An FPT Algorithm for Contraction to Cactus *

R. Krithika[1], Pranabendu Misra[2], and Prafullkumar Tale[1]

[1] The Institute of Mathematical Sciences, HBNI, Chennai, India.
{rkrithika|pptale}@imsc.res.in
[2] Department of Informatics, University of Bergen, Bergen, Norway.
Pranabendu.Misra@uib.no

Abstract. For a collection $\mathcal{F}$ of graphs, given a graph G and an integer k, the $\mathcal{F}$-Contraction problem asks whether we can contract k edges in G to obtain a graph in $\mathcal{F}$. $\mathcal{F}$-Contraction is well studied and known to be NP-complete for several classes $\mathcal{F}$. Heggerners et al. [Algorithmica (2014)] were the first to explicitly study contraction problems in the realm of parameterized complexity. They presented FPT algorithms for Tree-Contraction and Path-Contraction. In this paper, we study contraction to a class larger than trees, namely, *cactus graphs*. We present an FPT algorithm for Cactus-Contraction that runs in $c^k n^{\mathcal{O}(1)}$ time for some constant c.

1 Introduction

For a collection $\mathcal{F}$ of graphs, $\mathcal{F}$-Modification problem is to determine if an input graph G can be converted to some graph in $\mathcal{F}$ using at most k modifications. $\mathcal{F}$-Modification is an abstraction of practically well motivated problems like Vertex Cover, Feedback Vertex Set, Odd Cycle Transversal, Minimum Fill-In, to name a few. In recent times, there has been increasing interest in the study of Edge Contraction problems where the modification operation allowed is edge contraction. These problems generally turn out to be more difficult compared to their vertex/edge deletion/addition counterparts. For example, even determining whether a given graph G can be contracted to path of length four turns out to be NP-complete [4]. Formally, for a collection $\mathcal{F}$ of graphs, the $\mathcal{F}$-Contraction problem is to determine if an input graph G can be contracted to some graph in $\mathcal{F}$ using at most k edge contractions. For several choices of $\mathcal{F}$, early papers by Watanabe et al. and Asano and Hirata showed that $\mathcal{F}$-Edge Contraction is NP-complete even for several simple and well structured graph classes such as paths, stars, trees [2,3,18,19].

Graph contraction problems have received a lot of attention in parameterized complexity. It turns out that graph contraction problems are harder than their vertex/edge deletion/addition counterparts even in this setting. One of the intuitive reasons is that the classical *branching technique* does not work even for graph classes $\mathcal{F}$ that have a finite forbidden structure characterization. In case

* Due to space constraints, the proofs of results marked with $\star$ are moved to Appendix.

of vertex deletion or edge deletion/addition operations, to destroy a structure which forbids the input graph from being in $\mathcal{F}$, one needs to include at least one vertex (or edge) from that structure into the solution. This is not necessarily true in the case of contractions. Indeed, a forbidden structure may be destroyed by contracting edges which are not contained in the structure. Despite this inherent difficulty, there are several fixed-parameter tractability results known when the parameter is the solution size, i.e, the maximum number k of edges that can be contracted. To best of our knowledge, Hergerners et al. [12] were the first to explicitly study edge contraction problems in the realm of parameterized complexity. They presented a $4^k n^{\mathcal{O}(1)}$ algorithm for Tree Contraction and a $2^{k+o(k)} n^{\mathcal{O}(1)}$ algorithm for Path Contraction. When $\mathcal{F}$ is the set of graphs whose minimum degree is at least d, $\mathcal{F}$ is known to be FPT when parameterized by both k and d [10]. Golovach et al. proved that Planar Contraction is FPT [9]. Bipartite Contraction has been proved to be FPT by Heggernes et al. [13] and a faster algorithm was presented by Guillemot et al. [11]. Cai et al. [5] showed that Clique Contraction is FPT. On the negative side, it is known that $\mathcal{F}$-Contraction is W[2]-hard when $\mathcal{F}$ is either the family of $P_{\ell+1}$-free graphs or the family of C_ℓ-free cycles for some $\ell \geq 4$ [6,15]. Recently, Agarwal et al. [1] proved that Split Contraction is W[1]-hard.

In this paper, we present an algorithm for Cactus Contraction, adding it to the small list of graph classes for which FPT algorithms for contraction problems are known. A graph is called a *cactus* if every edge is a part of at most one simple cycle. Formally, the problem can be stated as follows.

Cactus Contraction	**Parameter:** k
Input: A graph G and an integer k	
Question: Does there exist $F \subseteq E(G)$ of size at most k whose contraction results in a cactus?	

It is easy to verify that the problem is in NP and its NP-completeness follows from [14]. As a cactus has treewidth at most 2, it follows that if a graph is k-contractible to a cactus, then its treewidth is at most $k+2$. Therefore, the problem is FPT by the celebrated result of Courcelle [7], as it is expressible in MSOL. However, this approach yields an impractical algorithm whose running time involves a large function of k. The main contribution of this work is a $c^k n^{\mathcal{O}(1)}$ algorithm for Cactus Contraction, where c is a fixed constant. Our algorithm builds upon ideas presented in [12], but requires a more involved structural analysis of the graph.

Outline of the algorithm: We can think of graph contraction problem as partition problem. The task is to find partition where each partition, called *witness set*, is connected and contracting all witness set to a point leads to desired graph. The idea is to color the graph with a small number of colors to "highlight" certain portions of the graph that contain the desired solution. This solution is then extracted via the structural properties of the graph. In first phase, we color $V(G)$ using three colors $\{1, 2, 3\}$ with the hope that all vertices of a *big witness set* (set with at least two vertices) receive the same color and that two distinct big witness sets with certain properties are "separated". We then identify some vertices that are not part of any big witness sets and recolor them using new colors 4 and 5. For instance, we identify certain induced paths that do not intersect with any minimal solution and are "adjacent" to only one big witness set (Lemma 3). The vertices of such paths are colored 4. After this we identify vertices that are not part of any big witness set and lie on a path between two big witness sets (Lemma 4) and color them using color 5. This completes the first phase. In the second phase, we extract the big witness sets from the components highlighted in the first phase. For this purpose, we define the notion of a *connected core* (Definition 4) which can be thought of as generalization of connected vertex cover. For every monochromatic component colored with $\{1, 2, 3\}$ by the first phase, we find connected vertex cover containing certain boundary vertices. The desired solution is the set of edges of a spanning forest of the corresponding connected cores.

The paper is organized as follows. In Section 2 we review some graph theoretic preliminaries. We present the properties of solution in Section 3 which are used in proving the correctness of algorithm. Following the approach of [12], we first give a randomized algorithm for the problem on 2-connected graphs, which is then used to give an algorithm in general graphs. Algorithm can be divided into two phases viz coloring phase (Section 4) and extracting a solution from colored graph (Section 5). Finally, in Section 6 we present overall algorithm and illustrate how this algorithm can be derandomized via (n, k)-universal sets. We remark that the main goal of this paper is to provide a $c^k n^{\mathcal{O}(1)}$ algorithm for Cactus Contraction, where c is a fixed constant. For the sake of simplicity, we have not attempted to optimize the running time.

2 Preliminaries

For graph theoretic terms and notation which are not explicitly defined here, we refer the reader to the book by Diestel [8]. An *undirected graph* is a pair consisting of a set V of vertices and a set E of edges where $E \subseteq V \times V$. An edge uv between vertices u and v is specified as an unordered pair of vertices. For a graph G, $V(G)$ and $E(G)$ denote the set of vertices and edges, respectively. Two vertices u, v are said to be *adjacent* if there is an edge uv in the graph. The *neighbourhood* of a vertex v, denoted by $N_G(v)$, is the set of vertices adjacent to v. The *degree* $d_G(v)$ of a vertex v is $|N_G(v)|$. The subscript in the notation for neighbourhood and degree are omitted if the graph under consideration is clear. For a set of edges F, $V(F)$ denotes the set of endpoints of edges in F. For a set $S \subseteq V(G)$, $G - S$ denotes the graph obtained by deleting S from G and $G[S]$ denotes the subgraph of G induced on the set S. For sets $X, Y \subseteq V(G)$, $E(X, Y)$ denotes the set of edges with one endpoint in X and other endpoint in Y. Similarly, $E(X)$ denotes the set of edges whose both endpoints are in X.

A *path* $P = (v_1, \ldots, v_l)$ is a sequence of distinct vertices in which there is an edge between any pair of consecutive vertices. The vertex set of P is the set $\{v_1, \ldots, v_l\}$ and is denoted by $V(P)$. The path P is called as a *cycle* if v_1 and

v_l are adjacent. An induced path (or cycle) is a path (or a cycle) in which no two non-consecutive vertices are adjacent. An induced path $P = (v_1, v_2, \dots, v_\ell)$ in G with $v_1 \neq v_\ell$ is called a simple path if $N_G(v_i) = \{v_{i-1}, v_{i+1}\}$ for each $2 \leq i \leq \ell - 1$. We define the neighborhood of such a path P as the set $N_G(P) = (N_G(v_1) \cup N_G(v_\ell)) \setminus V(P)$. We say that a set $X \subseteq V(G)$ is a simple path if there is an ordering of vertices in X that is a simple path. A graph is *connected* if there is a path between every pair of its vertices and it is *disconnected* otherwise. A set $S \subseteq V(G)$ is a *connected set of vertices* if $G[S]$ is connected. A *component* of a disconnected graph G is a maximal connected subgraph of G. A *cut-vertex* of a connected graph G is a vertex v such that $G - \{v\}$ is disconnected. A connected graph that has no cut-vertex is called *2-connected*. The operation of *subdividing* an edge uv results in the graph obtained by deleting uv and adding a new vertex w adjacent to both u and v. The operation of *short-circuiting* a degree two vertex v with neighbors u and w results in the graph obtained by deleting v and then adding the edge uw if it is not already present. A graph is called a *cactus* if every edge is a part of at most one cycle. Following properties of cactus are direct consequence of the definition.

Observation 1. *[14] The following statements hold for a cactus T.*

1. *The vertices of T can be properly colored using 3 colors.*
2. *Every vertex of degree at least 3 is a cut-vertex.*
3. *The graph obtained from T by subdividing any edge is a cactus.*
4. *The graph obtained from T by short-circuiting any degree 2 vertex is a cactus.*

The *contraction* operation of an edge $e = uv$ in G results in the deletion of u and v and the addition of a new vertex w adjacent to vertices that were adjacent to either u or v. The resulting graph is denoted by G/e. Formally, $V(G/e) = V(G) \cup \{w\} \setminus \{u, v\}$ and $E(G/e) = \{xy \mid x, y \in V(G) \setminus \{u, v\}, xy \in E(G)\} \cup \{wx \mid x \in N_G(u) \cup N_G(v)\}$. For a set of edges $F \subseteq E(G)$, G/F denotes the graph obtained from G by contracting the edges in F (in an arbitrary order). It is easy to see that G/F is oblivious to the contraction sequence.

A graph G is *contractible* to a graph T, if T can be obtained from G by a sequence of edge contractions. For graphs G and T with $V(T) = \{t_1, \dots, t_l\}$, G is said to have a *T-witness structure* $\mathcal{W}$ if $\mathcal{W}$ is a partition of $V(G)$ into l sets and there is a bijection $W : V(T) \mapsto \mathcal{W}$ such that the following properties hold.

- For each $t_i \in V(T)$, $G[W(t_i)]$ is connected.
- For a pair $t_i, t_j \in V(T)$, $t_i t_j \in E(T)$ if and only if there is an edge between a vertex in $W(t_i)$ and a vertex in $W(t_j)$ in G.

The sets $W(t_1), \dots, W(t_l)$ in $\mathcal{W}$ are called *witness sets* or *bags*. The bags $W(t)$ which contain a single vertex are called *small bags*, while the bags with more than one vertex are called *big bags*. For the sake of brevity, we omit curly brackets while denoting a singleton set. We associate a set $F \subseteq E(G)$ with a T-witness structure $\mathcal{W}$ of G, where F is the union of the set of edges of a spanning tree of the $G[W]$ for each $W \in \mathcal{W}$. Observe that $G/F = T$ and we say that G is $|F|$-contractible to T. Note that there is a unique T-witness structure of G

corresponding to a set F of edges. There are at most $|F|$ many big witness sets. Also the number of vertices which are contained in a big witness set is upper bounded by $|F| + 1$.

3 Key Properties of a Solution

In this section, we start with a simplifying assumption that let us concentrate on 2-connected graphs.

Proposition 1. *[14]. A graph is k-contractible to a cactus if and only if each of its 2-connected components is contractible to a cactus using at most k edges in total.*

Subsequently, we assume that the input graph G is 2-connected.

Observation 2 ($\star$)**.** *For a cactus T, let $\mathcal{W}$ be a T-witness structure of* 2-*connected graph G. If t is a cut-vertex in T, then $|W(t)| > 1$.*

Every big witness set need not be a cut vertex in T. We now define certain structures (or subgraphs) in T with respect to witness structure $\mathcal{W}$.

Definition 1 (Internal-Cactus). *The subgraph T_I of T obtained by removing any vertex which does not lie on a path between two distinct vertices in T corresponding to big bags is called as* internal-cactus *of T.*

We see that vertices of G which are not contained in witness sets corresponding to vertices in internal-cactus are easy to identify. For a given cactus T and its leaf t, if t does not corresponds to a big witness set then it can not be part of its internal cactus. We can say similar thing for cycles in T which have only one vertex which corresponds to one big witness set.

Definition 2 (Pendant Cycle). *A cycle in T is called as* pendant cycle *if there is exactly one vertex in cycle for which corresponds to a big witness set.*

If t is an unique vertex in cycle which corresponds to a big witness set, we say that pendant cycle is incident on t. To obtain an internal cactus, we need to delete all but one vertices in any pendant cycle. By Observation 2, every cut vertex in T is correspond to a big bag and hence it is a part of internal-cactus. In following observation, we bound the cardinality of neighborhood of such cut vertices in internal-cactus.

Observation 3 ($\star$)**.** *Let C_T be the set of cut-vertices in T. The number of neighbors of C_T in internal-cactus is at most $4|C_T|$. In other words, the number of vertices in $N_T(C_T)$ that are neither leaves nor part of a pendant cycle is at most $4|C_T|$.*

We end this section with following lemma which resolves a special instance of Cactus Contraction in polynomial time.

Lemma 1 ($\star$)**.** *If G is a 2-connected graph such that $V(G)$ can be partitioned into two simple paths P and Q in G, then we can solve the instance (G, k) of* Cactus Contraction *in polynomial time.*

4 Phase 1: The Coloring Phase

In coloring phase, we start with assigning uniformly at random one of colors $\{1, 2, 3\}$ to vertices of input graph G. Once we have obtained this coloring, we identify certain vertices of G which are contained in small witness sets. We re-color them using new colors $\{4, 5\}$ and move on to Phase 2 of algorithm to extract a solution from components of G which are colored $1, 2$ or 3.

We need notion of *compatible coloring* to argue the correctness of this coloring step. Let cactus T can be obtained from graph G by contracting edges F in G. Also, let $\mathcal{W}$ be T-witness structure of graph G. We determine whether a given coloring is compatible or not with respect to this witness structure. Informally speaking, for each big bag, a compatible coloring colors every vertex in this big bag with same color. It separates two big witness sets which shares an edge among them. If two big witness sets are connected by a path in G than the coloring gives different color to end points to this path.

Definition 3 (Compatible Coloring). *We say ϕ is* compatible *with $\mathcal{W}$ if the following three conditions are satisfied.*

- *For all $W(t) \in \mathcal{W}$, $W(t)$ is monochromatic.*
- *For all $t_x, t_y \in V(T)$ such that $|W(t_x)|, |W(t_y)| > 1$ and there is an edge in T between t_x and t_y, we have $\phi(W(t_x)) \neq \phi(W(t_y))$.*
- *For all $t_x, t_y \in V(T)$, such that $|W(t_x)|, |W(t_y)| > 1$ and there exists a simple path $P = (t_x, t_1, t_2, \ldots, t_q, t_y)$ in T such that $|W(t_i)| = 1$ for all $1 \leq i \leq q$, we have $\phi(W(t_x)) \neq \phi(W(t_1))$ and $\phi(W(t_y)) \neq \phi(W(t_q))$.*

We say that ϕ is compatible with solution F if ϕ is compatible with the witness structure $\mathcal{W}$ associated with F. We later argue that if (G, k) is an YES instance of Cactus Contraction than any random 3-coloring is compatible coloring with respect to an optimum solution with high probability. For this section, we assume that we are given a 3-coloring ϕ of G which is compatible with some optimum solution. Notice that we are not given the optimum solution. It is possible that same coloring can be compatible with different optimum solutions. In this section we prune coloring components and re-color them in order to move closer to obtain one of the optimum solution.

4.1 Properties of Coloring

We derive some structural properties of ϕ in G and use those properties to compute a solution. A set $X \subseteq V(G)$ is called a *colored component of* ϕ, if X is a maximal connected set of vertices that have the same color in ϕ. Let $\mathcal{X}$ be the set of all components of ϕ. Since $\mathcal{X}$ is a T-compatible partition and contracting an edge in a cactus graph results in another cactus graph, $\mathcal{X}$ is the witness structure of some cactus. For every color component X in $\mathcal{X}$, either all vertices of X are in small bags in $\mathcal{W}$ or X contains exactly one big witness set in $W(t) \in \mathcal{W}$ and the remaining vertices $X \setminus W(t)$ are in small bags. Given a coloring ϕ, we are only interested in finding an optimum solution which is compatible with this coloring. Hence, for any two components X, Y of ϕ, no edge uv in $E(X, Y)$ is in optimum solution.

We start with simple case when a connected component X in $\mathcal{X}$ is simple path in G. Lemma 2 states that X is either one big witness set or all vertices in X are singleton sets. The proof of the lemma is based on the observation that if two adjacent bags have only one edge crossing them then this edge is not incident vertex which has degree two.

Lemma 2 **($\star$).** *If colored component X in $\mathcal{X}$ is a simple path in G then either all vertices of X are in small bags or X is a big witness set in $\mathcal{W}$.*

4.2 Identifying Vertices in Pendant Cycles and Leaves

We now specify the criteria to identify vertices in G that are contained in pendant cycles in T or are leaves in T. We can not identify all such vertices in this phase.

Re-coloring I: For any colored component X in $\mathcal{X}$, if $G-X$ contains a vertex or a simple path as its connected component then recolor vertices in that connected component with color 4.

The re-coloring signifies that these vertices are part of pendant cycles or they are leaves in T. Notice that since v is not included in X and it is adjacent with X, initially vertex v had different color than X. We can say similar things for end points of path P. We argue that if vertices and simple paths in G are adjacent to only one colored components then they are either part of pendent cycles or leaves in T.

Lemma 3 **($\star$).** *For a colored component X in $\mathcal{X}$, let P be a connected component of $G - X$. If P is a simple path in G whose neighborhood is contained in X then P is either a part of a pendant cycle or it is a leaf in T.*

Notice that above Lemma also holds when P contains only one vertex. For a colored component X in $\mathcal{X}$, suppose there is an isolated vertex v which is connected component of $G - X$. Since ϕ is compatible with optimum solution, all big witness sets are monochromatic. This implies v can not be part of any big witness set and remains as singleton witness set. As it can have path to at most one big witness set, it is either part of some pendant cycle in T or it is a leaf.

4.3 Identifying Vertices in Simple Paths

We now identify vertices in G that correspond to paths in T that are between two big witness sets. Recall that in simple path no internal vertex is adjacent to any vertex outside this path. A simple path is *maximal* if it is not contained in any other simple path. In other words, in *maximal simple path* every internal vertex has degree exactly two and end points have degree strictly greater than two. We color vertices which are in maximal simple path which has neighbors in two different colored component.

Re-coloring II: For any two colored component Y, Z in $\mathcal{X}$, if $G - (Y \cup Z)$ contains a vertex or a maximal simple path as its connected component then recolor vertices in that connected component with color 5.

The re-coloring signifies that these vertices are part of simple paths in internal cactus of T. As in case of Re-coloring-I, a vertex and end points of paths have different color than either Y or Z. We prove the correctness of this coloring in following Lemma. We state this lemma when P is maximal simple path but it holds for a vertex.

Lemma 4 **($\star$).** *For two colored components Y, Z in $\mathcal{X}$, let P be a connected component of $G - (Y \cup Z)$. If P is a maximal simple path in G then no optimum solution contains a solution edge incident on vertices in P. Furthermore, both Y and Z contain big witness sets.*

4.4 Properties of Recoloring

By definition of compatible coloring, every colored component contains at most one big witness set. Before re-coloring, any colored component may or may not contain big witness set. In Lemma 5, we argue that after re-coloring, all colored components colored with $\{1, 2, 3\}$ must contains a big witness set. We can think of Lemma 5 as *(partial) completeness* part for Lemma 3 and 4. In other words, in Lemma 3 (in Lemma 4) we argue that vertices in G which satisfy some criteria are contained in witness sets which are part of pendent cycles or are leaves (in simple paths) of cactus T. In Lemma 5, we claim that *all* vertices in colored component which do not contain a big witness set and are part of pendent cycles or are leaves (simple paths) of cactus T satisfies the premise of Lemma 3 (Lemma 4).

Lemma 5 **($\star$).** *If a colored component X in $\mathcal{X}$ is monochromatic with color from $\{1, 2, 3\}$ after exhaustive application of two re-coloring rules then X contains a big witness set.*

5 Phase 2: Identifying Big Witness Sets

At the start of Phase 2, we have identified colored component which must contains big witness set. For a colored component X in $\mathcal{X}$, let $W(t)$ is a big witness set contained in X. Our objective in this section is to find subset X' of X which is *at least as good as* $W(t)$. Informally speaking, this means we can replace edges in spanning tree of $G[W(t)]$ by edges in spanning tree of $G[X']$ in any optimum solution F and we get another optimum solution F'. We examine what properties $W(t)$ has in graph $G[X]$. In fact, we consider a superset $\hat{X}$ of X and examine the properties of $W(t)$ with respect to graph $G[\hat{X}]$.

Let $\hat{X}$ be the superset of X which contains vertices in the connected components of $G - X$ that are either isolated vertices or a simple path in G whose neighborhood is contained in X. We now define the notion of connected core.

Definition 4 (Core). *A* core *of a graph G is a set $Z \subseteq V(G)$ such that every connected component of $G - Z$ is either an isolated vertex or a simple path whose neighborhood is contained in Z. If a core Z is a connected set in G, then we call it a* connected core *of G.*

Notice that any superset of a connected-core which induces a connected subgraph is also a connected core. We postpone discussion on how to find a connected core of given graph which contains specified vertex set and is of minimum size to Subsection 5.1. We claim that $W(t)$ is a connected core of graph $G[\hat{X}]$.

Lemma 6 ($\star$). *For a colored component X in $\mathcal{X}$, if $W(t)$ is the big witness set contained in X then $W(t)$ is a connected core of $G[\hat{X}]$.*

We point out that it is possible that there exists a proper superset of $W(t)$ which is a connected core of $G[\hat{X}]$. In other words, every vertex in $W(t)$ has at least one of the two responsibility: it is a part of connected core of $G[X]$ or it is in $W(t)$ because of external constraints. We introduce Marking Scheme 1 to mark vertices which are in $W(t)$ because of external constraints. Once we mark vertices which are present in big witness set because of external constraints, we can find *any* connected core of minimum cardinality which contains these vertices and this connected core is *as good as* $W(t)$ for our purposes. Marking scheme is as follows.

Marking Scheme 1. *For a colored component X in $\mathcal{X}$,*

1. *If there exists y in $N(X)$ such that $\phi(y) = 5$ then mark all the vertices in $N(y) \cap X$.*
2. *For a colored component X in $\mathcal{X}$ which contains a big witness set, mark all vertices in $N(X') \cap X$*

We now prove the soundness of this marking scheme. Lemma 7 and 8 argue that if X contains a big witness set $W(t)$ then all the vertices marked by marking scheme are contained in $W(t)$.

Lemma 7 ($\star$). *If there exists v in $N_G(X)$ such that v is colored 5 then $N_G(v) \cap X$ is contained in a big witness set of X.*

Lemma 8 ($\star$). *Let X, Y be two colored component in $\mathcal{X}$ which contain big witness sets, say, W_X and W_Y, respectively. Then, $N(X) \cap Y \subseteq W_Y$ and $N(Y) \cap X \subseteq W_X$.*

In the following Lemma we prove *completeness* of the marking scheme. We argue that all vertices which are present in big witness set because of external constraints has been marked by Marking Scheme 1. It is sufficient to argue that if t_1 is neighbor of t in internal-cactus of T_I then all vertices in $N_G(W(t_1)) \cap X$ has been marked. Completeness of Marking Scheme 1.1 and 1.2 follows when $|W(t_1)|$ is one and strictly greater than one, respectively.

Lemma 9 ($\star$). *For a colored component X in $\mathcal{X}$ let $W(t)$ be the big witness set contained in X. If t_1 is a neighbor of t in the internal cactus T_I of T then all the vertices in $N_G(W(t_1)) \cap X$ has been marked by Marking Scheme 1.*

We now prove how this marking scheme and connected core help us to identify a set in X which is as good as $W(t)$.

Pruning Operation: For a given collection of colored component X, consider another set $\mathcal{X}'$ obtained by performing following operations. For every colored-component Y in $\mathcal{X}$ of cardinality at least 2, if a vertex in v got recolored to 4 or 5, remove Y from $\mathcal{X}$ and add $Y \setminus \{v\}$ and $\{v\}$ to $\mathcal{X}$. For a colored component X in $\mathcal{X}$ which contains a big witness set, let M_X be set of marked vertices in X by Marking Scheme 1. Let Z_X be a connected core of $G[\hat{X}]$ of minimum cardinality which contains set M_X. For every colored component X in $\mathcal{X}$, if Z_X is proper subset of X then remove X and add Z_X to $\mathcal{X}$. For every vertex v in $\hat{X} \setminus Z_X$, add a singleton set $\{v\}$ to $\mathcal{X}$ (see Figure 1 in Appendix).

We stop the pruning operation when no colored component is replaced in $\mathcal{X}$. Notice that this pruning operation stops in polynomial time with respect to number of vertices in graph. As final lemma in this section, we argue that if we start applying pruning operation on set of colored classes obtained from compatible coloring ϕ, we end up with a witness structure corresponding with an optimum solution. Recall that F is a minimum set of edges such that G/F is a cactus and $\mathcal{W}$ is the G/F witness structure of G. Also, ϕ is coloring of $V(G)$ which is compatible coloring with respect to $\mathcal{W}$. Set $\mathcal{X}$ is collection of colored components of ϕ.

Lemma 10 ($\star$). *Let set $\mathcal{X}'$ be obtained from $\mathcal{X}$ by exhaustive application of Pruning Operations. If F^* be a union of spanning trees of graph induced on colored components in $\mathcal{X}^*$ then G/F^* is a cactus and $|F'| = |F|$.*

5.1 Finding Connected Cores

Recall that a connected-core of a graph G is subset Z of vertices such that, $G[Z]$ is connected and each connected component of $G - Z$ is either an isolated vertex or a simple path whose both end points have neighbors in Z. Here, we present a simple branching algorithm that determines if G has a connected core of size at most k or not. We use algorithm for Steiner Tree problem as subroutine. In Steiner Tree problem, we are given a graph G and set of vertices, called *terminals*, and a positive integer ℓ. The goal is to determine whether these is a tree with at most ℓ edges that connects all the terminals.

Lemma 11 ($\star$). *There is an algorithm that given a connected graph G and a subset X of its vertices, computes a minimum connected core of G which has at most k vertices and contains X in $\mathcal{O}^*(6^k)$ time if one such exists in the graph.*

6 Putting it all Together: The Overall Algorithm

The pseudo-code of the algorithm is presented as Algorithm 6.1 and Theorem 1 formally states our result.

Algorithm 6.1: Randomized Algorithm for Cactus Contraction

Input: A 2-connected graph G and an integer k
Output: A set F of k edges in G such that G/F is a cactus

Generate random coloring $\phi : V(G) \rightarrow \{1, 2, 3\}$ and construct $\mathcal{X}$.
for *each* $X \in \mathcal{X}$ **do**
    if *P is a simple path or a isolated vertex in* $G - X$ **then**
        for all $u \in P$: set color of u to 4
for *each pair* $X_1, X_2 \in \mathcal{X}$ **do**
    if *P is a simple path or a isolated vertex in* $G - (X_1 \cup X_2)$ **then**
        for all $u \in P$: set color of u to 5
for *each* $X \in \mathcal{X}$ **do**
    Apply Marking Scheme to obtain the set of marked vertices $Y_X \subseteq X$
    $Z_X \leftarrow$ minimum connected core of $(G[\hat{X}], Y_X)$
Construct $\mathcal{X}^*$ from $\mathcal{X}$ and $\{Z_X | \, X \in \mathcal{X}\}$.
if *a spanning forest F^* of $\mathcal{X}^*$ has $\leq k$ edges* **then**
    return F^*
else
    return NO

Theorem 1 **($\star$).** *There is an one-sided error Monte Carlo algorithm with false negatives which solves* Cactus Contraction *in time $c^k n^{\mathcal{O}(1)}$ on 2-connected graphs. It returns correct answer with constant probability.*

We apply the arguments presented in [12] to extend above theorem to solve Cactus Contraction on general graphs.

Theorem 2 **($\star$).** *There is an one-sided error Monte Carlo algorithm with false negatives which solves* Cactus Contraction *in time $c^k n^{\mathcal{O}(1)}$. It returns correct answer with constant probability.*

We can derandomize our algorithms by constructing a family of coloring function, that is derived from a perfect hash family. The details of the same are deferred to the appendix. This leads to the following result.

Theorem 3. Cactus Contraction *can be solved in $c^k n^{\mathcal{O}(1)}$ time.*

Acknowledgements: We would like to thank Prof. Saket Saurabh for invaluable advice and several helpful suggestions.

References

1. Agrawal, A., Lokshtanov, D., Saurabh, S., Zehavi, M.: Split contraction: The untold story. In: Symposium on Theoretical Aspects of Computer Science, STACS (2017)
2. Asano, T., Hirata, T.: Edge-contraction problems. Journal of Computer and System Sciences 26(2), 197 – 208 (1983)
3. Brouwer, A., Veldman, H.: Contractibility and NP-completeness. Journal of Graph Theory 11(1), 71–79 (1987)
4. Brouwer, A.E., Veldman, H.J.: Contractibility and np-completeness. Journal of Graph Theory 11(1), 71–79 (1987)
5. Cai, L., Guo, C.: Contracting Few Edges to Remove Forbidden Induced Subgraphs, pp. 97–109. Springer International Publishing (2013)
6. Cai, L., Guo, C.: Contracting few edges to remove forbidden induced subgraphs. In: IPEC. pp. 97–109. Springer (2013)
7. Courcelle, B.: The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. Information and Computation 85(1), 12 – 75 (1990)
8. Diestel, R.: Graph Theory. Springer-Verlag Berlin Heidelberg (2006)
9. Golovach, P.A., van t Hof, P., Paulusma, D.: Obtaining planarity by contracting few edges. Theoretical Computer Science 476, 38 – 46 (2013)
10. Golovach, P.A., Kamiński, M., Paulusma, D., Thilikos, D.M.: Increasing the minimum degree of a graph by contractions. Theoretical computer science 481, 74–84 (2013)
11. Guillemot, S., Marx, D.: A faster FPT algorithm for Bipartite Contraction. Information Processing Letters 113(2224), 906 – 912 (2013)
12. Heggernes, P., van't Hof, P., Lévêque, B., Lokshtanov, D., Paul, C.: Contracting Graphs to Paths and Trees. Algorithmica 68(1), 109–132 (2014)
13. Heggernes, P., van 't Hof, P., Lokshtanov, D., Paul, C.: Obtaining a bipartite graph by contracting few edges. SIAM Journal on Discrete Mathematics 27(4), 2143–2156 (2013)
14. Krithika, R., Misra, P., Rai, A., Tale, P.: Lossy kernels for graph contraction problems. In: 36th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2016, December 13-15, 2016, Chennai, India. pp. 23:1–23:14 (2016)
15. Lokshtanov, D., Misra, N., Saurabh, S.: On the hardness of eliminating small induced subgraphs by contracting edges. In: International Symposium on Parameterized and Exact Computation. pp. 243–254. Springer (2013)
16. Naor, M., Schulman, L.J., Srinivasan, A.: Splitters and near-optimal derandomization. In: Foundations of Computer Science, 1995. Proceedings., 36th Annual Symposium on. pp. 182–191. IEEE (1995)
17. Nederlof, J.: Fast polynomial-space algorithms using inclusion-exclusion. Algorithmica 65(4), 868–884 (2013)
18. Watanabe, T., Ae, T., Nakamura, A.: On the removal of forbidden graphs by edge-deletion or by edge-contraction. Discrete Applied Mathematics 3(2) (1981)
19. Watanabe, T., Ae, T., Nakamura, A.: On the NP-hardness of edge-deletion and -contraction problems. Discrete Applied Mathematics 6(1), 63 – 78 (1983)

Appendix

6.1 Missing Figure from Section 5

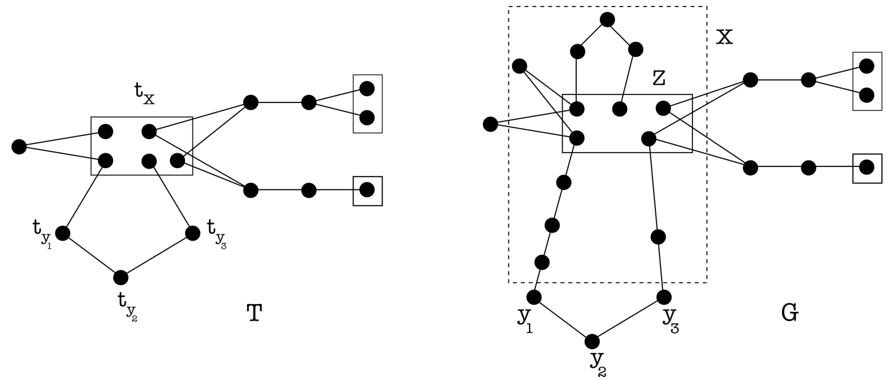

Fig. 1. Square represents a connected component in graph. Consider a colored component X in graph G on right hand side. Instead of contracting all of X to a vertex t_X (left side graph), we contract connected core Z of $G[\hat{H}]$ to a single vertex which require smaller edges to be contracted. We replace X by Z and singleton set for every vertex in $\hat{X} \setminus Z$ in $\mathcal{X}$.

6.2 Missing Proofs from Section 3

Observation. *For a cactus T, let $\mathcal{W}$ be a T-witness structure of 2-connected graph G. If t is a cut-vertex in T, then $|W(t)| > 1$.*

Proof. If t is a cut-vertex in T and $W(t) = \{u\}$ then we argue that u is a cut-vertex in G. Let T_1 and T_2 be any two connected components obtained by removing t from cactus T. Since t is a cut vertex there exits at least two such components which are not empty. Consider sets of vertices V_1, V_2 which are contained in witness sets corresponding to vertices in T_1, T_2 respectively. Formally, $V_1 = \{u | \ u \in W(t_1) \text{ and } t_1 \in W_1\}$ and V_2 defined in similar way. Since T_1, T_2 are non-empty, so are V_1, V_2. There is no edge between T_1, T_2 in T and since T is obtained from graph G by contracting edges, there is no edge between V_1, V_2 in G. This implies that $G - v$ has at least two connected component viz V_1, V_2. This contradicts the fact that G is a 2-connected graph. Hence for every cut vertex t in T, $|W(t)| > 1$. □

Observation. *Let C_T be the set of cut-vertices in T. The number of neighbors of C_T in internal-cactus is at most $4|C_T|$. In other words, the number of vertices in $N_T(C_T)$ that are neither leaves nor part of a pendant cycle is at most $4|C_T|$.*

Proof. We prove later part of the observation using induction hypothesis on number of cut vertices. If C_T is an empty set then the observation is vacuously true. By induction hypothesis, assume this observation is true when the number of cut vertices in graph are strictly smaller than $|C_T|$. Consider any cycle C of T which has at least two cut vertices. Let X_C be the set of all cut-vertices in $V(C)$. Consider cut vertices t, t_a, t_b in X_C such that t_a, t, t_b appear consecutively on cycle C. Vertices t_a, t_b need not be two different cut vertices. Since t is a cut vertex there exists a connected component T_1 of $T - t$ which contains vertices in t_a, t_b. Let T_2 be the graph obtained by adding a pendent vertex to the induced subgraph of T on vertices in $V(T) \setminus V(T_1)$. Let C_1, C_2 are cut vertices in graphs T_1, T_2 respectively. It is easy to see that $C_1 = C_T \cap V(T_1)$ and $C_2 = C_T \cap V(T_2)$. Any vertex in $N_{T_2}(C_2)$ that is neither leave nor part of a pendant cycle in graph T_2 is in $N_T(C_T)$ and it is not a leave nor part of pendant cycle in T. Similar statement is true for vertices in $N_{T_1}(C_1)$. The only vertices in $N_T(C_T)$ which are not leaves or part of pendent cycles and have not been counted in either $N_{T_1}(C_1)$ or $N_{T_2}(C_2)$ are at most four vertices between path t_a to t and t to t_b. Using Induction Hypothesis on graphs T_1, T_2, we get $|N_T(C_t)| \leq |N_{T_1}(C_1)| + |N_{T_2}(C_2)| + 4 \leq 4|C_1| + 4|C_2| + 4 \leq 4(|C \setminus V(T_2)| - 1) + 4(|C \cap V(T_2)| + 4 = 4|C_T|$. □

Lemma. *If G is a 2-connected graph such that $V(G)$ can be partitioned into two simple paths P and Q in G, then we can solve the instance (G, k) of* Cactus Contraction *in polynomial time.*

Proof. Let p_1, p_2 and q_1, q_2 be the endpoints of the simple paths P and Q, respectively. Observe that G has a hamiltonian cycle, as G is 2 connected and p_1, p_2, q_1, q_2 are the only vertices that may have degree greater than two. If G is an induced cycle, then the optimal solution is the empty set. Otherwise, G is a cycle with either one or two additional edges between p_1, p_2 and q_1, q_2. It follows that any optimal solution requires at most 3 edge contractions. □

6.3 Missing Proofs from Section 4

Lemma. *If colored component X in $\mathcal{X}$ is a simple path in G then either all vertices of X are in small bags or X is a big witness set in $\mathcal{W}$.*

Proof. Let X be the simple path $(v_1, v_2 \ldots, v_\ell)$. We consider a case when X contains a big witness set. We argue that if $W(t)$ is a big witness set contained in X then $X = W(t)$. For the sake of contradiction assume that there exists $v_i \in V(P) \setminus W(t)$. Notice that $W(t)$ is induces a connected subgraph and it is entirely contained in X. This implies that either v_1 or v_ℓ are not contained in $W(t)$. Without loss of generality, let v_1 is is not contained in $W(t)$. Let v_{i+1} be the least indexed vertex in $W(t)$. In other words, vertices $\{v_1, v_2, \ldots, v_i\}$ are not contained in $W(t)$ and remains as singleton witness set. Since $|W(t)| > 1$ we have $v_{i+1}, v_{i+2} \in W(t)$. Notice that v_{i+1} is a vertex in $W(t)$ such that $d_G(v_{i+1}) = 2$ and it has exactly one neighbor v_{i+2} in $W(t)$. The other neighbor of v_{i+1} is not in $W(t)$. We also note that there is no neighbor of v_i in $W(t)$ apart from v_{i+1}.

We now argue that such situation is not possible in a witness structure associated with a *minimal* solution. Let F be an optimum solution associated with witness structure $\mathcal{W}$. Since F contains a spanning tree of all big witness sets and there is unique edge $v_{i+1}v_{i+2}$ in $G[W(t)]$ which is incident on v_{i+1}, edge $v_{i+1}v_{i+2}$ is present in F. For a vertex u in $V(G)$, $W(t_u)$ denotes the witness set in $\mathcal{W}$ that contains u. Consider a graph T' obtained from G by contracting all edges in $F \setminus \{v_{i+1}v_{i+2}\}$ and let $\mathcal{W}'$ be T'-witness structure of G. Notice that we can obtain $\mathcal{W}'$ from $\mathcal{W}$ by first removing $W(t)$ and then adding two new sets $\{v_{i+1}\}$ and $W(t) \setminus \{v_{i+1}\}$. In graph T', vertex $t'_{v_{i+1}}$ is adjacent to exactly two vertices viz $t'_{v_i}, t'_{v_{i+2}}$ and $T'/t_{v_{i+1}}t_{v_{i+2}} = T$. Hence graph T' can be obtained from cactus T. By Observation 1 (3), T' is also a cactus. This contradicts to the fact that F is a minimal solution. Hence our assumption was wrong and this concludes the proof of lemma. □

Lemma. *For a colored component X in $\mathcal{X}$, let P be a connected component of $G - X$. If P is a simple path in G whose neighborhood is contained in X then P is either a part of a pendant cycle or it is a leaf in T.*

Proof. For the sake of contradiction assume the lemma is false. Recall that ϕ is a T-compatible coloring of G. Observe that P is a simple path in G such that, $N_G(P)$ is a subset of $X \in \mathcal{X}$. Therefore, for any $Y \in \mathcal{X}$ such that $Y \cap P \neq \emptyset$, we have $Y \subseteq P$. Hence by Lemma 2, it follows that no proper subset of edges in P is contained in the minimal solution F. Therefore all the edges of P are in F, which implies that $P \in \mathcal{X}$. Let t_p be the vertex corresponding to P in T, and observe that it is adjacent to $t \in T$ if and only if $W(t)$ contains a vertex from $N_G(P)$ (which is a subset of X). Now it is easy to see that, if $|E_G(P, X)| \leq 2$, then it is safe to un-contract all the edges of P, which contradicts the minimality of the solution F. Hence $|E_G(P, X)| \geq 3$. Before proceeding further, let us make a few observations about the graph. The set P is a big witness set in $\mathcal{W}$, and recall that either every vertex of X is a small bag in $\mathcal{W}$, or X contains exactly one big witness set in $\mathcal{W}$.

In the first case, every vertex of X defines a small bag, and X is a connected subset of T, and there are at least 3 edges in $E_G(P, X)$. If $N_G(P)$ corresponds to at least 3 vertices in T, then $T[X \cup t_p]$ contains two cycles with a common edge, i.e. T is not a cactus, which is a contradiction. Furthermore, $N_G(P)$ must contain at least two vertices, as P is a simple path in the 2-connected graph G and only the end-points of P have neighbors outside P. Hence, we may conclude that $N_G(P)$ contains exactly two vertices of X and $E_G(P, X)$ contains either 3 or 4 edges. Let v_1 and v_ℓ be the endpoints of P, and let x_1 and x_2 be the two vertices in $N_G(P)$. As G is 2 connected and every cut-vertex in T must correspond to a big witness set in $\mathcal{W}$, we conclude that no vertex of X is a cut-vertex in T or G. As X is a connected set, let Q be the path between x_1 and x_2 in G. As all the vertices of X lie in small bags, we have that Q is a path in T as well. Observe that $C = T[V(Q) \cup t_p]$ is a cycle in T with t_p being the only vertex corresponding to a big witness set in C. We claim that $T = C$. If not, then $V(T) \setminus (t_P \cup V(Q))$ is non-empty, and there is a vertex $y \in V(T) \setminus (V(P) \cup V(Q))$, such that there are two internally vertex disjoint paths between y and t_p in T. Indeed, we may start with a arbitrarily chosen y, and consider a minimum cut between y and t_p in T. If the minimum cut is a single vertex y', then observe that $y' \notin V(Q)$, as vertices of Q are not cut-vertices in T. We substitute y with y' and start over. Since the shortest path between y' and t_p in T is strictly shorter than the shortest path between y and t_p, we will obtain the vertex y in finitely many iterations, such that there are two internally vertex disjoint paths in T between y and t_P. Let R_1 and R_2 be those two paths. Observe that, we must have $x_1 \in R_1$ and $x_2 \in R_2$, and hence $R_1 \cup R_2$ contains a path between x_1 and x_2, say R in T. The path must be distinct from the path Q, because $Q_1 = V(Q) \cap V(R_1)$ and $Q_2 = V(Q) \cap V(R_2)$ are disjoint and therefore at least one edge of Q is absent from R. This implies that T contains three distinct paths between x_1 and x_2, namely $P_T = (x_1, t_p, x_2), Q$ and R. However, this contradicts the fact that T is a cactus. So, we conclude that $T = C$ and hence, $G = P \uplus Q$. Observe that, as both P and Q are simple paths in G, it is an instance that can be solved in polynomial time via Lemma 1. By our assumptions, this cannot be the case.

In the second case, let $Z \subseteq X$ be a big witness set in $\mathcal{W}$, and let t_Z be the vertex in T obtained by contracting Z. We claim that $N_G(P)$ is a subset of Z. Indeed, if this were not the case, let $v \in N_G(P) \setminus Z$, and let t_Z and t_P be the two cactus vertices corresponding to the big witness sets Z and P. Observe that v forms a small bag, and as X is a connected set, it lies on a path Q between t_Z and t_P in T, where the internal vertices of Q are associated with vertices of X. Observe that all internal vertices of Q are assigned the same color as the vertices of Z by ϕ. This contradicts the fact that ϕ is a compatible coloring. Hence $N_G(P) \subseteq Z$ and so we conclude that $N_T(t_p) = t_Z$ in T. Hence $G/(F-P)$ is also a cactus. This contradicts the minimality of the solution F. This concludes the proof of this lemma. □

Lemma. *For two colored components Y, Z in $\mathcal{X}$, let P be a connected component of $G - (Y \cup Z)$. If P is a maximal simple path in G then no optimum solution contains a solution edge incident on vertices in P. Furthermore, both Y and Z contain big witness sets.*

Proof. Let $P = (v_1, v_2, \dots, v_\ell)$ be a maximal simple path in G such that there exists $Y, Z \in \mathcal{X}$ with $N_G(v_1) \subseteq Y \cup v_2$, $N_G(v_\ell) \subseteq Z \cup v_{\ell-1}$. Let F be any optimum solution which is compatible with coloring. We argue that $E(P) \cap F = \emptyset$. Furthermore, both Y and Z contain big witness sets in $\mathcal{W}$.

Consider the first part of lemma and suppose that it is false, i.e $E(P) \cap F \neq \emptyset$. Now observe that, as ϕ is a compatible coloring, for any $A \in \mathcal{X}$, if $A \cap P \neq \emptyset$ then we have $A \subseteq P$. So by Lemma 2, P must be a witness set in $\mathcal{W}$. Let t_P be the vertex in T obtained from P. First, we claim that the neighborhood of P in Y, i.e. $Y_P = Y \cap N(P) = Y \cap N(v_1)$, must be in the same witness set of $\mathcal{W}$. Let $Z_P = Z \cap N(P) = Z \cap N(v_\ell)$. Consider Y_P and suppose that $y_1, y_2 \in Y_P$ are in two different witness sets in $\mathcal{W}$ corresponding to vertices $t(y_1), t(y_2)$ in T. As G is a 2-connected graph and P is a simple path in G, $G - V(P)$ is connected and so is $T - t_P$. Therefore the vertices $t(y_1), t(y_2)$ and t_z are connected in $T - t_P$

where t_z is a vertex in T such that $W(t_z) \cap Z_P \neq \emptyset$. Now observe that t_P is adjacent to $t(y_1), t(y_2)$ and t_z in T, which implies that T contains two cycles that have a common edge. However, this is a contradiction to the fact that T is a cactus. Hence, all of Y_P lies in the same witness set in $\mathcal{W}$. We can similarly show that Z_P is contained in the a witness set of $\mathcal{W}$.

Now, t_P, t_Y, t_Z be vertices in T where and $t_Y, t_Z \in T$ are formed by contracting the witness sets in $\mathcal{W}$ that contain Y_P and Z_P respectively. Clearly $(t_Y t_P)$ and $(t_P t_Z)$ are edges in T, and t_P is a vertex of degree 2 in T. Now observe that v_1 is not a cut-vertex in $G[P]$ as there is exactly one edge incident on it in this graph. Therefore, $G[V(P) \setminus \{v_1\}]$ is connected. Let $F' = F \setminus \{v_1 v_2\}$ and $P' = P - v_1$. Consider the graph $T' = G/F'$, and let t'_P and t_1 be the vertices in T' corresponding to P' and v_1. If T' is not a cactus, then there is a pair of cycles that have a edge in common. In particular, there is a pair of cycles that have the edges $(t_Y, t_1), (t_1, t'_P)$ and t'_P, t_Z, as t_1 and t'_P have degree two. But then $T = T/(t_1, t'_P)$ cannot be a cactus, which is a contradiction. On the other hand, if T' is a cactus it the minimality of the solution F. Hence, it is indeed true that $E(P) \cap F = \emptyset$.

Next, we argue that Y and Z must contain a big witness set in $\mathcal{W}$. Consider $v_1 \in P$, if v_1 has at least two neighbors in Y then then the above arguments imply that all these vertices must be in a single witness set of Y and therefore Y must contain a big witness set. Otherwise, v_1 only has only one neighbor in Y. If all the vertices in Y form small bags in $\mathcal{W}$, then the corresponding vertices in T are of degree two. As Y is a connected set in G, this implies that Y is a simple path in G. Then $Y \cup P$ is also a simple path in G, with endpoints $y_1 \in Y$ and $v_\ell \in P$. Let t_1 be the vertex of T corresponding to y_1. As t_1 is of degree two in T, $N(y_1) \setminus Y \subseteq W(t')$ for some $t' \in T$ that is a neighbor of t_1. Further, there is some $Y' \in \mathcal{X}$ such that $W(t') \subseteq Y'$. Now, observe that $(Y \cup P), Y'$ and Z satisfy the premise of this lemma, and this contradicts the maximality of P. We may similarly conclude that Z must contain a big witness set. This concludes the proof of the lemma. □

Lemma. *If a colored component X in $\mathcal{X}$ is monochromatic with color from $\{1, 2, 3\}$ after exhaustive application of two re-coloring rules then X contains a big witness set.*

Proof. Let T_B and T_S are set of vertices in T which corresponds to big witness sets and small witness set respectively.Consider internal-cactus T_I obtained from T by removing any vertex which does not lie on a path between two distinct vertices in T_B. By Observation 1 and 2, all vertices in $V(T) \setminus V(T_I)$ have maximum degree two in T. Hence $T - V(T_I)$ is a collection of isolated vertices and induced paths.

Let Q be a connected component of $T - V(T_I)$ which is adjacent to vertex t. We argue that vertices in G which are contained in witness set corresponding to vertices in Q are either isolated vertex in $G - W(t)$ or simple path P in G whose neighborhood is contained in $W(t)$. By the definition of T_I, every component of $T - V(T_I)$ has edges to exactly one vertex in T_B. Hence, for the component

Q there is some t in T_B such that Q is either a leaf incident on t, or $Q \cup t$ is a pendant cycle in T. Since every witness set correspond to vertices in Q is a singleton set, vertices in this witness set induces a path P in G. Note that P is a connected component of $G - W(t)$. All such vertices are re-colored to 4 by first rule of recoloring.

Let $Q = (t_i, x_1, x_2, \ldots, x_q, t_j)$ be a path in T_I, where $t_i, t_j \in T_B$ and $W(x_i)$ is a small bag for each x_i. As each $W(x_i)$ is a small bag and x_i have degree 2 in T, vertices in G which are contained in witness set corresponding to vertices in $V(Q) \setminus \{t_i, t_j\}$ induces a simple path P in G whose neighborhood is contained in $W(t_i)$ and $W(t_j)$ which are big witness set. All such vertices are re-colored to 5 by second rule of recoloring.

This implies exhaustive application of two re-coloring rules identify almost all the vertices in G that form small bags in T. The only exceptions being those vertices that are contained in some component X of ϕ which also contains a big witness set. After re-coloring if there is a colored component which have a color from $\{1, 2, 3\}$ then it have witness sets which corresponds to vertices in T_B and hence it must contain a big witness set. $\square$

6.4 Missing Proofs from Section 5

Lemma. *For a colored component X in $\mathcal{X}$, if $W(t)$ is the big witness set contained in X then $W(t)$ is a connected core of $G[\hat{X}]$.*

Proof. Since $W(t)$ is a witness set, by definition $G[W(t)]$ is connected. For the sake of contradiction assume that $W(t)$ is not a core of $G[\hat{X}]$. This implies that at least one connected component C of $G[\hat{X}] \setminus W(t)$ is neither a simple path nor a isolated vertex. Hence, C contains at least 3 vertices and there exists a vertex $x \in C$ such that $d_{G[\hat{X}]}(x) \geq 3$ and it is adjacent to at least two vertices in C. If $x \in \hat{X} \setminus X$, then by Lemma 3, it is contained in a small bag. Otherwise, $x \in X \setminus W(t)$ and it is again contained in a small bag. This implies that $W(t_x) = \{x\}$ and $d_T(t_x) \geq 3$. So by Observation 1, x is a cut-vertex in T. However, this contradicts Observation 2 which states that that every cut-vertex in T corresponds to big witness set. $\square$

Lemma. *If there exists v in $N_G(X)$ such that v is colored 5 then $N_G(v) \cap X$ is contained in a big witness set of X.*

Proof. If v is colored 5, by Lemma 4, v is contained in a simple path P in G between two components $X, X' \in \mathcal{X}$, such that all the vertices of P are in small bags in $\mathcal{W}$. Furthermore, both X and X' must contain big witness sets in $\mathcal{W}$. Let $W(t')$ be the big witness set contained in X'. Let P' be a path from $W(t')$ to an endpoint of P, whose internal vertices are in X'. Assume that there exists $x \in W(t) \setminus N(v)$. Consider a path Q from $W(t)$ to x which is contained entirely in $G[X]$. Since $xv \in E(G)$, we have that Q along with edge xv and paths P, P' form a path from $W(t)$ to $W(t')$ in G. This path in G gives a path between t and t' in T, such that all the internal vertices of this path correspond to small bags in $\mathcal{W}$. Notice that every vertex on the path from $W(t)$ to x, has same color as that of $W(t)$. All these vertices are in small bags. This is a contradiction to the fact that ϕ is compatible coloring with $\mathcal{W}$. Hence our assumption is wrong and the claim follows. $\square$

Lemma. *Let $X, Y \in \mathcal{X}$ be two components which contain big witness sets, say, W_X and W_Y, respectively. Then, $N(X) \cap Y \subseteq W_Y$ and $N(Y) \cap X \subseteq W_X$.*

Proof. If $E(X, Y) = \emptyset$ then the statement is vacuously true. Assume that $x \in (N(Y) \setminus W_X) \cap X$, and let t, t' be the vertices of T corresponding to the big witness sets W_X, W_Y respectively. Since X is connected, there exists a path between W_X and x which is entirely contained in X. As X may contain only one big witness set, x lies in a small bag in $\mathcal{W}$. This implies that there is path between t and t' in T (via x) such that the neighbor of t has the same color as vertices in W_X. This is a contradiction to the fact that ϕ is compatible coloring with $\mathcal{W}$. $\square$

Lemma. *For a colored component X in $\mathcal{X}$ let $W(t)$ be the big witness set contained in X. If t_1 is a neighbor of t in the internal cactus T_I of T then all the vertices in $N_G(W(t_1)) \cap X$ has been marked by Marking Scheme 1.*

Proof. If $W(t_1)$ is a small bag then by Lemma 5 it is recolored to 5. By Observation 1, t_1 is not a cut-vertex in T. As t_1 is part of the internal cactus of T, it must lie on some simple path in T between vertices t and t_2, where $W(t_2)$ is a big witness set in $\mathcal{W}$. By Lemma 4 and 5, $W(t_1)$ gets color 5. Therefore, $N_G(W(t_1)) \cap X$ has been marked. If $W(t_1)$ is a big witness set, then it is contained in a component $X' \in \mathcal{X}$ as ϕ is a compatible coloring. Therefore, it has also been marked. $\square$

Lemma. *Let set $\mathcal{X}'$ be obtained from $\mathcal{X}$ by exhaustive application of Pruning Operations. If F^* be a union of spanning trees of graph induced on colored components in $\mathcal{X}^*$ then G/F^* is a cactus and $|F'| = |F|$.*

Before moving to proof of the Lemma, we mention following two observations.

Observation 4. *Let X, Y are colored component in $\mathcal{X}$. If these colored components contain big witness sets in $\mathcal{W}$ then $\hat{X} \cap \hat{Y} = \emptyset$.*

Proof. As X and Y contain big witness sets and are components of ϕ, by definition they are disjoint. Now observe that X forms a separator between $\hat{X}$ and $V(G) - \hat{X}$, and similarly for Y. Hence, $\hat{X} \cap \hat{Y} = \emptyset$. $\square$

For $Y \subseteq V(G)$, let $\mathcal{W}_Y$ be the collection of witness sets in $\mathcal{W}$ that intersect Y.

Observation 5. *Let $X \in \mathcal{X}$ and $\hat{X}$ be as defined above. Then for all witness sets $W \in \mathcal{W}_{\hat{X}}$, we have $W \subseteq \hat{X}$.*

Proof. Suppose not. Then there is a vertex $y \in W \setminus \hat{X}$, and a vertex $x \in W \cap \hat{X}$. Hence, W is a big witness set, and there is $Y \in \mathcal{X}$ such that $W \subseteq Y$. As W is a connected set, there is a path P in $G[W]$ between y and x. And observe that X is a separator between $\hat{X}$ and $V(G) \setminus \hat{X}$ and it intersects the path P at a vertex x'. Hence $W \cap X \neq \emptyset$. Now, as ϕ is a compatible coloring we have $Y \in \mathcal{X}$ such that $W \subseteq Y$ and it is distinct from X. However, we have that $x' \in X \cap Y$, which contradicts the fact that any two components of ϕ are disjoint. $\square$

Proof (of Lemma 10). First consider the case when we do the operation for only one colored component X in $\mathcal{X}$. In other words, we fix a $X \in \mathcal{X}$ and let $\mathcal{W}' = (\mathcal{W} \setminus W_{\hat{X}}) \cup \mathcal{W}'_{\hat{X}}$ where $\mathcal{W}'_{\hat{X}} = \{Z\} \cup \{\{v\} | \ v \in \hat{X} \setminus Z\}$ and Z is connected core of $G[\hat{X}]$ which contains all marked vertices. Let T' be the graph obtained by contracting F' in G. We claim that T is a cactus and $|F'| \leq |F|$. Let t_z be the vertex in T' obtained from Z, and $W(t)$ be the big witness set in $\mathcal{W}$ that is a subset of X. Let T_X be the induced subgraph of T that contains all the vertices obtained from $\mathcal{W}_{\hat{X}}$, and by Observation 5 they form a partition of $\hat{X}$. Let T'_X be the induced subgraph of T' obtained from $\mathcal{W}'_{\hat{X}}$. Now observe that $\mathcal{W} \setminus \mathcal{W}_{\hat{X}} = \mathcal{W}' \setminus \mathcal{W}'_{\hat{X}}$, which implies that $T - V(T_X)$ is isomorphic to $T' - V(T'_X)$. By Lemma 9, Z contains every vertex in X that has a neighbor in $V(G) \setminus \hat{X}$, and therefore $N_{T'}(t_Z) \setminus V(T'_X) = N_T(t) \setminus V(T_X)$. Hence the induced subgraphs $(T - V(T_X)) \cup t$ and $(T' - V(T'_X)) \cup t_z$ are isomorphic as well. Now Z is a connected-core of $G[\hat{X}]$, hence T'_X is a cactus. And observe that t_z is the only vertex in T'_X that may have a neighbor in $T' - V(T'_X)$, i.e. it is a cut-vertex in T' that separates T'_X from $T' - V(T'_X)$. Therefore we conclude that $(V(T) \setminus V(T_X)) \cup V(T'_X)$ induces a cactus T', which means F' is also a solution. And by Lemma 6, $W(t)$ is a connected-core of $\hat{X}$, and F is a spanning forest of $\mathcal{W}$. As Z is a minimum connected core of $\hat{X}$ and F' forms a spanning forest of $\mathcal{W}'$, we have that $|F'| \leq |F|$.

Next, we consider all the sets $X \in \mathcal{X}$ and fix an arbitrary order among them. By Observation 4, for any two sets X, Y in $\mathcal{X}$, $\hat{X}, \hat{Y}$ are disjoint. Now, starting with a given solution F, we apply the above arguments for each X in $\mathcal{X}$ one by one. Here, we update the set F to F' each time, before proceeding to the next X. Observe that F' obtained at the end of the process, say F^*, is a solution, i.e. G/F^* is a cactus, and $|F^*| \leq |F|$. Since F was optimum solution it follows that $|F| \leq |F^*|$ which concludes the proof. $\square$

Theorem. *There is an algorithm that given a connected graph G and a subset X of its vertices, computes a minimum connected core of G which has at most k vertices and contains X in $\mathcal{O}^*(6^k)$ time if one such exists in the graph.*

Proof. Consider a connected core Z^* of G, and recall that G/Z^* is a cactus. Therefore, if (u, v, w) is a path in $G - Z^*$ then v must be a vertex of degree 2 in G. Furthermore, if (x, y) is an edge in $G - Z^*$, then it follows that x and y have degree 2 or more in G. Our algorithm is based on these observations. We first construct a core of G via a branching algorithm. Then at each leaf of the search tree, we extend the core constructed by the branching algorithm to a connected set by applying an algorithm for the Steiner Tree problem.

Let us delve into the details of our algorithm. Let Z denote a solution to the instance. Initially we set Z to X and decrease k by $|X|$. The following branching rule derived from the first observation.

Branching Rule 1. *If there is a path (u, v, w) in $G - Z$ such that $|N_G(v)| \geq 3$, then branch into three cases where each of u or v or w is added to Z. Decrease k by 1 in each of the branches.*

Observe that when this rule is no longer applicable, all vertices of $G - Z$ have degree at most 2. Hence the components of $G - Z$ are simple paths in G, or isolated vertices. Next, we have the following reduction rule that follows from the second observation.

Reduction Rule 1. *If there is an edge uv in $G - Z$ with $|N_G(v)| = 1$, then add u into Z and reduce k by 1.*

Since the only neighbor of v is u, the edge uv cannot be part of a simple path in G whose both endpoints have neighbours in Z. Now, if there exists an optimal solution Z^* that does not contain u, then $v \in Z^*$ and $Z' = (Z^* \setminus \{v\}) \cup \{u\}$ is also a connected core of G. This justifies the correctness of the rule.

We apply the above rules exhaustively, and consider the search tree constructed. Note that each node of the search tree is labeled with either a triple (u, v, w) indicating that the Branching rule 1 was applied, or an edge (x, y) indicating that Reduction rule 1 was applied at this node. If at any node in the search tree, k is 0 and the set Z is not a connected core of G, we abort that node. If all the leaves of the current search tree are aborted, then we output NO as a solution to this instance.

Next, we claim that if none of the rules are applicable at a leaf of the search tree, then the corresponding Z is a core of G. Assume to the contrary that Z is not a core of G. Then there is a component C of $G - Z$ that is neither an isolated vertex, nor it is a simple path in G whose both endpoints have neighbours in Z. Hence such a C has at least two vertices. Furthermore recall that the branching-rule is not applicable at this node of the search tree, and therefore all vertices in $G - Z$ have maximum degree 2. Consider the case when C is a cycle in $G - Z$. As G is connected, C has a vertex v that has a neighbour in Z. Let u and w be the neighbours of v in C. Then, it follows that (u, v, w) is a path in $G - Z$ with $|N_G(v)| \geq 3$. However, this leads to a contradiction as Branching rule 1 is not applicable. Now, consider the case when C is a path in $G - Z$ with end-points u and v. If there is an internal vertex on this path that has a neighbor in Z, then as before, we obtain a contradiction. Hence, C is a simple path in G, with end-points u and v. As Z is not a core of the connected graph G, one of u or v has no neighbour in Z, i.e. it is a vertex of degree 1 in G. But then, Reduction rule 1 is applicable, which is a contradiction. Hence Z must be a core of the graph G.

However, as Z may not be connected in G, we may have to add additional vertices to ensure connectivity. Observe that this can be achieved by computing a minimum Steiner Tree for Z in G. Given a graph G and a set S of vertices of G, the Steiner Tree problem is the task of computing a minimum cardinality connected subgraph that contains S. This problem is known to admit an algorithm with $\mathcal{O}^*(2^{|S|})$ running time [17]. The above algorithm computes a minimum cardinality connected set of vertices, $Z' \supseteq Z$, in time $\mathcal{O}(2^k)$. Observe that Z' is a connected-core of G, as $G - Z$ is a collection of isolated vertices and simple paths in G. Let $\bar{Z}$ be the minimum cardinality connected-core over all the leaves of the search-tree. If $|\bar{Z}| \leq k$, we output $\bar{Z}$ and otherwise we output NO as the solution to the instance.

Let us now argue the correctness of this algorithm. Assume Z^* is an optimal solution of size at most k. We claim that above algorithm finds a connected core $\bar{Z}$ such that $|\bar{Z}| \leq |Z^*|$. To argue this, we associate a path on the search tree of branching algorithm to the set Z^*.

Now consider an internal node in search tree that is labeled with (a, b, c). Since Branching rule 1 is applied at this node, we have that (a, b, c) is a path in $G - Z$ and $|N_G(b)| \geq 3$. As Z^* is a core of G, at least one of a, b, c must be present in it. Similarly, for any node labeled with an edge (x, y), one of these vertices, say y, is of degree 1 in G, and hence Z^* must contain one of them. Recall that, by previous arguments, we may assume $x \in Z^*$. Hence, we start from the root of the search tree and navigate to a leaf along the choices consistent with Z^*. If more than one choices are consistent with Z^*, we arbitrarily pick one of the them and proceed. Consider the set $\tilde{Z}$ obtained at the leaf via this navigation consistent with Z^* from the root-node of the search tree. Clearly $\tilde{Z} \subseteq Z^*$ and $\tilde{Z}$ is a core (not necessarily connected) of G. Let T be an optimal solution for an instance of $(H, \tilde{Z})$ of Steiner Tree as defined above. Since Z^* is a connected core of G and $\tilde{Z} \subseteq Z^*$ we know that $Z^* \setminus \tilde{Z}$ is a solution to this Steiner Tree instance. By the optimality of T, $|T| \leq Z^* \setminus \tilde{Z}$ and hence $\bar{Z} = \tilde{Z} \cup T$ is a desired solution.

Let us now consider the running time of this algorithm. We measure the progress of the algorithm measure is the solution size k. At each application of the Branching rule 1, we have a three-way branch and the measure drops by 1 branching vector is $(1, 1, 1)$. This leads to the recurrence $T(k) \leq 3T(k-1)$ which solution is $\mathcal{O}^*(3^k)$. Next, at each leaf of the search tree, we run the algorithm for finding a minimum Steiner tree, which runs in time $\mathcal{O}^*(2^k)$. Therefore, the overall running time is $\mathcal{O}^*(6^k)$. □

6.5 Missing Proofs from Section 6

Theorem. *There is an one-sided error Monte Carlo algorithm with false negatives which solves* Cactus Contraction *in time* $c^k n^{\mathcal{O}(1)}$ *on* 2*-connected graphs. It returns correct answer with constant probability.*

Proof. Consider an algorithm which uses Algorithm 6.1 as subroutine and runs it 3^{4k} many times. If any of these runs return a solution F, then the algorithm returns F otherwise after all iterations are over, it returns NO. This finishes the description of the algorithm.

Let us argue the correctness of this algorithm. Consider a graph G which is k-contractible to a cactus T and $\mathcal{W}$ is T-witness structure of G. Let ϕ is a 3-coloring of G which is compatible with $\mathcal{W}$. To argue the correctness, we first claim that given graph G and a compatible coloring ϕ, Algorithm 6.1 returns a correct answer. But this immediately follows from Lemma 10.

Since Algorithm 6.1 returns a solution only if it has found a witness structure with desired properties, it never returns false positives. We argue that our algorithm returns a solution, if there is any, with constant probability.

Let $\phi : V(G) \to \{1, 2, 3\}$ be a coloring where colors are chosen uniformly at random for each vertex. The total number of vertices contained in big witness sets of $\mathcal{W}$ is at most $2k$. Also, cactus T can have at most k big witness sets and hence at most k cut-vertices. (Note that, we consider all the vertices of T which correspond to big witness set as a cut-vertex, even if they are not actually a cut-vertex in T, as this doesn't affect our arguments in any way.) By Observation 3, there are at most $4k$ vertices which lie on a path between two cut-vertices and are adjacent to big witness sets. Therefore, by the definition of a compatible coloring, the probability that a random 3-coloring compatible with $\mathcal{W}$ is at least $\frac{1}{3^{6k}}$. Since the algorithm runs 3^{6k} many iterations of Algorithm 6.1, probability that none of these colorings which is generated uniformly at random is compatible with $\mathcal{W}$ is at most $(1 - \frac{1}{3^{6k}})^{3^{6k}} < 1/e$. Hence Algorithm 6.1 returns a solution on positive instances with probability at least $1 - 1/e$. Each iteration of Algorithm 6.1 takes $6^k \cdot n^{\mathcal{O}(1)}$ time and hence the total running time of the algorithm is $c^k \cdot n^{\mathcal{O}(1)}$ for a fixed constant c. □

Theorem. *There is an one-sided error Monte Carlo algorithm with false negatives which solves* Cactus Contraction *in time $c^k n^{\mathcal{O}(1)}$. It returns correct answer with constant probability.*

Proof. If input graph G is not 2-connected, we find its block decomposition [8] is $n^{\mathcal{O}(1)}$ time. If any any of 2-connected component of G is a cactus then we delete it and work in remaining graph. Let $G_1, G_2, \dots, G_q$ are two connected components of G such that G_i is not a cactus for all $i \in [q]$. If $q \geq k + 1$ then we return NO as at least one edge needs to be contracted in each of these 2-connected components. We now consider the case when $q \leq k$.

For each G_i and each possible values k_j between 1 and k, we run algorithm presented in Theorem 2 $3 \log k$ times on instance (G_i, k_j). Since there are at most k^2 such pairs, algorithm in Theorem 2 has been run at most $3k^2 \log k$ time. If algorithm returns NO for all the values of k_j for some G_i then we return NO. Otherwise let k'_i be the smallest value for which algorithm returns a solution for G_i. Since algorithm in Theorem 2 returns no false positive, G_i is k-contractible to a cactus. On the other hand if (G_i, k_i) is an YES instance of Cactus Contraction then probability that no run will output right answer is at most $(\frac{1}{e})^{3 \log k} = \frac{1}{k^3}$. Since there are at most k^2 pairs (G_i, k_j), and by the union bound on probabilities, the probability that there is a pair (G_i, k_j) for which the algorithm returns false negative is upper bounded by $k^2 \cdot \frac{1}{k^3} \geq \frac{1}{k}$. If such a failure does not occur, then for every i we have that k'_i is exactly the smallest value of k_j such that G_i is k_i-contractible to a cactus. Finally, the algorithm answers YES only if $\sum_{i=1}^{q} k_i' \leq k$, and answers NO otherwise. The correctness of this algorithm follows from Proposition 1. Consequently, the algorithm cannot give false positives, and it may give false negatives with probability at most $1/k \leq 1/q \leq 1/2$, where the two inequalities follows from the assumption that $2 \leq q \leq k$. $\square$

6.6 Derandomization

We can derandomize our algorithms by constructing a family of coloring function, that is derived from a perfect hash family.

Definition 5 ([16]). *A (n, k)-universal set is a family $\mathcal{H}$ of subsets of $[n]$ such that for any $S \subseteq [n]$ of size at most k, $\{S \cap H \mid H \in \mathcal{H}\}$ contains all subsets of S.*

Lemma 12 ([16]). *For any $n, k \geq 1$, we can construct a (n, k)-universal set of size $2^k k^{\mathcal{O}(\log k)} \log n$ in time $2^k k^{\mathcal{O}(\log k)} n \log n$.*

Now suppose that G is contractible to a cactus T and a coloring ϕ, that is compatible with T. Recall that, the total number of vertices contained in all big bags is at most $2k$. Further more, by Observation 3, the number of vertices that are adjacent to a big bag and are mapped to internal cactus of T is at most $4k$. Let S denote the set of all these vertices in G, and note that we can ensure $|S| = 6k$ by arbitrarily adding some extra vertices to it. Observe that, ϕ gives a partition of the set S into 3 parts, say S_1, S_2, S_3. For the remainder of this section, let us fix the cactus T and the compatible coloring ϕ. From the definition of compatible coloring, any coloring function ψ which partitions of S into S_1, S_2, S_3 is a compatible coloring. We say that the coloring function ϕ and ψ agree on S.

Lemma 13. *Let ϕ be a compatible coloring of G. Then there is a family of coloring functions, $\mathcal{F} = \{f : V(G) \to [3]\}$, such that there $\psi \in \mathcal{F}$ that agrees with ϕ on S. This family has size $4^{6k} k^{\mathcal{O}(\log k)} \log^2 n$ and it can be constructed in time $4^{6k} k^{\mathcal{O}(\log k)} n \log n$.*

Proof. Let $\mathcal{H}$ be a $(n, 6k)$-universal set, that is constructed by Lemma 12. We define a family of partitions of $V(G)$ as follows.

$$\mathcal{F}' = \{(A, B, C) \mid A \in \mathcal{H},\ B = Y \setminus A \text{ where } Y \in \mathcal{H},\ C = V(G) \setminus Y\}$$

Observe that $\mathcal{F}'$ can be constructed by considering each pair of sets in $\mathcal{H}$. We claim that there is a triple $(A, B, C) \in \mathcal{F}'$ such that $S \cap A = S_1$, $S \cap B = S_2$ and $S \cap C = S_3$. Indeed, since $\mathcal{H}$ is a $(n, 6k)$-universal-set, there is some set $Y \in \mathcal{H}$ such that $S \cap Y = S_1 \cup S_2$, and there is some $A \in \mathcal{H}$ such that $A \cap S = S_1$. Hence, $S \cap (Y - A) = (S \cap Y) \setminus (S \cap A) = S_2$. We can easily convert the family $\mathcal{F}'$ into a family of coloring functions, where for each $(A, B, C) \in \mathcal{F}'$ maps all vertices in A, B, C to $1, 2, 3$ respectively. Now it is clear that if ϕ partitions S into S_1, S_2, S_3, then there is a function $\psi \in \mathcal{F}$, which also partitions S into S_1, S_2, S_3. Observe that the family $\mathcal{F}$ has size $4^{6k} k^{\mathcal{O}(\log k)} \log^2 n$, since the family $\mathcal{H}$ has size $2^{6k} k^{\mathcal{O}(\log k)} \log n$. $\square$

Using the above coloring family (instead of a random coloring) in our algorithm for Cactus Contraction establishes Theorem 3.