{N}')$, and $(b_\mathcal{L}, \mathcal{N}'')$, and identify the "free" ends of ε' and ε'' with $\mathcal{N}'$ and $\mathcal{N}''$, respectively.) Observe that nonempty nodes of $\bar{\mathcal{H}}^2_S$ are not affected by such openings, i.e., the mapping remains the same. The size of the model does not increase.

The following situation may be considered as a degenerate case of the above construction. Let us consider a structural edge (b_1, b_2) of $\bar{\mathcal{H}}^2$ such that b_1, b_2 are halo nodes of local models $\mathcal{H}^2_{S_1}$ and $\mathcal{H}^2_{S_2}$, respectively (in this case, the edge (b_1, b_2) is the reminder of the edge $(\varphi(S_1), \varphi(S_2))$ of $\mathcal{T}^\lambda$). Observe that then the cuts defined by b_i in $\mathcal{H}^2_{S_i}$, $i = 1, 2$, model the same cut C of G as the structural edge (b_1, b_2), and thus $\delta(b_i)$ in $\mathcal{H}^2_{S_i}$, $i = 1, 2$, coincide, since each of them coincides with C. Assume that both b_1, b_2 are distinguished halo nodes, and that the corresponding balanced partitions of $\delta(b_1)$ and $\delta(b_2)$ coincide (see Fig. 16(d)). In this case, let us consider the edge (b_1, b_2) of the weight 2 as a cycle $\{\varepsilon', \varepsilon''\}$ of the length two in $\bar{\mathcal{H}}^2$ with the sets $E_{\varepsilon'}$ and $E_{\varepsilon''}$ equal to the parts of the balanced partition mentioned above. We define the opening of $\mathcal{H}^2_{S_1}$ and $\mathcal{H}^2_{S_2}$, one into the other, as the result of the two openings of $\mathcal{H}^2_{S_1}$ and $\mathcal{H}^2_{S_2}$ into that degenerate cycle (see Fig. 16(e),(f)) (in fact, to execute that opening, it is sufficient to remove b_1, b_2, and (b_1, b_2) and identify the corresponding pairs of red edges of $\mathcal{H}^2_{S_1}$ and $\mathcal{H}^2_{S_2}$).

We denote by $(\mathcal{H}^2, \varphi^2, \mathcal{F}^2)$, where $\mathcal{F}^2$ is the family of all 2- and 3-cuts of $\mathcal{H}^2$, the model obtained by executing all openings in $\bar{\mathcal{H}}^2$. Now, the following statement proves Theorem 4.1.

Lemma 4.16 *The model $(\mathcal{H}^2, \varphi^2)$ is a 2-skeleton for G, with connectivity 2 and of size $O(n_{\lambda+2})$.*

Proof: Let us consider an opening as described above. Observe that it does not produce bridges (hence, $\mathcal{H}^2$ is 2-connected) and retain all present cuts. Moreover, for any 3-cut cut C of $\mathcal{H}_S^2$ containing one of $(b_\mathcal{L}, \mathcal{N}')$ and $(b_\mathcal{L}, \mathcal{N}'')$, replacement of that edge by any edge ε in $\mathcal{L}$ results in a new 3-cut. For the modeled cuts, this corresponds to the replacement of the edge set $E_{\varepsilon'}$ or $E_{\varepsilon''}$ by E_ε, where $\varepsilon \in \mathcal{L}$. It is easy to see that those new cuts form exactly the bunch of type 1 generated by C and $\mathcal{L}$. Moreover, when C contains two such edges, the results of all pairs of such independent replacements form the bunch of type 2. Therefore, all the λ- and $(\lambda+1)$-cuts of G are φ^2-induced by the 2- and 3-cuts of $\mathcal{H}^2$.

Let us show that openings produce new 2- and 3-cuts only of the kind described above. By induction, let us assume that after some sequence of openings all the new 2- and 3-cuts are as required (for the initial graph this is trivial), and one more opening is executed; we keep notations as in the definition of an opening. Clearly, if a cut C of the new graph does not cross the cut $\{(b_\mathcal{L}, \mathcal{N}'), (b_\mathcal{L}, \mathcal{N}'')\} = \{\varepsilon', \varepsilon''\}$, it is a cut of the current graph and is, by the induction assumption, as required. Otherwise, let us consider the two corner cuts of the square of C and $\{\varepsilon', \varepsilon''\}$ that contain the edge ε'; both of them are cuts of the current graph. Clearly, their total cardinality is $|C|+2$. If C is a 2-cut, the only possibility is that both the corner cuts are 2-cuts. This is a contradiction, since, by the induction assumption, there are no 2-cuts dividing $\mathcal{H}_S^2$. If C is a 3-cut, the only possibility is that the corner cut C_1 dividing $\mathcal{H}_S^2$ is a 3-cut, and the other one is a 2-cut, which, by the induction assumption, must be of the form $\{\varepsilon', \varepsilon\}$, with $\varepsilon \in \mathcal{L}$. Observe that if C_1 is a cut of $\bar{\mathcal{H}}_S^2$, then C is in the bunch generated by C_1 and $\mathcal{L}$. If C_1 is a cut obtained by a replacement, say, in a cut C_0, corresponding to the opening of $\bar{\mathcal{H}}_S^2$ into a cycle $\mathcal{L}_1$, then by the induction assumption, the only possibility is that C_0 consists of a blue edge, ε', and a red edge of $\bar{\mathcal{H}}_S^2$ corresponding to $\mathcal{L}_1$. Clearly, in this case, C is in the bunch of type 2 generated by C_0, $\mathcal{L}$, and $\mathcal{L}_1$. Observe that a cut obtained by two replacements contains only one edge in $\bar{\mathcal{H}}_S^2$, which is blue; hence, C_1 cannot be of this kind.

It was mentioned above that the size of $\mathcal{H}^2$ is not greater than that of $\bar{\mathcal{H}}^2$, i.e., $O(n_{\lambda+2})$. □

Observe that the representation by $\mathcal{H}^2$ is not bijective. However, it can be shown that each λ-cut of G is represented at most twice (as in the cactus tree model) and, using Lemma 4.7, that each $(\lambda+1)$-cut of G is modeled by at most four cuts of $\mathcal{H}^2$.

From the last Lemma, Corollary 3.7, and [9] we deduce:

Theorem 4.17 *For λ even, the $(\lambda+2)$-classes of G can be maintained under a sequence of u updates* Insert-Edge *and q queries* Same-$(\lambda+2)$-Class? *in $O(u+q+n\log n)$ total time.*

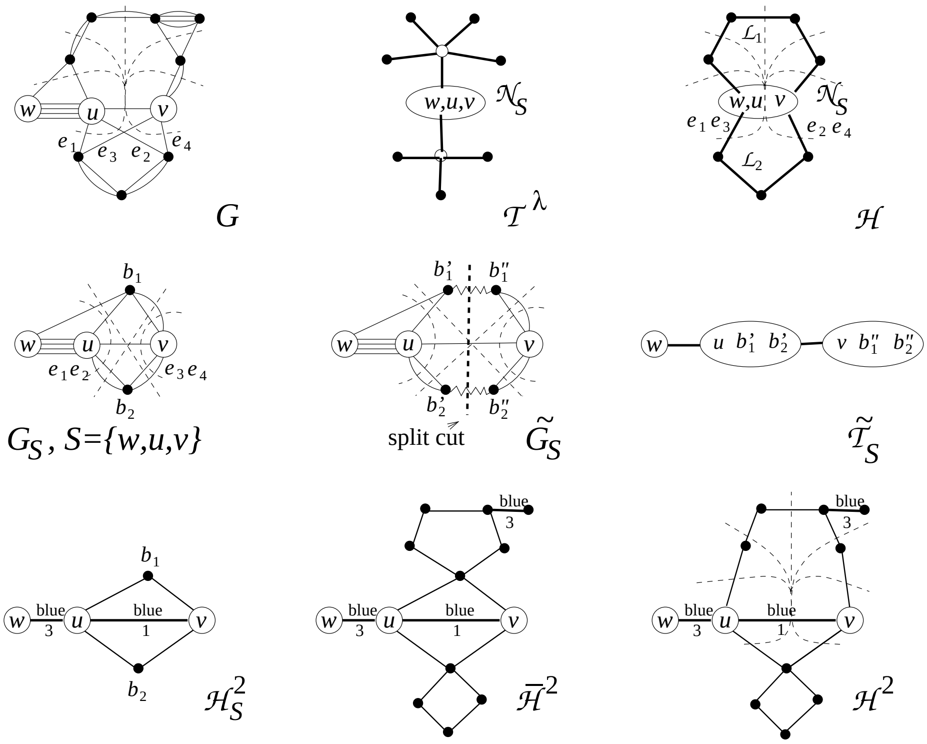

Figure 17: Example of the construction of the 2-level cactus tree model for λ even.

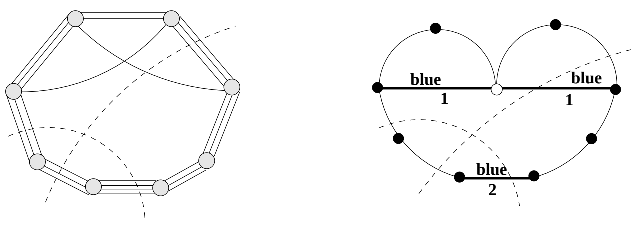

Figure 18: Example of the 2-level cactus tree model for a graph with $\lambda = 6$.

The worst case time for each query is $O(1)$. The initialization time is polynomial in n, and the space required is $O(n)$.

Relying on [17], the complexity of incremental maintaining can be reduced to $O((u + q + n)\alpha(u + q, n))$.

Notice that the above complexities of maintenance can be reduced substituting each instance of n by $n_{\lambda+2}$ 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 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 Concluding remarks

1. Observe that the properties mentioned in Theorem 4.1 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. It is likely that the 2-level cactus model can be a useful tool for handling edge-augmentation problems when the increase of the connectivity is 2. Another possible direction is to use our representation for effective maintenance of optimal augmentation sets of an incremental graph.

3. Let us call a cut of a weighted graph **subminimum** if its weight is the second minimum. It seems that the authors have obtained a generalization of 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 were combined; therefore, the new construction has the difficulties of both these cases simultaneously and 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.

References

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

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

[3] 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.

[4] 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).

[5] 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.

[6] 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 I: the odd case", a manuscript.

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

[8] 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).

[9] 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*).

[10] H. N. Gabow, "Applications of a poset representation to edge connectivity and graph rigidity", *Proc. 23rd Annual ACM Symp. on Theory of Computing*, 1991, 112–122.

[11] 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.

[12] 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.

[13] 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.

[14] 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, No. 2, 1986, 8–12).

[15] 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.

[16] H. Nagamochi, K. Nishimara and T. Ibaraki, "A tight upper bound on the number of small cuts in undirected graphs", Technical Report #94007, Kyoto Univ., Kyoto, Japan, 1994, 13p.

[17] 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.