

Computing Distances on Graph Associahedra Is Fixed-Parameter Tractable

Luís Felipe I. Cunha ✉ 🏠 

Instituto de Computação, Universidade Federal Fluminense, Brasil

Ignasi Sau ✉ 🏠 

LIRMM, Université de Montpellier, CNRS, France

Uéverton S. Souza ✉ 🏠 

Instituto de Computação, Universidade Federal Fluminense, Brasil

IMPA - Instituto de Matemática Pura e Aplicada, Rio de Janeiro, Brasil

Mario Valencia-Pabon ✉ 🏠 

Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

Abstract

An elimination tree of a connected graph G is a rooted tree on the vertices of G obtained by choosing a root v and recursing on the connected components of $G - v$ to obtain the subtrees of v . The graph associahedron of G is a polytope whose vertices correspond to elimination trees of G and whose edges correspond to tree rotations, a natural operation between elimination trees. These objects generalize associahedra, which correspond to the case where G is a path. Ito et al. [ICALP 2023] recently proved that the problem of computing distances on graph associahedra is NP-hard. In this paper we prove that the problem, for a general graph G , is fixed-parameter tractable parameterized by the distance k . Prior to our work, only the case where G is a path was known to be fixed-parameter tractable. To prove our result, we use a novel approach based on a marking scheme that restricts the search to a set of vertices whose size is bounded by a (large) function of k .

2012 ACM Subject Classification Theory of computation → Fixed parameter tractability; Mathematics of computing → Combinatorics

Keywords and phrases graph associahedra, elimination tree, rotation distance, parameterized complexity, fixed-parameter tractable algorithm, combinatorial shortest path, reconfiguration

Digital Object Identifier 10.4230/LIPIcs.ICALP.2025.63

Category Track A: Algorithms, Complexity and Games

Related Version The full version of this article is permanently available here: <http://arxiv.org/abs/2504.18338>

Funding *Luís Felipe I. Cunha*: FAPERJ-JCNE (E-26/201.372/2022) and CNPq-Universal (406173/2021-4).

Ignasi Sau: French project ELiT (ANR-20-CE48-0008-01).

Uéverton S. Souza: CNPq (312344/2023-6), and FAPERJ (E-26/201.344/2021).

Mario Valencia-Pabon: French project ABYSM (ANR-23-CE48-0017).

1 Introduction

Given a connected and undirected graph G , an *elimination tree* T of G is any rooted tree that can be defined recursively as follows. If $V(G) = \{v\}$, then T consists of a single root vertex v . Otherwise, a vertex $v \in V(G)$ is chosen as the root of T , and an elimination tree is created for each connected component of $G - v$. Each root of these elimination trees of $G - v$ is a child of v in T . For a disconnected graph G , an *elimination forest* of G is the disjoint union of elimination trees of the connected components of G . Equivalently, an elimination



© Luís Felipe I. Cunha, Ignasi Sau, Uéverton S. Souza, and Mario Valencia-Pabon; licensed under Creative Commons License CC-BY 4.0

52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025).

Editors: Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis


Article No. 63; pp. 63:1–63:19



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany





■ **Figure 1** A graph G and two of its elimination trees T and T' , where the second one is obtained from the first one by the rotation of edge uv (in red).

forest of a graph G is a rooted forest F (that is, a forest with a root in every connected component) on vertex set $V(G)$ such that for each edge $uv \in E(G)$, vertex u is an ancestor of vertex v in F , or vice versa.

Figure 1 illustrates an example of two elimination trees T and T' of a graph G . With slight (and standard) abuse of notation, we use the same labels for the vertices of a graph G and any of its elimination trees. Note that an elimination tree is unordered, i.e., there is no ordering associated with the children of a vertex in the tree. Similarly, there is no ordering among the elimination trees in an elimination forest.


Elimination trees have been studied extensively in various contexts, including graph theory, combinatorial optimization, polyhedral combinatorics, data structures, or VLSI design; see the recent paper by Cardinal, Merino, and Mütze [7] and the references therein. In particular, elimination trees play a prominent role in structural and algorithmic graph theory, as they appear naturally in several contexts. As a relevant example, the *treedepth* of a graph G is defined as the minimum height of an elimination forest of G [26].

Given a class of combinatorial objects and a “local change” operation between them, the corresponding *flip graph* has as vertices the combinatorial objects, and its edges connect pairs of objects that differ by the prescribed change operation. In this article, we focus on the case where this class of combinatorial objects is the set of elimination forests of a graph G . For these objects, the commonly considered “local change” operation is that of *edge rotation* defined as follows, where we suppose for simplicity that G is connected. Given an elimination tree T of a graph G , the *rotation* of an edge $uv \in E(T)$, with u being the parent of v , creates another elimination of G , denoted by $\text{rot}(T, uv)$, obtained, informally, by just swapping the choice of u and v in the recursive definition of T (that is, in the so-called *elimination ordering*), and updating the parent of the subtrees rooted at v accordingly; see Figure 2 for an illustration. The formal definition can be found in Subsection 3.1 (cf. Definition 2).

For example, in Figure 1, $T' = \text{rot}(T, uv)$. The definition of the rotation operation clearly implies that it is self-inverse with respect to any edge, that is, for any elimination tree T of a graph G and any edge $uv \in E(T)$, it holds that $T = \text{rot}(\text{rot}(T, uv), vu)$. The *rotation distance* between two elimination trees T, T' of a graph G , denoted by $\text{dist}(T, T')$, is the minimum number of rotations it takes to transform T into T' . The self-invertibility property of rotations discussed above implies that $\text{dist}(T, T') = \text{dist}(T', T)$.

It is well known that for any graph G , the flip graph of elimination forests of G under tree rotations is the skeleton of a polytope, referred to as the *graph associahedron* $\mathcal{A}(G)$ and that was introduced by Carr, Devadoss, and Postnikov [11, 16, 29]. For the particular cases of G being a complete graph, a cycle, a path, a star, or a disjoint union of edges, $\mathcal{A}(G)$ is the permutahedron, the cyclohedron, the (standard) associahedron, the stellohedron, or the hypercube, respectively; see the introduction of [7] for nice figures to illustrate these objects.

Graph associahedra naturally generalize associahedra, which correspond to the particular



■ **Figure 2** On the left: An elimination tree T of a graph G with adjacent vertices u and v . Vertex v has four subtrees, and two of them, namely T_2 and T_3 , contain vertices adjacent to vertex u in G . On the right: Elimination tree resulting from T by applying the rotation of uv . Since both $G[V(T_2) \cup \{u\}]$ and $G[V(T_3) \cup \{u\}]$ are connected, T_2 and T_3 become subtrees of u in $\text{rot}(T, uv)$.

case where G is a path. As mentioned in [7], the associahedron has a rich history and literature, connecting computer science, combinatorics, algebra, and topology [18, 22, 30, 32]. See the introduction of the paper by Ceballos, Santos, and Ziegler [12] for a historical account. In an associahedron, each vertex corresponds to a binary tree over a set of n elements, and each edge corresponds to a rotation operation between two binary trees, an operation used in standard online binary search tree algorithms [1, 17, 33]. Binary trees are in bijection with many other Catalan objects such as triangulations of a convex polygon, well-formed parenthesis expressions, Dyck paths, etc. [34]. For instance, in triangulations of a convex polygon, the rotation operation maps to another simple operation, known as a flip, which removes the diagonal of a convex quadrilateral formed by two triangles and replaces it by the other diagonal.

Related work. Distances on graph associahedra have been object of intensive study. Probably, the most studied parameter is the diameter, that is, the maximum distance between two vertices of $\mathcal{A}(G)$. A number of influential articles either determine the diameter exactly, or provide lower and upper bounds, or asymptotic estimates, for the cases where the underlying graph G is a path [30, 32], a star [25], a cycle [31], a tree [6, 25], a complete bipartite graph [8], a caterpillar [3], a trivially perfect graph [8], a graph in which some width parameter (such as treedepth or treewidth) is bounded [8], or a general graph [25].

Our focus is on the algorithmic problem of determining the distance between two vertices of $\mathcal{A}(G)$, or equivalently, determining the rotation distance between two given elimination trees of a graph G . There are very few cases where this problem is known to be solvable in polynomial time, namely when G is a complete graph (folklore), a star [9], or a complete split graph [9]. The complexity of the case where G is a path is a notorious long-standing open problem. On the positive side, for G being a path, there exist a polynomial-time 2-approximation algorithm [15] and several fixed-parameter tractable (FPT) algorithms when the distance is the parameter [14, 20, 21, 23, 24]. It is worth mentioning that there are some hardness results on generalized settings [2, 23, 27] and polynomial-time algorithms for some type of restricted rotations [13].

Cardinal et al. [5] asked whether computing distances on general graph associahedra is NP-hard. Very recently, this question was answered positively by Ito et al. [19].

Our result. The NP-hardness result of Ito et al. [19] (see also [10]) paves the way for studying the parameterized complexity of the problem of computing distances on graph

associahedra. Thus, in this article we are interested in the following parameterized problem, where we consider the natural parameter, that is, the desired distance.

ROTATION DISTANCE

Instance: A graph G , two elimination trees T and T' of G , and a positive integer k .
Parameter: k .
Question: Is the rotation distance between T and T' at most k ?

As mentioned above, ROTATION DISTANCE was known to be polynomial-time solvable on complete graphs, stars, and complete split graphs [9], and FPT algorithms were only known on paths [14, 20, 21, 23, 24]. In this article we vastly generalize the known results by providing an FPT algorithm to solve the ROTATION DISTANCE problem for a general input graph G . More precisely, we prove the following theorem.

► **Theorem 1.** *The ROTATION DISTANCE problem can be solved in time $f(k) \cdot |V(G)|$, with $f(k) = k^{k \cdot 2^{2^{\dots 2^{\mathcal{O}(k^2)}}}}$, where the tower of exponentials has height at most $(3k + 1)4k = \mathcal{O}(k^2)$.*

In particular, Theorem 1 yields a linear-time algorithm to solve ROTATION DISTANCE for every fixed value of the distance k . To the best of our knowledge, this is the first positive algorithmic result for the general ROTATION DISTANCE problem (i.e., with no restriction on the input graph G), and we hope that it will find algorithmic applications in the many contexts where graph associahedra arise naturally [7, 11, 16, 19, 25, 29]. Our result can also be seen through the lens of the very active area of the parameterized complexity of graph reconfiguration problems; see [4] for a recent survey.

Organization. In Section 2 we present an overview of the main ideas of the algorithm of Theorem 1, which may serve as a road map to read the rest of the article. In Subsection 3.1 we provide some preliminaries and fix our notation, in Section 3 we formally present our FPT algorithm (split into several subsections), and in Section 4 we discuss several directions for further research. Due to space limitations, the proofs of the statements marked with “(★)” can be found in the full version available online.

2 Overview of the main ideas of the algorithm

Our approach to obtain an FPT algorithm to solve ROTATION DISTANCE is novel, and does not build on previous work. Given two elimination trees T and T' of a connected graph G and a positive integer k , our goal is to decide whether there exists what we call an ℓ -rotation sequence σ from T to T' , for some $\ell \leq k$, that is, an ordered list of ℓ edges to be rotated in order to obtain T' from T , going through the intermediate elimination trees $T_1, \dots, T_{\ell-1}$ (all of the same graph G); see Subsection 3.1 for the formal definition. At a high level, our approach is based on identifying a subset of *marked vertices* $M \subseteq V(T)$, of size bounded by a function of k , so that we can assume that the desired rotation sequence σ uses only vertices in M . Once this is proved, an FPT algorithm follows directly by applying brute force and guessing all possible rotations using vertices in M .

A crucial observation (cf. Observation 3) is that a rotation may change the set of children of at most three vertices (but the parent of arbitrarily many vertices, such as the roots of T_2 and T_3 in Figure 2). Motivated by this, we say that a vertex $v \in V(T)$ is (T, T') -children-bad

if its set of children in T is different from its set of children in T' . By Observation 3, we may assume (cf. Observation 5) that we are dealing with an instance in which the number of (T, T') -children-bad vertices is at most $3k$.

In a first step, we prove (cf. Lemma 7) that we can assume that the desired sequence σ of at most k rotations to transform T into T' uses only vertices lying in the union of the balls of radius $2k$ around (T, T') -children-bad vertices of T , which we denote by B_{cb} . The proof of Lemma 7 exploits, in particular, the fact that a rotation may increase or decrease vertex distances (in the corresponding trees) by at most one (cf. Equation 1). This is then used to show that if a rotation uses some vertex outside of B_{cb} , then it can be “simplified” into another one that does not (cf. Figure 3).

By Lemma 7, we restrict henceforth to rotations using only vertices in B_{cb} . We can consider each connected component Z of $T[B_{cb}]$, since it can be easily seen that we can assume that there are at most k of them. By definition of B_{cb} , the diameter of such a component Z is $\mathcal{O}(k^2)$ (cf. Equation 3). Thus, the “only” obstacle to obtain the desired FPT algorithm is that the vertices in B_{cb} can have an arbitrarily large degree. Note that in the particular case where the underlying graph G has bounded degree, the maximum degree of any elimination tree of G is bounded, and therefore in that case $|B_{cb}|$ is bounded by a function of k , and an FPT algorithm follows immediately. To the best of our knowledge, this result was not known for graphs other than paths.

Our strategy to deal with high-degree vertices in B_{cb} is as follows. Fix one connected component Z of $T[B_{cb}]$. Our goal is to identify a subset $M_Z \subseteq V(Z)$ of size bounded by a function of k , such that we can restrict our search to rotations using only vertices in M_Z . To find such a “small” set $M_Z \subseteq V(Z)$, we define the notion of *type* of a vertex $v \in V(Z)$, in such a way that the number of different types is bounded by a function of k . Then, we will prove via our marking algorithm that it is enough to keep in M_Z , for each type, a number of vertices bounded again by a function of k .

Before defining the types, we need to define the *trace* of a vertex v in Z . To get some intuition, look at the rotation depicted in Figure 2. Note that, for each of the subtrees $T_1, \dots, T_4$ that are children of v in T , what determines whether they are children of u or v in the resulting subtree is whether some vertex in T_i is adjacent to u or not. Iterating this idea, if we are about to perform at most k rotations starting from T , then the behavior of such a subtree T_i , assuming that no vertex of it is used by a rotation, is determined by its neighborhood in a set of ancestors of size at most the diameter of Z , and this is what the trace is intended to represent. That is, the trace of a vertex v in Z , denoted by $\text{trace}(T, Z, v)$, captures “abstractly” the neighborhood of the whole subtree rooted at v among (the ordered set of) its ancestors within the designated vertex set $Z \subseteq V(T)$; see Definition 11 for the formal definition of trace and Figure 4 for an example. We stress that, when considering the neighborhood in the set of ancestors, we look at the whole subtree $T(v)$ rooted at v , and not only at its restriction to the set Z .

Equipped with the definition of trace, we can define the notion of type, which is somehow involved (cf. Definition 12) and whose intuition behind is the following. For our marking algorithm to make sense, we want that if two vertices v, v' with the same parent (called T -siblings) have the same type (within Z), denoted by $\tau(T, Z, v) = \tau(T, Z, v')$, and an ℓ -rotation sequence σ from T to T' uses some vertex from $T(v)$ but uses no vertex in $T(v')$, then there exists another ℓ -rotation sequence σ' from T to T' that uses vertices in $T(v')$ instead of those in $T(v)$. To guarantee this replacement property, we need a stronger condition than just v and v' having the same trace. Informally, we need them to have the same “variety of traces among their children within Z ”. More formally, this leads to a recursive definition where, in

the leaves of Z (that are not necessarily leaves of T), the type corresponds to the trace, and for non-leaves, the type is defined by the trace and by the number of children of each possible lower type. Note that, a priori, the number of children of a given type may be unbounded, which would rule out the objective of bounding the number of types as a function of k . To overcome this obstacle, the crucial and simple observation is that at most k subtrees rooted at a vertex of T contain vertices used by the desired rotation sequence σ (cf. Lemma 16). This implies that if there are at least $k + 1$ T -siblings of the same type, necessarily the whole subtree of at least one of them, say u , will not be used by σ , implying that u (and its whole subtree) achieves the desired parent in T' without being used by σ , and the same occurs to any other T -sibling of the same type. Thus, keeping track of the existence of at least $k + 1$ such children (regardless of their actual number) is enough to capture this “static” behavior, and allows us to shrink the possible distinct numbers to keep track of to a function of k (cf. Equation 4, where the “min” is justified by the previous discussion). Finally, for technical reasons we also incorporate into the type of a vertex its desired parent in T' , in case it defers from its parent in T' (cf. function `want-parent`($T, T', \cdot$)). See Definition 12 for the formal definition of type and Figure 5 for an example with $k = 2$.

We prove (cf. Lemma 13) that the number of types is indeed bounded by a (large) function $g(k)$ depending only on k , and this function is what yields the upper bound on the asymptotic running time of the FPT algorithm of Theorem 1. Moreover, we show (cf. Observation 14) that the type of a vertex can be computed in time $g(k) \cdot |V(G)|$. We then use the notion of type and the bound given by Lemma 13 to define the desired set $M_Z \subseteq Z$ of size bounded by a function of k . In order to find M_Z , we apply a marking algorithm on Z , that first identifies a set $M_Z^{\text{pre}} \subseteq V(Z)$ of *pre-marked* vertices, whose size is not necessarily bounded by a function of k , and then “prunes” this set M_Z^{pre} in a root-to-leaf fashion to find the desired set of *marked* vertices $M_Z \subseteq M_Z^{\text{pre}}$ of appropriate size. See Figure 6 for an example of the marking algorithm for $k = 1$. We define $M = \cup_{Z \in \text{cc}(T[B_{\text{cb}}])} M_Z$ (where cc denotes the set of connected components), and we call it the set of *marked vertices* of T . We prove (cf. Lemma 15) that the size of M is roughly equal to the number of types, and that the set M can be computed in time FPT.

Once we have our set of marked vertices M at hand, it remains to prove that we can restrict the rotations to use only vertices in M . This is proved in our main technical result (cf. Lemma 17), whose proof critically exploits the recursive definition of the types. In a nutshell, we consider an ℓ -rotation sequence σ from T to T' , for some $\ell \leq k$, minimizing, among all ℓ -rotation sequences from T to T' , the number of used vertices in $V(T) \setminus M$. Our goal is to define another ℓ -rotation sequence σ' from T to T' using strictly less vertices in $V(T) \setminus M$ than σ , contradicting the choice of σ . To this end, let $v \in V(T) \setminus M$ be a lowest (with respect to the distance to $\text{root}(T)$) non-marked vertex of T that is used by σ . We distinguish two cases.

In Case 1, we assume that v has a marked T -sibling v' with $\tau(T, Z, v) = \tau(T, Z, v')$ (cf. Figure 8). It is not difficult to prove that we can define σ' from σ by just replacing v with v' in all the rotations of σ involving v (cf. Claim 18 and Claim 19).

In Case 2, all T -siblings v' of v with $\tau(T, Z, v) = \tau(T, Z, v')$, if any, are non-marked. In this case, in order to define another ℓ -rotation sequence σ' from T to T' that uses more marked vertices than σ , we need to modify σ in a more *global* way than in Case 1. Namely, in order to define σ' , we need a more global (and involved) replacement, which we achieve via what we call a *representative function* ρ . To define ρ , we first guarantee the existence of a very helpful vertex v^* that is a non-marked ancestor of v having a marked T -sibling v' of the same type such that no vertex in $T(v')$ is used by σ ; see Claim 20 and Figure 9.

Exploiting the recursive definition of type, we then define our representative function ρ , mapping vertices used by σ in $T(v^*)$ to vertices in $T(v')$ of the same type (cf. Claim 21), and prove that we can define σ' from σ by replacing each vertex v used by σ in $T(v^*)$ by its image via ρ in $T(v')$ in all the rotations of σ involving v (cf. Claim 22 and Claim 23).

3 Formal description of the FPT algorithm

In this section we present our FPT algorithm to solve the ROTATION DISTANCE problem. We start in Subsection 3.1 by providing some preliminaries and useful observations about what we call *good* and *bad* vertices. In Subsection 3.2 we show that we can assume that all the rotations involve vertices within balls of small radius around bad vertices. In Subsection 3.3 we describe our marking algorithm, using the definition of type, and show that the set of marked vertices can be computed in FPT time. In Subsection 3.4 we prove our main technical result (Lemma 17), stating that we can restrict the desired rotations to involve only marked vertices. In Subsection 4.5 of the full version we wrap up the previous results to prove Theorem 1. We stress that, as discussed in Section 2, once we have at hand the set of marked vertices M of bounded size, the algorithm consists just in a naive brute-force search over all possible rotation using only vertices in M , and this is why we omit the wrap up in this extended abstract.

3.1 Preliminaries and useful observations

We use standard notation about graphs and parameterized complexity, and we refer to the full version for additional preliminaries. Here we mostly focus on rooted trees and rotations.

For a positive integer p , we let $[p]$ denote the set $\{1, 2, \dots, p\}$. If $f : A \rightarrow B$ is a function between two sets A and B and $A' \subseteq A$, we denote by $f|_{A'}$ the restriction of f to A' .

Rooted trees. For a rooted tree T , we use $\text{root}(T)$ to denote its root. For a vertex $v \in V(T)$, we denote by $\text{parent}(T, v)$ the unique parent of v in T (or the empty set if v is the root), by $\text{children}(T, v)$ the set of children of v in T , by $\text{ancestors}(T, v)$ the set of ancestors of v in T (including v itself), and by $\text{desc}(T, v)$ the set of descendants of v in T (including v itself). The *strict* ancestors (resp. descendants) of v are the vertices in the set $\text{ancestors}(T, v) \setminus \{v\}$ (resp. $\text{desc}(T, v) \setminus \{v\}$). We denote by $T(v)$ the subtree of T rooted at v . Two vertices $v, v' \in V(T)$ are *T-siblings* if $\text{parent}(T, v) = \text{parent}(T, v')$.

Rotation of an edge in an elimination tree. We provide the formal definition of the rotation operation, which has been already informally defined in the introduction (cf. Figure 2).

► **Definition 2** (rotation operation). *Let T be an elimination tree of a graph G and let $uv \in E(T)$ with $\text{parent}(T, v) = u$. The rotation of uv in T creates another elimination tree $\text{rot}(T, uv)$ of G defined as follows, where for better readability we use $T' = \text{rot}(T, uv)$:*

1. $\text{parent}(T', u) = v$.
2. $u \in \text{children}(T', v)$.
3. If $u \neq \text{root}(T)$, let $z = \text{parent}(T, u)$. Then $\text{children}(T', z) = (\text{children}(T, z) \setminus \{u\}) \cup \{v\}$.
4. $\text{children}(T', u) \subseteq \text{children}(T, u)$.
5. Let $w \in \text{children}(T, v)$. If u is adjacent in G to some vertex in $T(w)$, then $w \in \text{children}(T', u)$; otherwise $w \in \text{children}(T', v)$.
6. For every vertex $s \in V(G) \setminus \{u, v, z\}$, $\text{children}(T', s) = \text{children}(T, s)$.

A k -rotation sequence from an elimination tree T to another elimination tree T' (of the same graph G) is an ordered set $(e_1, \dots, e_k)$ of edges such that, letting inductively $T_0 := T$ and, for $i \in [k]$, $T_i := \text{rot}(T_{i-1}, e_i)$ with $e_i \in E(T_{i-1})$, we have that $T_k = T'$. In other words, a k -rotation sequence consists of the ordered list of the k edges to be rotated in order to obtain T' from T , going through the intermediate elimination trees $T_1, \dots, T_{k-1}$ (of the same graph G). Clearly, $\text{dist}(T, T') \leq k$ if and only if there exists an ℓ -rotation sequence from T to T' for some $\ell \leq k$. We say that a vertex $v \in V(T)$ is *used* by a rotation sequence σ if it is an endpoint of some of the edges that are rotated by σ .

Throughout the paper, we assume that all the considered elimination trees are of a same fixed input graph G . For simplicity, we may assume henceforth that G is connected.

Good and bad vertices. Our algorithm exploits how a rotation in an elimination tree T may affect the parents and the children of its vertices. Note that a single rotation of an edge $uv \in E(T)$, yielding an elimination tree T' , may change the parent of arbitrarily many vertices. Indeed, these vertices are the roots of the red subtrees in Figure 2, and the considered vertex v may be adjacent to the root of arbitrarily many subtrees containing at least one vertex adjacent to u : for each such root r , $\text{parent}(T, r) = v$ but $\text{parent}(T', r) = u$. As a concrete example, in Figure 1, $\text{parent}(T, s) = v$ but $\text{parent}(T', s) = u$. On the other hand, item 6 of Definition 2 implies that there are at most three vertices whose children set changes from T to T' , namely u, v, z as depicted in Figure 2. (Note that the sets of children of u and v always change, and that of z changes provided that this vertex exists.) We state this observation formally, since it will be extensively used afterwards.

► **Observation 3.** *One rotation may change the set of children of at most three vertices.*

The above discussion motivates the following definition.

► **Definition 4 (bad vertices).** *Given two elimination trees T and T' , a vertex $v \in V(T)$ is (T, T') -children-bad (resp. (T, T') -parent-bad) if $\text{children}(T, v) \neq \text{children}(T', v)$ (resp. $\text{parent}(T, v) \neq \text{parent}(T', v)$). A vertex $v \in V(T)$ is (T, T') -bad if it is (T, T') -children-bad, or (T, T') -parent-bad, or both. A vertex $v \in V(T)$ is (T, T') -good if it is not (T, T') -bad.*

Note that T contains no (T, T') -children-bad (or (T, T') -parent-bad, or just (T, T') -bad) vertices if and only if $T = T'$, that is, if and only if $\text{dist}(T, T') = 0$. Also, note that a vertex $v \in V(T)$ is (T, T') -children-bad, with $\text{children}(T, v) \neq \emptyset$, if and only if at least one of its children is (T, T') -parent-bad. Observation 3 directly implies the following necessary condition for the existence of a solution.

► **Observation 5.** *Given two elimination trees T and T' , if $\text{dist}(T, T') \leq k$ then the number of (T, T') -children-bad vertices is at most $3k$.*

Observation 5 is equivalent to saying that we can safely conclude that any instance (G, T, T', k) of ROTATION DISTANCE with at least $3k + 1$ (T, T') -children-bad vertices is a no-instance. Thus, we can assume henceforth that we are dealing with an instance of ROTATION DISTANCE containing at most $3k$ (T, T') -children-bad vertices.

3.2 Restricting the rotations to small balls around bad vertices

Our next goal is to prove (Lemma 7) that we can assume that the desired sequence of at most k rotations to transform T into T' uses only edges whose both endvertices lie in the union of all the balls of appropriate radius (depending only on k) around (T, T') -children-bad vertices of T , whose number is bounded by a function of k by Observation 5.

In the next definition, for the sake of notational simplicity we omit T, T' , and k from the notation B_{cb} , as we assume that they are already given, and fixed, as the input of our problem. We include $\text{root}(T)$ in the considered set for technical reasons, namely in the proof of Claim 9.

► **Definition 6** (union of balls of children-bad vertices). *Let $C \subseteq V(T)$ be the set of (T, T') -children-bad vertices. We define $B_{cb} = N_T^{2k}[C \cup \text{root}(T)]$.*

► **Lemma 7.** *If $\text{dist}(T, T') \leq k$, then there exists an ℓ -rotation sequence from T to T' , with $\ell \leq k$, using only vertices in B_{cb} .*

Proof. Let σ be an ℓ -rotation sequence from T to T' , with $\ell \leq k$. Let us denote by $\text{out}(\sigma)$ the number of edges in σ with at least one endvertex not in B_{cb} . Assuming that $\text{out}(\sigma) \geq 1$, we proceed to construct another ℓ' -rotation sequence σ' from T to T' , with $\ell' \leq \ell$, such that $\text{out}(\sigma') < \text{out}(\sigma)$. Repeating this procedure eventually yields a sequence as claimed in the statement of the lemma.

For $i \in [\ell]$, let $u_i v_i$ be the i -th edge of σ and let T_i be the elimination tree obtained after the rotation of $u_i v_i$. Let also $T_0 = T$. For $i \in [\ell]$, we say that a vertex $w \in V(T)$ is *affected* by the rotation of $u_i v_i$ if $\text{children}(T_{i-1}, w) \neq \text{children}(T_i, w)$, and it is σ -*affected* if it is affected by the rotation of some edge in σ . Recall that a rotation affects at most three vertices, and that these vertices are within distance at most two in the original tree (cf. vertices u, v, z in Figure 2). Moreover, any rotation may increase or decrease vertex distances by at most one, that is, for any $i \in [\ell]$ and any two vertices $x, y \in V(T)$, it holds that

$$|\text{dist}_{T_{i-1}}(x, y) - \text{dist}_{T_i}(x, y)| \leq 1. \quad (1)$$

Let $A \subseteq V(T)$ be the set of σ -affected vertices, and note that $|A| \leq 3k$. The observation above about the fact that the (two or three) vertices affected by a rotation are within distance at most two (in the tree where the rotation is done), together with Equation 1, imply that for every $Z \in \text{cc}(T[A])$, it holds that its diameter is linearly bounded by k , namely

$$\text{diam}(Z) \leq 2k. \quad (2)$$

Since by assumption $\text{out}(\sigma) \geq 1$, there exists $j \in [\ell]$ such that $u_j \notin B_{cb}$ or $v_j \notin B_{cb}$ (or both); assume without loss of generality that $u_j \notin B_{cb}$. Let $Z^* \in \text{cc}(T[A])$ be the connected component of $T[A]$ containing vertices u_j and v_j (note that they indeed lie in the same component of $T[A]$ since edge $u_j v_j$ belongs to σ). Let $C \subseteq V(T)$ be the set of (T, T') -children-bad vertices, and recall that $B_{cb} = N_T^{2k}[C \cup \text{root}(T)]$. See Figure 3.

► **Claim 8.** No vertex in Z^* is (T, T') -children-bad.


Proof. Suppose towards a contradiction that the statement does not hold, and let $x \in Z^* \cap C$. Since $u_j \notin B_{cb}$ (see Figure 3), the definition of B_{cb} implies that $\text{dist}_T(x, u_j) \geq 2k + 1$, contradicting Equation 2 because both x and u_j belong to Z^* . ◀

► **Claim 9** (★). All vertices in Z^* are (T, T') -good.

Relying on Claim 9, we define from σ an ℓ' -rotation sequence σ' from T to T' , with $\ell' \leq \ell$, as follows: σ' consists of the (ordered) edges appearing in σ , except from those with both endvertices lying in the connected component Z^* of $T[A]$.

► **Claim 10** (★). σ' is an ℓ' -rotation sequence from T to T' with $\ell' \leq \ell$ and $\text{out}(\sigma') < \text{out}(\sigma)$.

The above claim concludes the proof of the lemma. ◀



■ **Figure 3** Illustration of the proof of Lemma 7. (T, T') -children-bad vertices are depicted in red, and the balls of radius $2k$ around them are shown with orange bubbles. The connected component $Z^* \in \text{cc}(T[A])$ containing both vertices u_j and v_j is depicted with thick blue edges. Distances in the figure are not meant to be accurate, and an extremity of an edge without a vertex means that T continues in that direction.

By Lemma 7, we focus henceforth on trying to find an ℓ -rotation sequence from T to T' , with $\ell \leq k$, consisting only of edges with both endvertices in B_{cb} . First, we will consider each of the at most $3k + 1$ connected components of $T[B_{cb}]$ separately. In fact, we can get a better bound, as if $T[B_{cb}]$ has at least $k + 1$ connected components, we can immediately conclude that we are dealing with a no-instance, since at least one rotation is needed in each component. Thus, we may assume that $T[B_{cb}]$ has at most k connected components. On the other hand, since $T[B_{cb}]$ is defined as the union of at most $3k + 1$ balls of radius $2k$, it follows that every $Z \in \text{cc}(T[B_{cb}])$ satisfies


$$\text{diam}(Z) \leq (3k + 1)4k. \quad (3)$$

Thus, by Equation 3, the “only” obstacle to obtain the desired FPT algorithm is that the vertices in B_{cb} can have an arbitrarily large degree. Note that in the particular case where the underlying graph G has bounded degree, for instance if G is a path [14, 20, 21, 23], the maximum degree of any elimination tree of G is bounded, and therefore in that case $|B_{cb}|$ is bounded by a function of k , and an FPT algorithm follows immediately. To the best of our knowledge, this result was not known for graphs other than paths (albeit, with a better running time than the one that results from just brute-forcing on the set B_{cb} , which is of the form $2^{2^{O(k)}} \cdot |V(G)|$).

3.3 Description of the marking algorithm

Let henceforth Z be a connected component of $T[B_{cb}]$, which we consider as a rooted tree with its own set of leaves, which are not necessarily leaves in T . We define $\text{root}(Z)$ to be the vertex in $V(Z)$ closest to $\text{root}(T)$ in T . Before defining the types, we need to define the *trace* of a vertex v in a designated vertex set $Z \subseteq V(T)$.

► **Definition 11** (trace of a vertex in a component Z). *Let T be an elimination tree (of a graph G), let Z be a rooted subtree of T corresponding to a connected component of B_{cb} , and let $v \in V(Z)$. The trace of v in Z , denoted by $\text{trace}(T, Z, v)$, is a binary vector of dimension $\text{dist}_T(v, \text{root}(Z))$ defined as follows (note that if $v = \text{root}(Z)$, then its trace is empty). For $i \in [\text{dist}_T(v, \text{root}(Z))]$, let $u_i \in \text{ancestors}(T, v)$ be the ancestor of v in T such that $\text{dist}_T(v, u_i) = i$. Then the i -th coordinate of $\text{trace}(T, Z, v)$ is 1 if $wu_i \in E(G)$ for some vertex $w \in V(T(v))$, and 0 otherwise.*



■ **Figure 4** A component Z of $T[B_{\text{cb}}]$ and the trace of some of its vertices $v_1, \dots, v_7$. Red dotted edges represent adjacencies in G . Note the $\text{trace}(T, Z, v_3) = \text{trace}(T, Z, v_6)$, even if v_3 and v_6 are not siblings in Z .

See Figure 4 for an example of the trace of some vertices in a component Z .


For a vertex $v \in V(T)$, let $\text{want-parent}(T, T', v)$ be equal to $\emptyset$ if $\text{parent}(T, v) = \text{parent}(T', v)$, and to $\text{parent}(T', v)$ otherwise. Note that, by Observation 5, the function $\text{want-parent}(T, T', v)$ can take up to $3k + 1$ distinct values when ranging over all $v \in V(T)$.

► **Definition 12** (type of a vertex in a component Z). *Let T be an elimination tree (of a graph G), let Z be a rooted subtree of T corresponding to a connected component of B_{cb} , and let $v \in V(Z)$. The type of vertex v , denoted by $\tau(T, Z, v)$, is recursively defined as follows, where $\text{type-children}(T, Z, v) := \{\tau(T, Z, u) \mid u \in \text{children}(Z, v)\}$ is the set of types occurring in the children of v :*

- If v is a leaf of Z , then $\tau(T, Z, v)$ consists of the pair $(\text{want-parent}(T, T', v), \text{trace}(T, Z, v))$.
- Otherwise, $\tau(T, Z, v)$ consists of a tuple $(\text{want-parent}(T, T', v), \text{trace}(T, Z, v), f_v)$, where $f_v : \text{type-children}(T, Z, v) \rightarrow [k + 1]$ is a mapping defined such that, for every $\tau \in \text{type-children}(T, Z, v)$,

$$f_v(\tau) = \min\{k + 1, |\{u \in \text{children}(Z, v) \mid \tau(T, Z, u) = \tau\}|\}. \quad (4)$$

See Figure 5 for an example for $k = 2$ of how the types are computed in a component Z .



■ **Figure 5** A component Z of $T[B_{\text{cb}}]$ and the types of its vertices, for an instance with $k = 2$. For the sake of simplicity, different types are depicted with different numbers. Assume that the leaves have only two possible types, namely 1 and 2, and that all non-leaf vertices at the same distance from the root have the same trace and the same function $\text{want-parent}(T, T', \cdot)$. Note that the red vertices have the same type (namely, 4) because they both have one child of type 1 and at least $k + 1 = 3$ children of type 2. Note also that the blue vertices have the same type (namely, 6) because they both have one child of type 3 and one child of type 4.

► **Lemma 13.** *The set $\{\tau(T, Z, v) \mid v \in V(Z)\}$ has size bounded by a function $g(k)$, with*

$$g(k) = k^{2^{2^{\dots 2^{\mathcal{O}(k^2)}}}}, \text{ where the tower of exponentials has height } \text{diam}(Z) = \mathcal{O}(k^2). \quad (5)$$

Proof. Let $d = \text{diam}(Z)$, and note that $d \leq (3k+1)4k = \mathcal{O}(k^2)$ by Equation 3. For $i \in [d]$, let τ_i be the number of distinct types among the vertices in $V(Z)$ that are at distance exactly i from $\text{root}(Z)$ in T . Formally,

$$\tau_i = |\{\tau(T, Z, v) \mid \text{dist}_T(v, \text{root}(Z)) = i\}|. \quad (6)$$

By the definition of type, $\tau_d \leq 2^d \cdot (3k+1)$ since, on the one hand, all vertices at distance d from $\text{root}(Z)$ are leaves of Z , and the number of possible traces among leaves is at most 2^d , and on the other hand the term $3k+1$ corresponds to the possible distinct values of the function $\text{want-parent}(T, T', v)$. For every $i \in [d-1]$, Equation 4 implies that

$$\tau_i \leq (3k+1) \cdot 2^i \cdot \sum_{j=i+1}^d (k+2)^{\tau_j}, \quad (7)$$

where the term $3k+1$ again comes from the possible distinct values of the function $\text{want-parent}(T, T', v)$, the term 2^i comes from the possible different traces within distance i from $\text{root}(Z)$, and the term $k+2$ follows from the fact that, for every $v \in V(Z)$, the function f_v can take up to $k+1$ values for each type τ of a children of v , together with the possibility that a type is not present among the children of v . Note that a vertex $v \in V(T)$ with $\text{dist}(v, \text{root}(Z)) = i$ may have children being roots of any possible subtree with diameter at most $d-i$, justifying the sum in Equation 7. Clearly, the upper bound of Equation 7 is maximized for $i = 1$, that is, for the children of $\text{root}(Z)$, yielding the bound claimed in Equation 5. ◀

Note that, in order to compute the type of a vertex in a component Z , the recursive definition of the types together with Lemma 13 easily imply the following observation, where the term $|V(G)|$ comes from checking the neighborhood of $T(v)$ within the set $\text{ancestors}(T, v)$ in the computation of the trace (cf. Definition 11).

► **Observation 14.** *Let T be an elimination tree of a graph G , let Z be a rooted subtree of T corresponding to a connected component of B_{cb} , and let $v \in V(Z)$. Then $\tau(T, Z, v)$ can be computed in time $g(k) \cdot |V(G)|$, where $g(k)$ is the function from Lemma 13.*


We will now use the notion of type and the bound given by Lemma 13 to define the desired set $M_Z \subseteq Z$ of size bounded by a function of k . In order to find M_Z , we apply a marking algorithm on Z , that first identifies a set $M_Z^{\text{pre}} \subseteq V(Z)$ of *pre-marked* vertices, whose size is not necessarily bounded by a function of k , and then “prunes” this set M_Z^{pre} in a root-to-leaf fashion to find the desired set of *marked* vertices $M_Z \subseteq M_Z^{\text{pre}}$ of appropriate size.

Start with $M_Z^{\text{pre}} = \emptyset$. For every vertex $v \in V(Z)$ and every $\tau \in \text{type-children}(Z, v)$, do the following:

- If $|\{u \in \text{children}(Z, v) \mid \tau(Z, u) = \tau\}| \leq k+1$, add the whole set $\{u \in \text{children}(Z, v) \mid \tau(Z, u) = \tau\}$ to M_Z^{pre} .
- Otherwise, add to M_Z^{pre} an arbitrarily chosen subset of $\{u \in \text{children}(Z, v) \mid \tau(Z, u) = \tau\}$ of size $k+1$.

Finally, add $\text{root}(Z)$ to M_Z^{pre} . We define $M_{\text{pre}} = \cup_{Z \in \text{cc}(T[B_{\text{cb}}])} M_Z^{\text{pre}}$ and we call it the set of *pre-marked vertices* of T .

We are now ready to define our bounded-size set $M_Z \subseteq M_Z^{\text{pre}}$. Start with $M_Z = \{\text{root}(Z)\}$ and for $i = 0, \dots, \text{diam}(Z) - 1$, proceed inductively as follows: if $v \in V(Z)$ is a vertex with $\text{dist}_Z(v, \text{root}(Z)) = i$ that already belongs to M_Z , add to M_Z the set $\text{children}(Z, v) \cap M_Z^{\text{pre}}$. Finally, for every (T, T') -children-bad vertex v of T that belongs to Z , we add to M_Z the set $\text{ancestors}(Z, v)$. This concludes the construction of M_Z . Note that if a vertex $v \in V(Z)$ belongs to M_Z , then the whole set $\text{ancestors}(Z, v)$ belongs to M_Z as well. We define $M = \bigcup_{Z \in \text{cc}(T[B_{\text{cb}}])} M_Z$, and we call it the set of *marked vertices* of T . See Figure 6 for an example of the marking algorithm.



■ **Figure 6** Example of the marking algorithm applied to a component Z of $T[B_{\text{cb}}]$, for an instance with $k = 1$. As in Figure 5, different types are depicted with different numbers. Vertices inside blue squares belong to M_Z^{pre} , and red vertices belong to M_Z .

► **Lemma 15** (*). *The set $M \subseteq V(T)$ of marked vertices has size bounded by a function $h(k)$, where $h(k)$ has the same asymptotic growth as the function $g(k)$ given by Lemma 13. Moreover, M can be computed in time $h(k) \cdot |V(G)|$.*

3.4 Restricting the rotations to marked vertices


We now prove our main technical result (Lemma 17), which immediately yields the desired FPT algorithm combined with Lemma 15 (whose proof uses Lemma 7), as discussed in Section 2. We first need an easy lemma that will be extensively used in the proof of Lemma 17; see Figure 7.

► **Lemma 16.** *Let σ be an ℓ -rotation sequence from T to T' , for some $\ell \leq k$. For every vertex $v \in V(T)$, there are at most k vertices $u_1, \dots, u_k \in \text{children}(T, v)$ such that σ uses a vertex in each of the rooted subtrees $T(u_1), \dots, T(u_k)$.*

Proof. Assume towards a contradiction that there is a vertex $v \in V(T)$ having $k+1$ children, say $u_1, \dots, u_{k+1}$, such that σ uses $k+1$ vertices $u'_1, \dots, u'_{k+1}$ with $u'_r \in T(u_r)$ for $r \in [k+1]$. Then, since $\sigma = (e_1, \dots, e_\ell)$ is made of at most k rotations, by the pigeonhole principle necessarily there exist two children u_i, u_j of v and an integer $p \in [\ell]$ such that $e_p = u'_i u'_j$, and none of u'_i, u'_j occurs in any other rotation of σ other than e_p . Since $e_p = u'_i u'_j$, it means that in the elimination tree where this rotation takes place, namely T_{p-1} , it holds that $u'_i u'_j \in E(T_{p-1})$. See Figure 7 for an illustration.

On the other hand, note that if T_2 is an elimination tree resulting from an elimination tree T_1 after the rotation of an edge $wz \in E(T_1)$, and a, b are vertices of T_1 (and T_2) such that $ab \notin E(T_1)$ and $ab \in E(T_2)$, then necessarily $\{a, b\} \cap \{w, z\} \neq \emptyset$, that is, necessarily the rotation involves at least one of a and b . See Figure 2 for a visualization of this claim, where the new edges that appear after the rotation of wz are zv and the edges between u and some of the red subtrees: each of these new edges contains u or v .

Since $u'_i u'_j \notin E(T_0) = E(T)$ because $u'_i \in T(u_i), u'_j \in T(u_j)$, and u_i, u_j are T -siblings, and $u'_i u'_j \in E(T_{p-1})$, by the above paragraph there exists some integer $q \in [\ell]$, with $q < p$, such that the rotation e_q of σ contains at least one of u'_i and u'_j . This contradicts that fact that none of u'_i, u'_j occurs in any other rotation of σ other than e_p . ◀



■ **Figure 7** Illustration of the statement and the proof of Lemma 16.

We are now ready to prove our main lemma.

► **Lemma 17.** *If $\text{dist}(T, T') \leq k$, then there exists an ℓ -rotation sequence from T to T' , with $\ell \leq k$, using only vertices in M .*


Proof. Let σ be an ℓ -rotation sequence from T to T' , for some $\ell \leq k$, minimizing, among all ℓ -rotation sequences from T to T' , the number of vertices in $V(T) \setminus M$ (that is, the non-marked vertices) used by σ . Note that σ exists by the hypothesis that $\text{dist}(T, T') \leq k$. If there are no vertices in $V(T) \setminus M$ used by σ , then we are done, so assume that there are. Our goal is to define another ℓ -rotation sequence σ' from T to T' using strictly less vertices in $V(T) \setminus M$ than σ , contradicting the choice of σ and concluding the proof.

To this end, let $v \in V(T) \setminus M$ be a furthest (with respect to the distance to $\text{root}(T)$) non-marked vertex of T that is used by σ . By Lemma 7, we can assume that $v \in B_{\text{cb}}$. Let Z be the connected component of $T[B_{\text{cb}}]$ such that $v \in Z$. For technical reasons, it will be helpful to assume that v is not a leaf of Z . (This can be achieved, for instance, by observing that the analysis of the size of the components Z in Lemma 7 is not tight. Alternatively, we can just “artificially” increase their diameter by one –i.e., replacing $2k$ with $2k + 1$ in the definition of B_{cb} – so that we can safely assume that the leaves of Z are never used by a rotation sequence.) Note that the choice of v as a lowest (i.e., furthest) non-marked vertex of T used by σ implies that for every vertex $u \in \text{children}(T, v)$, the whole subtree $T(u)$ remains intact throughout σ , meaning that it appears as a rooted subtree in all the intermediate elimination trees generated by the ℓ -rotation sequence σ . We distinguish two cases, the second one being considerably more involved, but that will benefit from the intuition developed in the first one.

Case 1: v has a marked T -sibling v' with $\tau(T, Z, v) = \tau(T, Z, v')$.

Since v is non-marked and by assumption it has some marked T -sibling, the definition of M (namely, that up to $k + 1$ vertices of each type are recursively marked) and Lemma 16 imply that v has some marked T -sibling of the same type that is *not* used by σ . Let without loss of generality v' be such a T -sibling of v . Note that the choice of v as a lowest non-marked vertex used by σ implies that the whole subtree $T(v')$ remains intact throughout σ . See Figure 8 for an illustration.

In this case, we define σ' from σ by just replacing v with v' in all the rotations of σ involving v . We need to prove that σ' is well-defined (that is, that the edges to be rotated do exist in the intermediate elimination trees) and that it is an ℓ -rotation sequence from T to T' . Once this is proved, this case is done, as σ' uses strictly more marked vertices than σ .



■ **Figure 8** Illustration of Case 1 in the proof of Lemma 17. Vertex v is non-marked and used by σ , and its T -sibling v' is marked (in red) and not used by σ .

▷ **Claim 18** (★). σ' is a well-defined ℓ -rotation sequence.

▷ **Claim 19** (★). σ' is an ℓ -rotation sequence from T to T' .

Case 2: all T -siblings v' of v with $\tau(T, Z, v) = \tau(T, Z, v')$, if any, are non-marked.

In this case, in order to define another ℓ -rotation sequence σ' from T to T' that uses more marked vertices than σ , we need to modify σ in a more global way than what we did in Case 1, where it was enough to replace vertex v with a marked T -sibling of the same type. To this end, the following claim guarantees the existence of a very helpful vertex v^* . See Figure 9 for an illustration.


▷ **Claim 20.** There exists a unique vertex $v^* \in \text{ancestors}(Z, v)$ such that

- v^* is non-marked,
- v^* has a marked T -sibling v' such that
 - $\tau(T, Z, v^*) = \tau(T, Z, v')$, and
 - no vertex in $T(v')$ is used by σ , and
- v^* is the vertex closest to v satisfying the above properties.

Proof. Since $\text{root}(Z) \in M$ and $v \notin M$, the definition of the marking algorithm implies that there exists a marked T -sibling v' of v^* with $\tau(T, Z, v^*) = \tau(T, Z, v')$. Moreover, the fact that M is defined by recursively marking up to $k + 1$ vertices of each type implies, together with Lemma 16 and the fact that the desired vertex v^* is non-marked, that there is some T -sibling v' of v^* with $\tau(T, Z, v^*) = \tau(T, Z, v')$ such that no vertex in $T(v')$ is used by σ . Finally, we can clearly choose v^* in a unique way as being the vertex closest to v satisfying the above properties. $\triangleleft$

Note that Case 1 of the proof corresponds to the particular case where v^* is equal to v itself, but we prefer to separate both cases for the sake of readability. Intuitively, we will apply recursively the argument of Case 1 to the rooted subtrees $T(v^*)$ and $T(v')$, starting with v^* and v' , exploiting the definition of types to appropriately define the desired replacement of vertices to construct σ' from σ .

Formally, let U be the set of vertices in $V(T(v^*)) \cap Z$ used by σ (so in Case 1, $U = \{v\}$), and let $\rho : U \rightarrow V(T(v')) \cap Z$ be the injective function defined as follows. For $i = 0, \dots, \text{dist}_T(v^*, v)$, let $U_i \subseteq U$ be the set of vertices in $V(T(v^*)) \cap Z$ used by σ that are at distance exactly



■ **Figure 9** Illustration of Case 2 in the proof of Lemma 17. All vertices in $T(v^*)$ are non-marked, and (at least) vertex v is used by σ . No vertex in $T(v')$ is used by σ , and (at least) vertex v' is marked (in red). The sets $U_{i,\tau}$ of vertices used by σ in $T(v^*) \cap Z$ are depicted with squares, as well as their images in $T(v') \cap Z$ via the representative function ρ .

i from vertex v^* in T . Note that some of the sets U_i may be empty, and that $U_{\text{dist}_T(v^*,v)}$ contains v . For every type τ occurring in a vertex in U_i , let $U_{i,\tau}$ be the set of vertices of type τ in U_i . Note that, if a set U_i is non-empty, then $\{U_{i,\tau} \mid \tau \text{ occurs in } U_i\}$ defines a partition of U_i into non-empty sets. Let $U'_{i,\tau}$ be a set of marked vertices of type τ in $V(T(v')) \cap Z$ of size $|U_{i,\tau}|$ (we shall prove in Claim 21 that it exists). Then we define $\rho|_{U_{i,\tau}}$ as any bijection between $U_{i,\tau}$ and $U'_{i,\tau}$. See Figure 9 for an illustration.

▷ **Claim 21** (★). The function ρ is well-defined and injective.

For every vertex $u \in U$ (recall that U is the set of vertices in $V(T(v^*)) \cap Z$ used by σ), the vertex $\rho(u)$ is called the *representative* of u . Note that the function ρ is also defined on v , since it is used by σ . We now define σ' from σ by replacing, in the rotations defining the sequence, every vertex $u \in U$ by its representative $\rho(u)$.

▷ **Claim 22** (★). σ' is a well-defined ℓ -rotation sequence.

▷ **Claim 23** (★). σ' is an ℓ -rotation sequence from T to T' .

By Claim 23, σ' is an ℓ -rotation sequence from T to T' , and it uses strictly more marked vertices than σ , because no vertex of $T(v^*)$ is marked (by the conditions in Claim 20), and within $T(v')$ there is at least one marked vertex, namely v' . This concludes Case 2.

In both cases, we have defined from σ another ℓ -rotation sequence σ' from T to T' using strictly less vertices in $V(T) \setminus M$ than σ , contradicting the choice of σ . ◀

■ **Table 1** Known results and open problems about the (parameterized) complexity of the ROTATION DISTANCE problem, both on paths and general graphs.

ROTATION DISTANCE	paths	general graphs
NP-hard	open	✓ [19]
FPT	✓ [14, 20, 21, 23, 24]	✓ [Theorem 1]
Polynomial kernel	✓ [14, 24]	open

4 Further research

We proved that the ROTATION DISTANCE problem, for a general graph G , can be solved in time $f(k) \cdot |V(G)|$, where $f(k)$ is the function given by Theorem 1. This function is quite large, and it is worth trying to improve it. The growth of $f(k)$ is mainly driven by the number of different types of vertices (cf. Definition 12) that we consider in our marking algorithm. We need this recursive definition of type to guarantee that, when two vertices v, v' have the same type, then for each possible type τ and every integer d at most the bound given in Equation 3, vertices v and v' have the same number (up to $k + 1$) of descendants of type τ within distance d . This is exploited, for instance, in Case 2 of the proof of Lemma 17 to apply a recursive argument. It may be possible to find a simpler argument in the replacement operation performed in the proof of Lemma 17 (using the representative function ρ), and in that case, one may allow for a less refined notion of type, leading to a better bound.

Another natural direction is to investigate whether ROTATION DISTANCE admits a polynomial kernel parameterized by k . So far, this is only known when the considered graph G is a path, where even linear kernels are known [14, 24]; see Table 1. As an intermediate step, one may consider graphs of bounded degree, for which it seems plausible that Lemma 7 (restriction to few balls of bounded diameter) provides a helpful opening step.

Finally, Ito et al. [19] also proved the NP-hardness of a related problem called COMBINATORIAL SHORTEST PATH ON POLYMATROIDS, relying on the fact that graph associahedra can be realized as the base polytopes of some polymatroids [28]. To the best of our knowledge, the parameterized complexity of this problem has not been investigated.

References

- 1 Georgy M. Adel'son-Vel'skii and Evgenii M. Landis. An algorithm for the organization of information. *Soviet Mathematics Doklady*, 3:1259–1263, 1962. URL: <https://zhjwpku.com/assets/pdf/AED2-10-av1-paper.pdf>.
- 2 Oswin Aichholzer, Wolfgang Mulzer, and Alexander Pilz. Flip Distance Between Triangulations of a Simple Polygon is NP-Complete. *Discrete & Computational Geometry*, 54(2):368–389, 2015. doi:10.1007/S00454-015-9709-7.
- 3 Benjamin Aram Berendsohn. The diameter of caterpillar associahedra. In *Proc. of the 18th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT)*, volume 227 of *LIPIcs*, pages 14:1–14:12, 2022. doi:10.4230/LIPICS.SWAT.2022.14.
- 4 Nicolas Bousquet, Amer E. Mouawad, Naomi Nishimura, and Sebastian Siebertz. A survey on the parameterized complexity of reconfiguration problems. *Computer Science Review*, 53:100663, 2024. doi:10.1016/J.COSREV.2024.100663.
- 5 Jean Cardinal, Linda Kleist, Boris Klemz, Anna Lubiw, Torsten Mütze, Alexander Neuhaus, and Lionel Pournin. Working group 4.2: Rotation distance between elimination trees. *Report from Dagstuhl Seminar 22062: Computation and Reconfiguration in Low-*

- Dimensional Topological Spaces*, page 35, 2022. URL: <https://drops.dagstuhl.de/storage/04dagstuhl-reports/volume12/issue02/22062/DagRep.12.2.17/DagRep.12.2.17.pdf>.
- 6 Jean Cardinal, Stefan Langerman, and Pablo Pérez-Lantero. On the diameter of tree associahedra. *Electronic Journal of Combinatorics*, 25(4):4, 2018. doi:10.37236/7762.
 - 7 Jean Cardinal, Arturo Merino, and Torsten Mütze. Efficient generation of elimination trees and graph associahedra. In *Proc. of the 33rd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2128–2140, 2022. doi:10.1137/1.9781611977073.84.
 - 8 Jean Cardinal, Lionel Pournin, and Mario Valencia-Pabon. Diameter estimates for graph associahedra. *Annals of Combinatorics*, 26:873–902, 2022. doi:10.1007/s00026-022-00598-z.
 - 9 Jean Cardinal, Lionel Pournin, and Mario Valencia-Pabon. The rotation distance of brooms. *European Journal of Combinatorics*, 118:103877, 2024. doi:10.1016/J.EJC.2023.103877.
 - 10 Jean Cardinal and Raphael Steiner. Shortest paths on polymatroids and hypergraphic polytopes. *CoRR*, abs/2311.00779, 2023. doi:10.48550/arXiv.2311.00779.
 - 11 Michael Carr and Satyan L. Devadoss. Coxeter complexes and graph-associahedra. *Topology and its Applications*, 153(12):2155–2168, 2006. doi:10.1016/j.topol.2005.08.010.
 - 12 Cesar Ceballos, Francisco Santos, and Günter M. Ziegler. Many non-equivalent realizations of the associahedron. *Combinatorica*, 35(5):513–551, 2015. doi:10.1007/s00493-014-2959-9.
 - 13 Sean Cleary. Restricted rotation distance between k -ary trees. *Journal of Graph Algorithms and Applications*, 27(1):19–33, 2023. doi:10.7155/JGAA.00611.
 - 14 Sean Cleary and Katherine St. John. Rotation distance is fixed-parameter tractable. *Information Processing Letters*, 109(16):918–922, 2009. doi:10.1016/J.IPL.2009.04.023.
 - 15 Sean Cleary and Katherine St. John. A linear-time approximation algorithm for rotation distance. *Journal of Graph Algorithms and Applications*, 14(2):385–390, 2010. doi:10.7155/JGAA.00212.
 - 16 Satyan L. Devadoss. A realization of graph associahedra. *Discrete Mathematics*, 309(1):271–276, 2009. doi:10.1016/J.DISC.2007.12.092.
 - 17 Leonidas J. Guibas and Robert Sedgwick. A dichromatic framework for balanced trees. In *Proc. of the 19th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 8–21, 1978. doi:10.1109/SFCS.1978.3.
 - 18 Christophe Hohlweg and Carsten E. M. C. Lange. Realizations of the associahedron and cyclohedron. *Discrete & Computational Geometry*, 37(4):517–543, 2007. doi:10.1007/S00454-007-1319-6.
 - 19 Takehiro Ito, Naonori Kakimura, Naoyuki Kamiyama, Yusuke Kobayashi, Shun-ichi Maezawa, Yuta Nozaki, and Yoshio Okamoto. Hardness of Finding Combinatorial Shortest Paths on Graph Associahedra. In *Prof. of the 50th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 261 of *LIPICs*, pages 82:1–82:17, 2023. doi:10.4230/LIPICs.ICALP.2023.82.
 - 20 Iyad Kanj, Eric Sedgwick, and Ge Xia. Computing the flip distance between triangulations. *Discrete & Computational Geometry*, 58(2):313–344, 2017. doi:10.1007/S00454-017-9867-X.
 - 21 Haohong Li and Ge Xia. An $O(3.82^k)$ time FPT algorithm for convex flip distance. In *Proc. of the 40th International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 254 of *LIPICs*, pages 44:1–44:14, 2023. doi:10.4230/LIPICs.STACS.2023.44.
 - 22 Jean-Louis Loday. Realization of the Stasheff polytope. *Archiv der Mathematik*, 83:267–278, 2004. doi:10.1007/s00013-004-1026-y.
 - 23 Anna Lubiw and Vinayak Pathak. Flip distance between two triangulations of a point set is NP-complete. *Computational Geometry*, 49:17–23, 2015. doi:10.1016/J.COMGEO.2014.11.001.
 - 24 Joan M. Lucas. An improved kernel size for rotation distance in binary trees. *Information Processing Letters*, 110(12-13):481–484, 2010. doi:10.1016/J.IPL.2010.04.022.
 - 25 Thibault Manneville and Vincent Pilaud. Graph properties of graph associahedra. *CoRR*, abs/1409.8114, 2010. arXiv:1409.8114.

- 26 Jaroslav Nesetril and Patrice Ossona de Mendez. *Sparsity - Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and combinatorics*. Springer, 2012. doi:10.1007/978-3-642-27875-4.
- 27 Alexander Pilz. Flip distance between triangulations of a planar point set is APX-hard. *Computational Geometry*, 47(5):589–604, 2014. doi:10.1016/J.COMGEO.2014.01.001.
- 28 Alex Postnikov, Victor Reiner, and Lauren Williams. Faces of generalized permutohedra. *Documenta Mathematica*, 13:207–273, 2008. doi:10.4171/DM/248.
- 29 Alexander Postnikov. Permutohedra, associahedra, and beyond. *International Mathematics Research Notices*, 2009(6):1026–1106, 2009. doi:10.1093/imrn/rnn153.
- 30 Lionel Pournin. The diameter of associahedra. *Advances in Mathematics*, 259:13–42, 2014. doi:10.1016/j.aim.2014.02.035.
- 31 Lionel Pournin. The asymptotic diameter of cyclohedra. *Israel Journal of Mathematics*, 219(2):609–635, 2017. doi:10.1007/s11856-017-1492-0.
- 32 Daniel D. Sleator, Robert E. Tarjan, and William P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. *Journal of the American Mathematical Society*, 1(3):647–681, 1988. doi:10.1090/S0894-0347-1988-0928904-4.
- 33 Daniel Dominic Sleator and Robert Endre Tarjan. Self-adjusting binary search trees. *Journal of the ACM*, 32(3):652–686, 1985. doi:10.1145/3828.3835.
- 34 Richard P. Stanley. *Catalan numbers*. Cambridge University Press, 2015. doi:10.1017/CB09781139871495.