

Reconfiguration of Spanning Trees with Degree Constraint or Diameter Constraint *

Nicolas Bousquet ^{†1}, Takehiro Ito ^{‡2}, Yusuke Kobayashi ^{§3}, Haruka Mizuta ², Paul Ouvrard ^{¶4}, Akira Suzuki ², and Kunihiro Wasa ⁵

¹CNRS, LIRIS, Université de Lyon, Université Claude Bernard Lyon 1, Lyon, France

²Graduate School of Information Sciences, Tohoku University, Japan

³Research Institute for Mathematical Sciences, Kyoto University, Japan

⁴Univ. Bordeaux, Bordeaux INP, CNRS, LaBRI, UMR5800, F-33400 Talence, France ^{||}

⁵Toyohashi University of Technology, Japan ^{**}

January 26, 2022

Abstract

We investigate the complexity of finding a transformation from a given spanning tree in a graph to another given spanning tree in the same graph via a sequence of edge flips. The exchange property of the matroid bases immediately yields that such a transformation always exists if we have no constraints on spanning trees. In this paper, we wish to find a transformation which passes through only spanning trees satisfying some constraint. Our focus is bounding either the maximum degree or the diameter of spanning trees, and we give the following results. The problem with a lower bound on maximum degree is solvable in polynomial time, while the problem with an upper bound on maximum degree is PSPACE-complete. The problem with a lower bound on diameter is NP-hard, while the problem with an upper bound on diameter is solvable in polynomial time.

1 Introduction

Given an instance of some combinatorial search problem and two of its feasible solutions, a *reconfiguration problem* asks whether one solution can be transformed into the other in a step-by-step fashion, such that each intermediate solution is also feasible. Reconfiguration problems capture dynamic situations, where some solution is in place and we would like to move to a desired alternative solution without becoming infeasible. A systematic study of the complexity of reconfiguration problems was initiated in [13]. Recently the topic has gained a lot of attention in the context of CSP and graph problems, such as the independent set problem, the matching problem, and the dominating set problem. For an overview of recent results on reconfiguration problems, the reader is referred to the surveys of van den Heuvel [10] and Nishimura [17].

*This work is partially supported by JSPS and MEAE-MESRI under the Japan-France Integrated Action Program (SAKURA).

[†]Partially supported by ANR project GrR (ANR-18-CE40-0032).

[‡]Partially supported by JSPS KAKENHI Grant Numbers JP18H04091, JP19K11814 and JP20H05793, Japan.

[§]Partially supported by JSPS KAKENHI Grant Numbers 18H05291, JP20K11692, and 20H05795, Japan.

[¶]Partially supported by ANR project GrR (ANR-18-CE40-0032).

^{||}Partially supported by JSPS KAKENHI Grant Numbers JP18H04091, JP20K11666 and JP20H05794, Japan.

^{**}Partially supported by JST CREST Grant Numbers JPMJCR18K3 and JPMJCR1401, and JSPS KAKENHI Grant Numbers 19K20350 and JP20H05793, Japan.




Figure 1: A reconfiguration sequence from T_s to T_t (with no constraint on spanning trees). There is no reconfiguration sequence from T_s to T_t if we restrict spanning trees either with maximum degree at least three or with diameter at most two.

In this paper, our reference problem is the spanning tree problem. Let $G = (V, E)$ be a connected graph on n vertices. A *spanning tree* of G is a subgraph of G which is a tree (connected acyclic subgraph) and includes all the vertices in G . Spanning trees naturally arise in various situations such as routing or discrete geometry. In order to define a valid step-by-step transformation, an adjacency relation on the set of feasible solutions is needed. Let T_1 and T_2 be two spanning trees of G . We say that T_1 and T_2 are *adjacent* by an *edge flip* if there exist $e_1 \in E(T_1)$ and $e_2 \in E(T_2)$ such that $E(T_2) = (E(T_1) \setminus \{e_1\}) \cup \{e_2\}$. For two spanning trees T_s and T_t of G , a *reconfiguration sequence* (or simply a *transformation*) from T_s to T_t is a sequence of spanning trees $\langle T_0 := T_s, T_1, \dots, T_\ell := T_t \rangle$ such that two consecutive spanning trees are adjacent. Ito et al. [13] remarked that any spanning tree can be transformed into any other via a sequence of edge flips, which easily follows from the exchange property of the matroid bases.

In practice, we often need that spanning trees satisfy some additional desirable properties. Even if finding a spanning tree can be done in polynomial time, the problem becomes often NP-complete when additional constraints are added. In this paper, we consider spanning tree reconfiguration with additional constraints. More formally, we study the following questions: 1) does a transformation always exist when we add some constraints on the spanning trees all along the transformation? 2) If not, is it possible to decide efficiently if such a transformation exists? This question was already studied for spanning trees with restrictions on the number of leaves [2] or vertex modification between Steiner trees [16] for instance. If the answer to the first question is positive, it means that we can sample uniformly at random constrained spanning trees via a simple Monte Carlo Markov Chain. When the answer is negative, we might still want to find a transformation if possible between a fixed pair of solutions, for instance for updating a routing protocol in a network step by step without breaking the network and not over-requesting nodes during the transformation.

In this paper, we study RECONFIGURATION OF SPANNING TREES (RST) with degree constraints or with diameter constraints (See Figure 1.) We first describe the problem with degree constraints.

RST WITH SMALL (RESP. LARGE) MAXIMUM DEGREE

Input: A graph G , a positive integer d , and two spanning trees T_s and T_t in G with maximum degree at most (resp. at least) d .

Question: Is there a reconfiguration sequence from T_s to T_t such that any spanning tree in the sequence is of maximum degree at most (resp. at least) d ?

Bounding the maximum degree of spanning trees has applications for routing problems when we send data (i.e., a flow) along a spanning tree in a communication network. In this setting, the degree of a node is a measure of its load, and hence it is natural to bound the maximum degree in the spanning tree. In a complex dynamic networks, we want to reconfigure spanning trees on the fly to keep this property on the dynamic setting, which motivates us to study the reconfiguration problem.

The problem of finding a spanning tree with degree bounds is studied also from the theoretical point of view. Notice that spanning trees with bounds on the maximum degree include Hamiltonian paths that are spanning trees of maximum degree two. This implies that finding a spanning tree with maximum degree at most d is NP-hard. For restricted graph classes, this search problem is investigated in [4]. It is shown in [5] that if we relax the degree bound by one, then the search problem can be solved in polynomial time. Its

optimization variants are also studied in [6, 19].

We also study the problem with diameter constraints, which is formally stated as follows.

RST WITH SMALL (RESP. LARGE) DIAMETER

Input: A graph G , a positive integer d , and two spanning trees T_s and T_t in G with diameter at most (resp. at least) d .

Question: Is there a reconfiguration sequence from T_s to T_t such that any spanning tree in the sequence is of diameter at most (resp. at least) d ?

Spanning trees with largest possible diameter are Hamiltonian paths which receive a considerable attention. Spanning trees with upper bound on the diameter are for instance desirable in high-speed networks like optical networks since they minimize the worst-case propagation delay to all the nodes of the graphs, see e.g. [12]. We can find a spanning tree with minimum diameter in polynomial time [8], and some related problems have been studied in the literature [7, 20]

The problem of updating minimum spanning trees to maintain a valid spanning tree in dynamic networks is an important problem that received a considerable attention in the last decades, see for instance [1, 11]. In this situation, the graph is dynamic and is dynamically updated at each time step. The solution at time t , which might not be a solution anymore at time $t + 1$ (e.g. if edges of the spanning has been deleted from the graph), has to be modified with as few modifications as possible into a valid solution as good as possible. Spanning tree reconfiguration lies between the static situation (since the graph is fixed) and the dynamic situation (since the solution has to be modified).

Our Results

The contribution of this paper is to study the computational complexity of RST WITH SMALL (or LARGE) MAXIMUM DEGREE and RST WITH SMALL (or LARGE) DIAMETER.

Theorem 1. RST WITH LARGE MAXIMUM DEGREE *can be decided in polynomial time.*

Our proof for theorem 1 is in two steps. First we show that if there exists a vertex that has degree at least d in both T_s and T_t , then there is a reconfiguration sequence between them. Then, for two vertices u and v , we prove that we can decide in polynomial time if there exists a pair of adjacent spanning trees T and T' such that u has degree at least d in T and v has degree at least d in T' . These results together will imply theorem 1.

While the existence of a spanning tree with maximum degree at least d can be decided in polynomial time, it is NP-complete to find a spanning tree of maximum degree at most 2 (that is a Hamiltonian path). A similar behavior holds for RST with degree constraints.

Theorem 2. For every $d \geq 3$, RST WITH SMALL MAXIMUM DEGREE *is PSPACE-complete.*

The proof for theorem 2 consists of a reduction from NCL (Nondeterministic Constraint Logic), known to be PSPACE-complete [9]. This result is tight in the following sense: if at least one of T_s and T_t has maximum degree at most $d - 1$, then the problem becomes polynomial-time solvable (shown in Theorem 15). It is worth noting that this behavior is similar to the result for the search problem shown in [5]; while finding a spanning tree with maximum degree at most d is NP-hard, if we relax the degree bound by one, then the problem can be solved in polynomial time.

In the second part of the paper, we study RST WITH SMALL or LARGE DIAMETER.

Theorem 3. RST WITH LARGE DIAMETER *is NP-hard even restricted to planar graphs.*

The proof for theorem 3 consists of a reduction from the HAMILTONIAN PATH problem, which is not a reconfiguration problem but the original search problem. We note that since the length of a reconfiguration sequence is not necessarily bounded by a polynomial in the input size, it is unclear whether RST WITH LARGE DIAMETER belongs to the class NP. In a similar way to RST WITH SMALL MAXIMUM DEGREE, we conjecture that RST WITH LARGE DIAMETER is PSPACE-complete.

Finally, the main technical result of the paper is the following positive result.

Theorem 4. RST WITH SMALL DIAMETER is polynomial-time solvable.

The proof for theorem 4 follows a similar scheme to theorem 1. First we show that all the spanning trees with the same “center” can be transformed into any other. Therefore, it suffices to consider the transformation of the centers. However, for two vertices u and v , it is hard to determine whether there exists a pair of adjacent spanning trees T and T' such that u and v are centers of T and T' , respectively. Indeed, we do not know whether it can be done in polynomial time. The core of the proof is to focus on only “good” pairs of centers for which the existence of a desired pair of spanning trees can be tested in polynomial time (see theorem 23). A key ingredient of our proof consists in proving that if there is a reconfiguration sequence between the spanning trees, then there exists a sequence of centers from the initial center to the final center in which any consecutive centers form a good pair (see theorem 24).

Organization

The rest of this paper is organized as follows. We first give some preliminaries in Section 2. Next, Sections 3 and 4 are devoted to RST WITH LARGE MAXIMUM DEGREE (Theorem 1) and RST WITH SMALL MAXIMUM DEGREE (Theorem 2), respectively. Then, Sections 5 and 6 are devoted to RST WITH LARGE DIAMETER (Theorem 3) and RST WITH SMALL DIAMETER (Theorem 4), respectively. Some technical parts in the proof of Theorem 4 are deferred to Sections 7 and 8. Finally, we conclude this paper by giving some remarks in Section 9.

2 Preliminaries

Throughout this paper, we consider graphs that are simple and loopless. Let $G = (V, E)$ be a graph. For a vertex $v \in V$, we denote by $d_G(v)$ the degree of v in G , by $N_G(v)$ the (open) neighborhood of v in G , and by $\delta_G(v)$ the set of edges incident to v in G . Since G is simple, $d_G(v) = |N_G(v)| = |\delta_G(v)|$. For a tree T , a vertex v is a leaf if its degree is one, and is an internal node otherwise. A branching node is a vertex of degree at least three.

For a subgraph H of G and an $F \subseteq E$, we denote by $H - F$ the graph $(V(H), E(H) \setminus F)$ and by $H + F$ the graph $(V(H), E(H) \cup F)$. To avoid cumbersome notation, if $e \in E$, $H - \{e\}$ and $H + \{e\}$ will be denoted by $H - e$ and $H + e$, respectively.

For $u, v \in V$, the distance $\ell_G(u, v)$ between u and v is defined as the minimum number of edges in a shortest u - v path. For $v \in V$, the eccentricity $\epsilon_G(v)$ of v in G is the maximum distance between v and any vertex in G , that is, $\epsilon_G(v) := \max \{\ell_G(v, u) \mid u \in V\}$. The diameter $\text{diam}(G)$ of G is the maximum eccentricity among V . That is, $\text{diam}(G) := \max \{\epsilon_G(v) \mid v \in V\} = \max \{\ell_G(u, v) \mid u, v \in V\}$.

For two spanning trees T and T' , we denote $T \leftrightarrow T'$ if $|E(T) \setminus E(T')| = |E(T') \setminus E(T)| \leq 1$, that is, either $T = T'$ or T and T' are adjacent. We say that T_s is reconfigurable to T_t if there exists a reconfiguration sequence from T_s to T_t such that any spanning tree in the sequence satisfies a given degree/diameter constraint. When we have no degree/diameter constraints, since spanning trees form a base family of a matroid, the exchange property of the matroid bases ensures that there always exists a reconfiguration sequence between any pair of spanning trees.

Lemma 5 (see Proposition 1 in [13]). *Let G be a graph and T and T' be two spanning trees of G . There exists a reconfiguration sequence $\langle T = T_0, T_1, \dots, T_\ell = T' \rangle$ between T and T' such that for all $i \in \{0, 1, \dots, \ell\}$, the spanning tree T_i contains all the edges in $E(T) \cap E(T')$.*

3 Large Maximum Degree (Proof of Theorem 1)

In this section, we prove theorem 1, which we restate here.

Theorem 1. RST WITH LARGE MAXIMUM DEGREE can be decided in polynomial time.

Algorithm 1: Algorithm for RST WITH LARGE MAXIMUM DEGREE

Input: A graph G and two spanning trees T_s and T_t in G with \max . degree $\geq d$.

Output: Is T_s reconfigurable to T_t ?

- 1 Compute $\text{large}(T_s)$ and $\text{large}(T_t)$, and construct \mathcal{G} ;
 - 2 **if** there is a path between $\text{large}(T_s)$ and $\text{large}(T_t)$ in \mathcal{G} **then return YES**;
 - 3 **else return NO**;
-

Let (G, d, T_s, T_t) be an instance of RST WITH LARGE MAXIMUM DEGREE. For a spanning tree T in G , let $\text{large}(T) \subseteq V$ be the set of all the vertices of degree at least d in T , that is, $\text{large}(T) := \{v \in V \mid d_T(v) \geq d\}$. Note that T has maximum degree at least d if and only if $\text{large}(T) \neq \emptyset$. The following lemma is easy but is essential to prove theorem 1.

Lemma 6. *Let T_1 and T_2 be spanning trees in G with maximum degree at least d . If there exists a vertex $u \in \text{large}(T_1) \cap \text{large}(T_2)$, then T_1 is reconfigurable to T_2 .*

Proof. We show that T_1 is reconfigurable to T_2 by induction on $d - |\delta_{T_1}(u) \cap \delta_{T_2}(u)|$.

Suppose that $d - |\delta_{T_1}(u) \cap \delta_{T_2}(u)| \leq 0$ holds. By lemma 5, there exists a reconfiguration sequence from T_1 to T_2 in which all the spanning trees contain $\delta_{T_1}(u) \cap \delta_{T_2}(u)$. This shows that, for any spanning tree T' in the sequence, $|\delta_{T'}(u)| \geq |\delta_{T_1}(u) \cap \delta_{T_2}(u)| \geq d$. Hence, T_1 is reconfigurable to T_2 .

Suppose that $d - |\delta_{T_1}(u) \cap \delta_{T_2}(u)| \geq 1$ holds. Since $|\delta_{T_2}(u)| \geq d$ and $|\delta_{T_1}(u) \cap \delta_{T_2}(u)| \leq d - 1$, there exists an edge $e \in \delta_{T_2}(u) \setminus \delta_{T_1}(u)$. Since $T_1 + e$ contains a unique cycle C and T_2 contains no cycle, there exists an edge $f \in E(C) \setminus E(T_2)$. Then, we have that $f \in E(T_1) \setminus E(T_2)$ and $T'_1 := T_1 + e - f$ is a spanning tree in G . Observe that $|\delta_{T'_1}(u)| \geq |\delta_{T_1}(u) \cup \{e\}| - 1 \geq |\delta_{T_1}(u)| \geq d$, which shows that $u \in \text{large}(T'_1)$. We also see that $d - |\delta_{T'_1}(u) \cap \delta_{T_2}(u)| = d - |\delta_{T_1}(u) \cap \delta_{T_2}(u)| - 1$. Therefore, by the induction hypothesis, T'_1 is reconfigurable to T_2 . This shows that T_1 is reconfigurable to T_2 as T_1 and T'_1 are adjacent. \square

Our algorithm is based on testing the reachability in an auxiliary graph \mathcal{G} , which is defined as follows. The vertex set of \mathcal{G} is defined as V , where each vertex v in $V(\mathcal{G})$ corresponds to the set of spanning trees T with $v \in \text{large}(T)$. For any pair u, v of distinct vertices in $V(\mathcal{G})$, there is an edge $uv \in E(\mathcal{G})$ if and only if there exist spanning trees T and T' such that $u \in \text{large}(T)$, $v \in \text{large}(T')$, and $T \leftrightarrow T'$ (possibly $T = T'$). Then, by definition of the auxiliary graph and Lemma 6, we have the following lemma.

Lemma 7. *Let T_s and T_t be spanning trees with maximum degree at least d . Then, T_s is reconfigurable to T_t if and only if \mathcal{G} contains a path from $\text{large}(T_s)$ to $\text{large}(T_t)$.*

Proof. We first show the “only if” part. Suppose that there exists a reconfiguration sequence $\langle T_s = T_0, T_1, \dots, T_k = T_t \rangle$ from T_s to T_t , where T_i is a spanning tree of maximum degree at least d for any $i \in \{0, 1, \dots, k\}$ and T_i and T_{i+1} are adjacent for any $i \in \{0, 1, \dots, k-1\}$. For each i , let v_i be a vertex in $\text{large}(T_i)$. By the definition of \mathcal{G} , for $i \in \{0, 1, \dots, k-1\}$, we have either $v_i = v_{i+1}$ or \mathcal{G} contains an edge $v_i v_{i+1}$. Since $v_0 \in \text{large}(T_s)$ and $v_k \in \text{large}(T_t)$, \mathcal{G} contains a path from $\text{large}(T_s)$ to $\text{large}(T_t)$.

To show the “if” part, suppose that \mathcal{G} contains a path (v_0, v_1, \dots, v_k) from $\text{large}(T_s)$ to $\text{large}(T_t)$. For $i \in \{0, 1, \dots, k-1\}$, $v_i v_{i+1} \in E(\mathcal{G})$ means that there exist two spanning trees T_i^+ and T_{i+1}^- such that $v_i \in \text{large}(T_i^+)$, $v_{i+1} \in \text{large}(T_{i+1}^-)$, and $T_i^+ \leftrightarrow T_{i+1}^-$. Let $T_0^- := T_s$ and $T_k^+ := T_t$. Then, for $i \in \{0, 1, \dots, k\}$, since $v_i \in \text{large}(T_i^-) \cap \text{large}(T_i^+)$, T_i^- is reconfigurable to T_i^+ by lemma 6. This together with $T_i^+ \leftrightarrow T_{i+1}^-$ shows that T_s is reconfigurable to T_t . \square

By this lemma, we can solve RST WITH LARGE MAXIMUM DEGREE by detecting a path from $\text{large}(T_s)$ to $\text{large}(T_t)$ in \mathcal{G} (see algorithm 1 for a pseudocode of our algorithm). Our remaining task is to construct the auxiliary graph \mathcal{G} in polynomial time which is possible by the following lemma.




Figure 2: Case when $uv \notin E(G)$.




Figure 3: Cycle C and edge e' .

Lemma 8. For two distinct vertices $u, v \in V$, there exists an edge $uv \in E(\mathcal{G})$ if and only if $|N_G(u)| \geq d$, $|N_G(v)| \geq d$, and

$$|N_G(u) \cup N_G(v)| \geq \begin{cases} 2d - 1 & \text{if } uv \in E(G), \\ 2d - 2 & \text{otherwise.} \end{cases} \quad (1)$$

Proof. We first prove the “only-if” direction. Suppose that \mathcal{G} contains an edge uv , that is, there exist spanning trees T and T' such that $u \in \text{large}(T)$, $v \in \text{large}(T')$, and $T \leftrightarrow T'$ (possibly $T = T'$). Then, $|N_G(u)| \geq d$ and $|N_G(v)| \geq d$ are obvious. Since T contains no cycle, we know that $N_T(u)$ and $N_T(v)$ contain at most one common vertex. Then, we obtain

$$\begin{aligned} |N_G(u) \cup N_G(v)| &\geq |N_T(u) \cup N_T(v)| \\ &\geq |N_T(u)| + |N_T(v)| - 1 && \text{(by } |N_T(u) \cap N_T(v)| \leq 1) \\ &\geq |N_T(u)| + (|N_{T'}(v)| - 1) - 1 && \text{(by } |E(T') \setminus E(T)| \leq 1) \\ &\geq 2d - 2. \end{aligned} \quad (2)$$

Similarly, if $uv \in E(G) \setminus E(T)$, then we obtain

$$|N_G(u) \cup N_G(v)| \geq |N_T(u) \cup N_T(v) \cup \{u, v\}| \geq |N_T(u)| + |N_T(v)| + 1 \geq 2d. \quad (3)$$

If $uv \in E(T)$, then $N_T(u) \cap N_T(v) = \emptyset$ holds, and hence we obtain

$$|N_G(u) \cup N_G(v)| \geq |N_T(u) \cup N_T(v)| = |N_T(u)| + |N_T(v)| \geq 2d - 1. \quad (4)$$

By (2), (3), and (4), we obtain (1).

We next prove the “if” direction. Suppose that $|N_G(u)| \geq d$, $|N_G(v)| \geq d$, and (1) hold. For each of the following two cases, we define an edge set $F \subseteq E$.

- Suppose that $uv \notin E(G)$ holds (Figure 2). Let $S_u \subseteq N_G(u)$ be a vertex set with $|S_u| = d$ that maximizes $|S_u \setminus N_G(v)|$. Then, we have either $S_u \subseteq N_G(u) \setminus N_G(v)$ or $S_u \supseteq N_G(u) \setminus N_G(v)$. If $S_u \subseteq N_G(u) \setminus N_G(v)$, then let $S_v \subseteq N_G(v)$ be a vertex set with $|S_v| = d - 1$. Otherwise, let $S_v \subseteq N_G(v)$ be a vertex set such that $|S_v| = d - 1$ and $|S_u \cap S_v| = 1$, where such S_v exists because $|N_G(v) \setminus S_u| = |(N_G(u) \cup N_G(v)) \setminus S_u| \geq d - 2$ and $|N_G(v) \cap S_u| \geq 1$. In either case, we obtain $S_u \subseteq N_G(u)$ and $S_v \subseteq N_G(v)$ such that $|S_u| = d$, $|S_v| = d - 1$, and $|S_u \cap S_v| \leq 1$. Define $F := \{uw \mid w \in S_u\} \cup \{vw \mid w \in S_v\}$.
- Suppose that $uv \in E(G)$ holds. Since $|N_G(u) \setminus \{v\}| \geq d - 1$, $|N_G(v) \setminus \{u\}| \geq d - 1$, and $|(N_G(u) \setminus \{v\}) \cup (N_G(v) \setminus \{u\})| \geq 2d - 3$, by the same argument as above, we can take $S_u \subseteq N_G(u) \setminus \{v\}$ and $S_v \subseteq N_G(v) \setminus \{u\}$ such that $|S_u| = d - 1$, $|S_v| = d - 2$, and $S_u \cap S_v = \emptyset$. Define $F := \{uw \mid w \in S_u\} \cup \{vw \mid w \in S_v\} \cup \{uv\}$.

In both cases, it holds that $|F \cap \delta_G(u)| = d$, $|F \cap \delta_G(v)| = d - 1$, and F contains no cycle. Therefore, there exists a spanning tree T with $E(T) \supseteq F$ such that $|\delta_T(u)| \geq |F \cap \delta_G(u)| = d$ and $|\delta_T(v)| \geq |F \cap \delta_G(v)| = d - 1$. If $|\delta_T(v)| \geq d$, then we obtain $\{u, v\} \subseteq \text{large}(T)$, which shows that $uv \in E(\mathcal{G})$. Therefore, it suffices to consider the case when $|\delta_T(v)| = d - 1$. Since $|\delta_G(v)| \geq d$, there exists an edge $e \in \delta_G(v) \setminus \delta_T(v)$. Let C be the unique cycle in $T + e$ and e' be an edge in $E(C) \setminus \delta_T(v)$ (see Figure 3). Then, $T' := T + e - e'$ is a spanning tree such that $|\delta_{T'}(v)| = |\delta_T(v) \cup \{e\}| = d$, which means that $v \in \text{large}(T')$. Since T and T' are adjacent, we obtain $uv \in E(\mathcal{G})$. \square

Since we can easily check the inequality (1) for each pair of vertices u and v , Lemma 8 ensures that the auxiliary graph \mathcal{G} can be constructed in polynomial time. Therefore, algorithm 1 correctly decides RST WITH LARGE MAXIMUM DEGREE in polynomial time, which completes the proof of Theorem 1. Note that all the proofs are constructive, and hence we can find a desired reconfiguration sequence from T_s to T_t in polynomial time if it exists.

4 Small Maximum Degree

In this section, we consider RST WITH SMALL MAXIMUM DEGREE. We first show the PSPACE-completeness in Section 4.1. In contrast, we show in Section 4.2 that if at least one of T_s and T_t has maximum degree at most $d - 1$, then an instance (G, d, T_s, T_t) of RST WITH SMALL MAXIMUM DEGREE is a YES-instance.

4.1 PSPACE-Completeness (Proof of Theorem 2)

In this subsection, we prove theorem 2, i.e., we show that RST WITH SMALL MAXIMUM DEGREE is PSPACE-complete. The problem is indeed in PSPACE. We prove the PSPACE-hardness by giving a polynomial reduction from *Reconfiguration of Nondeterministic Constraint Logic on AND/OR graphs*, which we call NCL RECONFIGURATION for short.

Suppose that we are given a cubic graph with edge-weights such that each vertex is either incident to three weight-2 edges (“OR vertex”) or one weight-2 edge and two weight-1 edges (“AND vertex”), which we call an AND/OR graph. An NCL configuration is an orientation of the edges in the graph such that the total weights of incoming arcs at each vertex is at least two. Two NCL configurations are *adjacent* if they differ in a single edge direction. In NCL RECONFIGURATION, we are given an AND/OR graph and its two NCL configurations, and the objective is to determine whether there exists a sequence of adjacent NCL configurations that transforms one into the other. It is shown in [9] that NCL RECONFIGURATION is PSPACE-complete. In what follows, we give a polynomial reduction from NCL RECONFIGURATION to RST WITH SMALL MAXIMUM DEGREE.

Construction of the graph. Suppose that we are given an instance of NCL RECONFIGURATION, that is, an AND/OR graph $H = (V(H), E(H))$ and two configurations σ_s and σ_t of H . Fix $d \geq 3$. We first construct a graph $G' = (V', E')$, a vertex subset $L \subseteq V'$, and an integer $b(v) \in \{1, 2, 3\}$ for each $v \in V'$, and then construct a graph $G = (V, E)$ by using G', L , and b . We consider an instance (G, d, T_s, T_t) of RST WITH SMALL MAXIMUM DEGREE, where T_s and T_t will be defined later. The construction of G', L , and b is described as follows.

- We initialize $G' = (V', E')$ and L as the empty graph and the empty set, respectively.
- For a vertex $u \in V(H)$ and an edge $e \in \delta_H(u)$, we introduce a vertex $v_{u,e}$ in V' . Let $b(v_{u,e}) = 2$.
- For an edge $e \in E(H)$ connecting u and u' , we introduce a vertex v_e in V' and two edges $v_e v_{u,e}$ and $v_e v_{u',e}$ in E' (Figure 4). Let $b(v_e) = 1$.
- For an OR vertex $u \in V(H)$ with $\delta_H(u) = \{e_1, e_2, e_3\}$, we introduce a vertex r_u in V' and an edge $r_u v_{u,e_i}$ in E' for $i \in \{1, 2, 3\}$ (Figure 5). Let $b(r_u) = 1$. Add v_{u,e_i} to L for $i \in \{1, 2, 3\}$.
- For an AND vertex $u \in V(H)$ with $\delta_H(u) = \{e_0, e_1, e_2\}$, where e_0 is a weight-2 edge and e_1 and e_2 are weight-1 edges, we introduce four vertices $r_u, w_u, x_u,$ and y_u in V' , and seven edges $v_{u,e_0} r_u, r_u w_u, w_u x_u, w_u y_u, x_u v_{u,e_1}, y_u v_{u,e_1} v_{u,e_2}$ in E' (Figure 6). We denote by E'_u the set of these seven edges. Let $b(r_u) = 1, b(w_u) = 3,$ and $b(x_u) = b(y_u) = 2$. Add v_{u,e_0} and w_u to L .

We next construct $G = (V, E)$ by adding new vertices and edges to $G' = (V', E')$ as follows (see Figure 7 for an illustration).




Figure 4: The gadget for an edge.




Figure 5: The gadget for an OR vertex.




Figure 6: The gadget for an AND vertex.




Figure 7: Construction of G from G' .

- We construct a tree $T^* = (V(T^*), E(T^*))$ of maximum degree at most three such that $V(T^*) \cap V' = L$, $E(T^*) \cap E' = \emptyset$, and L is the set of all the leaves of T^* . Then, we attach T^* to G' . We denote the obtained graph by $G' + T^*$.
- For each vertex $v \in V'$, we add $d - b(v)$ new vertices $\bar{v}_1, \dots, \bar{v}_{d-b(v)}$ and new edges $v\bar{v}_1, \dots, v\bar{v}_{d-b(v)}$.

Correspondence between solutions. In order to see the correspondence between NCL configurations in H and spanning trees in G with maximum degree at most d , we begin with the following easy lemma.

Lemma 9. *Any spanning tree T in G with maximum degree at most d satisfies the following properties: (a) $v\bar{v}_i \in E(T)$ for $v \in V'$ and for $i \in \{1, 2, \dots, d - b(v)\}$; (b) $|\delta_{G'+T^*}(v) \cap E(T)| \leq b(v)$ for $v \in V'$; (c) T contains exactly one of $v_e v_{u,e}$ and $v_e v_{u',e}$ for $e = uu' \in E(H)$; and (d) $E(T^*) \subseteq E(T)$.*

Proof. Since T is a spanning tree with maximum degree at most d , (a), (b), and (c) are obvious. By (b), $T - E(T^*)$ contains no path connecting two distinct components of $G' - \{v \in V' \mid b(v) = 1\}$. Since each connected component of $G' - \{v \in V' \mid b(v) = 1\}$ contains exactly one vertex in L , for any pair of vertices $v_1, v_2 \in L$, the unique v_1 - v_2 path in T^* must be contained in T . This shows that $E(T^*) \subseteq E(T)$, because L is the set of all the leaves of T^* . \square

For a spanning tree T in G with maximum degree at most d , we define an orientation σ_T of H as follows: an edge $e = uu' \in E(H)$ is inward for u if $v_e v_{u',e} \in E(T)$, and it is outward for u if $v_e v_{u,e} \in E(T)$. This defines an orientation of H by lemma 9 (c). The following two lemmas show the correspondence between NCL configurations in H and spanning trees in G with maximum degree at most d .

Lemma 10. *For any spanning tree T in G with maximum degree at most d , the orientation σ_T is an NCL configuration of H .*

Proof. It suffices to show that, for any $u \in V(H)$, the total weights of incoming arcs at u is at least two in σ_T .

First, let $u \in V(H)$ be an OR-vertex with $\delta_H(u) = \{e_1, e_2, e_3\}$. Since T is a spanning tree, it holds that $r_u v_{u,e_i} \in E(T)$ for some $i \in \{1, 2, 3\}$. Then, since $|\delta_{G'}(v_{u,e_i}) \cap E(T)| \leq b(v_{u,e_i}) - 1 = 1$ by (b) and (d) in lemma 9, it holds that $v_{e_i} v_{u,e_i} \notin E(T)$. This means that e_i is inward for u in σ_T , and hence the total weights of incoming arcs at u is at least two.

Second, let $u \in V(H)$ be an AND-vertex with $\delta_H(u) = \{e_0, e_1, e_2\}$, where e_0 is a weight-2 edge and e_1 and e_2 are weight-1 edges. Since T is a spanning tree, we have either $r_u v_{u,e_0} \in E(T)$ or $r_u w_u \in E(T)$. If $r_u v_{u,e_0} \in E(T)$, then e_0 is inward for u in σ_T , which implies that the total weights of incoming arcs at u is at least two. Therefore, it suffices to consider the case when $r_u w_u \in E(T)$. Since $|\delta_{G'}(v) \cap E(T)| \leq 2$ for $v \in \{w_u, x_u, y_u, v_{u,e_1}, v_{u,e_2}\}$ by (b) and (d) in lemma 9, we have either $\{r_u w_u, w_u x_u, x_u v_{u,e_1}, v_{u,e_1} v_{u,e_2}, v_{u,e_2} y_u\} \subseteq E(T)$ or $\{r_u w_u, w_u y_u, y_u v_{u,e_2}, v_{u,e_2} v_{u,e_1}, v_{u,e_1} x_u\} \subseteq E(T)$. In either case, $v_{e_i} v_{u,e_i} \notin E(T)$ for $i \in \{1, 2\}$, because $|\delta_{G'}(v_{u,e_i}) \cap E(T)| \leq 2$. This means that e_i is inward for u in σ_T for $i \in \{1, 2\}$, and hence the total weights of incoming arcs at u is at least two.

Therefore, σ_T is an NCL configuration of H . \square

Lemma 11. *For any NCL configuration σ of H , we can construct a spanning tree T in G with maximum degree at most d such that $\sigma_T = \sigma$ in polynomial time.*

Proof. Given an NCL configuration σ of H , we construct a spanning subgraph T of G such that

$$E(T) := E(T^*) \cup \{v\bar{v}_i \mid v \in V', i \in \{1, 2, \dots, d - b(v)\}\} \cup \{f_e \mid e \in E(H)\} \cup \bigcup_{u \in V(H)} F_u,$$

where an edge f_e for $e \in E(H)$ and an edge set F_u for $u \in V(H)$ are defined as follows.

- For an edge $e = uu' \in E(H)$, define $f_e := v_e v_{u',e}$ if e is inward for u in σ and define $f_e := v_e v_{u,e}$ otherwise.
- For an OR-vertex $u \in V(H)$ with $\delta_H(u) = \{e_1, e_2, e_3\}$, choose an arbitrarily edge e_i that is inward for u in σ and define $F_u := \{r_u v_{u,e_i}\}$. Note that such e_i exists as σ is an NCL configuration.
- For an AND-vertex $u \in V(H)$ with $\delta_H(u) = \{e_0, e_1, e_2\}$, where e_0 is a weight-2 edge and e_1 and e_2 are weight-1 edges, define $F_u = E'_u \setminus \{r_u w_u, v_{u,e_1} v_{u,e_2}\}$ if e_0 is inward for u in σ , and define $F_u := E'_u \setminus \{r_u v_{u,e_0}, w_u y_u\}$ otherwise.

Then, T is a spanning tree in G with maximum degree at most d such that $\sigma_T = \sigma$, which completes the proof. \square

For two NCL configurations σ_s and σ_t of H , by lemma 11, we can construct spanning trees T_s and T_t in G with maximum degree at most d such that $\sigma_{T_s} = \sigma_s$ and $\sigma_{T_t} = \sigma_t$. This yields an instance (G, d, T_s, T_t) of RST WITH SMALL MAXIMUM DEGREE.

Correctness. In order to show the PSPACE-hardness of RST WITH SMALL MAXIMUM DEGREE, we show that the original instance (H, σ_s, σ_t) of NCL RECONFIGURATION is equivalent to the obtained instance (G, d, T_s, T_t) of RST WITH SMALL MAXIMUM DEGREE, that is, we prove that (H, σ_s, σ_t) is a YES-instance if and only if (G, d, T_s, T_t) is a YES-instance. To this end, we use the following lemma.

Lemma 12. *Let T_1 and T_2 be spanning trees in G with maximum degree at most d . If σ_{T_1} and σ_{T_2} are adjacent, then there is a reconfiguration sequence from T_1 to T_2 in which all the spanning trees have maximum degree at most d .*

Proof. Let $e^* \in E(H)$ be the unique edge in H whose direction is different in σ_1 and σ_2 . We prove the existence of a reconfiguration sequence by induction on $|E(T_1) \setminus E(T_2)|$. If $|E(T_1) \setminus E(T_2)| = 1$, then T_1 and T_2 are adjacent, and hence the claim is obvious. Suppose that $|E(T_1) \setminus E(T_2)| \geq 2$. Since σ_{T_1} and σ_{T_2} are adjacent, there exists a vertex $u \in V(H)$ such that T_1 and T_2 contain different edge sets in the gadget corresponding to u . That is, $(E(T_1) \setminus E(T_2)) \cap \delta_{G'}(r_u) \neq \emptyset$ for an OR-vertex $u \in V(H)$ or $(E(T_1) \setminus E(T_2)) \cap E'_u \neq \emptyset$ for an AND-vertex $u \in V(H)$. We fix such a vertex $u \in V(H)$.

Suppose that u is an OR-vertex such that $(E(T_1) \setminus E(T_2)) \cap \delta_{G'}(r_u) \neq \emptyset$. In this case, $E(T_1) \cap \delta_{G'}(r_u) = \{r_u v_{u,e_i}\}$ and $E(T_2) \cap \delta_{G'}(r_u) = \{r_u v_{u,e_j}\}$ for some distinct $i, j \in \{1, 2, 3\}$. By changing the roles of T_1 and T_2 if necessary, we may assume that either $e^* \notin \delta_H(u)$ or e^* is inward for u in σ_1 . Then, $T'_1 := T_1 - r_u v_{u,e_i} + r_u v_{u,e_j}$ is a spanning tree with maximum degree at most d such that T'_1 is adjacent to T_1 ,

$\sigma_{T'_1} = \sigma_{T_1}$, and $|E(T'_1) \setminus E(T_2)| = |E(T_1) \setminus E(T_2)| - 1$. By induction hypothesis, T'_1 is reconfigurable to T_2 , and hence T_1 is reconfigurable to T_2 .

Suppose that u is an AND-vertex such that $|(E(T_1) \setminus E(T_2)) \cap E'_u| = 1$. By changing the roles of T_1 and T_2 if necessary, we may assume that either $e^* \notin \delta_H(u)$ or e^* is inward for u in σ_{T_1} . Then, $T'_1 := T_1 - (E(T_1) \cap E'_u) + (E(T_2) \cap E'_u)$ is a spanning tree with maximum degree at most d such that T'_1 is adjacent to T_1 , $\sigma_{T'_1} = \sigma_{T_1}$, and $|E(T'_1) \setminus E(T_2)| = |E(T_1) \setminus E(T_2)| - 1$. By induction hypothesis, T'_1 is reconfigurable to T_2 , and hence T_1 is reconfigurable to T_2 .

The remaining case is that u is an AND-vertex such that $|(E(T_1) \setminus E(T_2)) \cap E'_u| \geq 2$. Since each of T_1 and T_2 contains exactly one edge in $\delta_{G'}(r_u)$ and exactly four edges in $E'_u \setminus \delta_{G'}(r_u)$, we have that $|(E(T_1) \setminus E(T_2)) \cap \delta_{G'}(r_u)| = 1$ and $|(E(T_1) \setminus E(T_2)) \cap (E'_u \setminus \delta_{G'}(r_u))| = 1$. By changing the roles of T_1 and T_2 if necessary, we may assume that $E(T_1) \cap \delta_{G'}(r_u) = \{r_u v_{u, \epsilon_0}\}$ and $E(T_2) \cap \delta_{G'}(r_u) = \{r_u w_u\}$. This implies that $v_{u, \epsilon_0} v_{\epsilon_0} \notin E(T_1)$, and hence ϵ_0 is inward for u in σ_{T_1} . We consider the following two cases separately.

- Suppose that ϵ_0 is inward for u in σ_{T_2} . In this case, $T'_2 := T_2 - r_u w_u + r_u v_{u, \epsilon_0}$ is a spanning tree with maximum degree at most d such that T'_2 is adjacent to T_2 , $\sigma_{T'_2} = \sigma_{T_2}$, and $|E(T_1) \setminus E(T'_2)| = |E(T_1) \setminus E(T_2)| - 1$. By induction hypothesis, T_1 is reconfigurable to T'_2 , and hence T_1 is reconfigurable to T_2 .
- Suppose that ϵ_0 is outward for u in σ_{T_2} . In this case, ϵ_1 and ϵ_2 are inward for u in σ_{T_2} by lemma 10. Furthermore, since $e^* = \epsilon_0$ holds, ϵ_1 and ϵ_2 are inward for u also in σ_{T_1} , that is, $v_{u, \epsilon_1} v_{\epsilon_1}, v_{u, \epsilon_2} v_{\epsilon_2} \notin E(T_1)$. Then, $T'_1 := T_1 - (E(T_1) \cap (E'_u \setminus \delta_{G'}(r_u))) + (E(T_2) \cap (E'_u \setminus \delta_{G'}(r_u)))$ is a spanning tree with maximum degree at most d such that T'_1 is adjacent to T_1 , $\sigma_{T'_1} = \sigma_{T_1}$, and $|E(T'_1) \setminus E(T_2)| = |E(T_1) \setminus E(T_2)| - 1$. By induction hypothesis, T'_1 is reconfigurable to T_2 , and hence T_1 is reconfigurable to T_2 .

By the above argument, there is a reconfiguration sequence from T_1 to T_2 . \square

We are now ready to show the equivalence of (H, σ_s, σ_t) and (G, d, T_s, T_t) .

Lemma 13. *Let (H, σ_s, σ_t) be an instance of NCL RECONFIGURATION and (G, d, T_s, T_t) be an instance of RST WITH SMALL MAXIMUM DEGREE obtained by the above construction. Then, (H, σ_s, σ_t) is a YES-instance if and only if (G, d, T_s, T_t) is a YES-instance.*

Proof. We first show the “if” part. Suppose that there exists a reconfiguration sequence $\langle T_s = T_0, T_1, \dots, T_k = T_t \rangle$ from T_s to T_t , where T_i is a spanning tree in G with maximum degree at most d for any $i \in \{0, 1, \dots, k\}$ and T_i and T_{i+1} are adjacent for any $i \in \{0, 1, \dots, k-1\}$. Then, σ_{T_i} is an NCL configuration of H for $i \in \{0, 1, \dots, k\}$ by lemma 10, and we have either $\sigma_{T_i} = \sigma_{T_{i+1}}$ or σ_{T_i} and $\sigma_{T_{i+1}}$ are adjacent for $i \in \{0, 1, \dots, k-1\}$ as $|E(T_i) \setminus E(T_{i+1})| \leq 1$. Since $\sigma_{T_0} = \sigma_{T_s} = \sigma_s$ and $\sigma_{T_k} = \sigma_{T_t} = \sigma_t$, there exists a sequence of adjacent NCL configurations from σ_s to σ_t .

To show the “only-if” part, suppose that there exists a reconfiguration sequence $\langle \sigma_s = \sigma_0, \sigma_1, \dots, \sigma_k = \sigma_t \rangle$, where σ_i is an NCL configuration of H for any $i \in \{0, 1, \dots, k\}$ and σ_i and σ_{i+1} are adjacent for any $i \in \{0, 1, \dots, k-1\}$. For $i \in \{1, 2, \dots, k-1\}$, let T_i be a spanning tree in G with maximum degree at most d such that $\sigma_{T_i} = \sigma_i$, whose existence is guaranteed by Lemma 11. Let $T_0 := T_s$ and $T_k := T_t$. Since lemma 12 shows that there is a reconfiguration sequence from T_i to T_{i+1} for $i \in \{0, 1, \dots, k-1\}$, T_s is reconfigurable to T_t . \square

This lemma shows that the above construction gives a polynomial reduction from NCL RECONFIGURATION to RST WITH SMALL MAXIMUM DEGREE. Therefore, RST WITH SMALL MAXIMUM DEGREE is PSPACE-hard, which completes the proof of theorem 2.

4.2 A Solvable Special Case

In this subsection, we show a sufficient condition for the reconfigurability of instances. The condition is as follows; at least one of T_s and T_t has maximum degree at most $d - 1$. Without loss of generality, we may assume that T_t satisfies the condition. Under this assumption, we have the following lemma.

Lemma 14. *Suppose that (G, d, T_s, T_t) is an instance of RST WITH SMALL MAXIMUM DEGREE such that T_t has maximum degree at most $d - 1$. There exists an edge $e = xy \in E(T_t) \setminus E(T_s)$ such that $d_{T_s}(x) \leq d - 1$ and $d_{T_s}(y) \leq d - 1$.*

Proof. To derive a contradiction, assume that Lemma 14 does not hold, that is, for any $e = xy \in E(T_t) \setminus E(T_s)$, we have $d_{T_s}(x) = d$ or $d_{T_s}(y) = d$. Let $T_s^* := T_s - E(T_t)$ and $T_t^* := T_t - E(T_s)$. Note that $|E(T_s^*)| = |E(T_t^*)|$, because both T_s and T_t are spanning trees in G . Let $S := \{v \in V \mid d_{T_s}(v) = d, d_{T_t^*}(v) \geq 1\}$. With this notation, the assumption means that S forms a vertex cover of T_t^* . In what follows, we compare $|E(T_s^*)|$ and $|E(T_t^*)|$.

Let $X_1 := \{v \in V \mid d_{T_s^*}(v) = 1\}$ and $X_{\geq 2} := \{v \in V \mid d_{T_s^*}(v) \geq 2\}$. Then, we see that

$$\frac{1}{2} \sum_{v \in V} d_{T_s^*}(v) = |E(T_s^*)| < |X_1 \cup X_{\geq 2}|, \quad (5)$$

because T_s^* is a forest. We also see that, for any $v \in S$,

$$d_{T_s^*}(v) = d - |\delta_{T_s}(v) \cap \delta_{T_t}(v)| \geq d_{T_t}(v) + 1 - |\delta_{T_s}(v) \cap \delta_{T_t}(v)| = d_{T_t^*}(v) + 1 \quad (6)$$

holds, and hence $S \subseteq X_{\geq 2}$. With these observations, we obtain

$$\begin{aligned} |E(T_s^*)| &= \sum_{v \in V} d_{T_s^*}(v) - \frac{1}{2} \sum_{v \in V} d_{T_s^*}(v) \\ &> \sum_{v \in V} d_{T_s^*}(v) - |X_1 \cup X_{\geq 2}| && \text{(by (5))} \\ &= \sum_{v \in X_{\geq 2}} (d_{T_s^*}(v) - 1) \\ &\geq \sum_{v \in S} (d_{T_s^*}(v) - 1) && \text{(by } S \subseteq X_{\geq 2}\text{)} \\ &\geq \sum_{v \in S} d_{T_t^*}(v) && \text{(by (6))} \\ &\geq |E(T_t^*)|, && \text{(because } S \text{ is a vertex cover of } T_t^*\text{)} \end{aligned}$$

which is a contradiction to $|E(T_s^*)| = |E(T_t^*)|$. Therefore, Lemma 14 holds. \square

Let $e \in E(T_t) \setminus E(T_s)$ be the edge as in the lemma and let $e' \in E(T_s) \setminus E(T_t)$ be an edge such that $T'_s := T_s + e - e'$ is a spanning tree in G . Note that the maximum degree of T'_s is at most d . Since $|E(T'_s) \setminus E(T_t)| = |E(T_s) \setminus E(T_t)| - 1$, (G, d, T'_s, T_t) is a YES-instance by induction. This implies that T'_s is reconfigurable to T_t , and thus the following theorem holds.

Theorem 15. *If at least one of T_s and T_t has maximum degree at most $d - 1$, then an instance (G, d, T_s, T_t) of RST WITH SMALL MAXIMUM DEGREE is a YES-instance.*

Note that the above discussion shows that we can find a reconfiguration sequence $\langle T_s = T_0, T_1, \dots, T_k = T_t \rangle$ with $k = |E(T_s) \setminus E(T_t)|$, which is a shortest reconfiguration sequence, in polynomial time. Moreover, Theorem 15 implies the following corollary.

Corollary 16. *Let G be a graph and d be a positive integer. If G contains a spanning tree with maximum degree at most $d - 1$, then any instance (G, d, T_s, T_t) of RST WITH SMALL MAXIMUM DEGREE is a YES-instance.*

Proof. Let T^* be a spanning tree in G with maximum degree at most $d - 1$. Then, Theorem 15 shows that T_s is reconfigurable to T^* and T^* is reconfigurable to T_t . Hence, T_s is reconfigurable to T_t , which completes the proof. \square

We note that it is not easy to determine whether or not G contains a spanning tree with maximum degree at most $d - 1$ even when $d = 3$, because finding a Hamiltonian path in cubic graphs is NP-hard.

5 Large Diameter (Proof of Theorem 3)

In this section, we prove theorem 3, which we restate here.

Theorem 3. RST WITH LARGE DIAMETER is NP-hard even restricted to planar graphs.

To prove the theorem, we give a polynomial reduction from HAMILTONIAN PATH problem to RST WITH LARGE DIAMETER. A *Hamiltonian path* of a graph G is a path that visits each vertex of G exactly once. Given a graph $G = (V, E)$ and two vertices $s, t \in V$, the HAMILTONIAN PATH problem asks to determine whether or not G has a Hamiltonian path whose endpoints are s and t , which is known to be NP-hard [14].

Reduction. Let (G', s', t') be an instance of HAMILTONIAN PATH. We may assume that G' is connected, since otherwise (G', s', t') is trivially a NO-instance. We construct a corresponding instance (G, d, T_s, T_t) of RST WITH LARGE DIAMETER as follows (see Figure 8).

Let n' be the number of vertices in G' , that is, $n' = |V(G')|$. We first add three vertices t_1, t_2 , and t_3 , and five edges $t't_1, t't_2, t_1t_2, t_1t_3$, and t_2t_3 to G' . Let D be the subgraph induced by $\{t', t_1, t_2, t_3\}$, which is isomorphic to the so-called diamond graph. We then add three paths $P_x = (x_{3n'}, x_{3n'-1}, \dots, x_1, s')$, $P_y = (s', y_1, y_2, \dots, y_{n'-3}, t')$, and $P_z = (t_3, z_1, z_2, \dots, z_{3n'})$, where all the vertices in P_x, P_y , and P_z except for s', t' , and t_3 are distinct new vertices. Note that $|E(P_x)| = |E(P_z)| = 3n'$ and $|E(P_y)| = n' - 2$. Let $G = (V, E)$ be the obtained graph and set $d = 7n' + 1$.

Let F' be an arbitrary spanning forest in G' such that F' consists of two connected components (trees) of which one contains s' and the other contains t' . Then, define spanning trees T_s and T_t in G by

$$\begin{aligned} E(T_s) &= E(P_x) \cup E(P_y) \cup E(P_z) \cup E(F') \cup \{t't_1, t_1t_2, t_2t_3\}, \\ E(T_t) &= E(P_x) \cup E(P_y) \cup E(P_z) \cup E(F') \cup \{t't_2, t_1t_2, t_1t_3\}. \end{aligned}$$

We notice that $E(T_s) \setminus E(F')$ forms a path in T_s of length $d = 7n' + 1$, and hence $\text{diam}(T_s) \geq d$. Similarly, $E(T_t) \setminus E(F')$ forms a path in T_t of length d , and hence $\text{diam}(T_t) \geq d$. This completes the construction of the instance (G, d, T_s, T_t) of RST WITH LARGE DIAMETER.

Correctness. In the following, we show that G' contains a Hamiltonian path from s' to t' if and only if (G, d, T_s, T_t) is a YES-instance. The following lemma shows that the diameter of a spanning tree is dominated by the distance between $x_{3n'}$ and $z_{3n'}$.

Lemma 17. For any spanning tree T in G , $\text{diam}(T) = \bar{\ell}_T(x_{3n'}, z_{3n'})$.

Proof. Let T be a spanning tree in G and P^* be a longest path in T . For $x, y \in V$ and for a spanning tree T in G , we denote by $T[x, y]$ the unique path between x and y in T . Since $T[x_{3n'}, z_{3n'}]$ contains all the edges in P_x and P_z , the length of $T[x_{3n'}, z_{3n'}]$ is at least $6n'$, and hence $|E(P^*)| \geq |E(T[x_{3n'}, z_{3n'}])| \geq 6n'$. Since each of $G - \{x_1, \dots, x_{3n'}\}$ and $G - \{z_1, \dots, z_{3n'}\}$ contains at most $5n'$ vertices, we obtain $V(P^*) \cap \{x_1, \dots, x_{3n'}\} \neq \emptyset$ and $V(P^*) \cap \{z_1, \dots, z_{3n'}\} \neq \emptyset$. This shows that $P^* = T[x_i, z_j]$ for some $i, j \in \{1, 2, \dots, 3n'\}$. Since $T[x_i, z_j]$ is a subpath of $T[x_{3n'}, z_{3n'}]$, P^* must be equal to $T[x_{3n'}, z_{3n'}]$, that is, $\text{diam}(T) = \bar{\ell}_T(x_{3n'}, z_{3n'})$. \square

Thus, intuitively speaking, to keep the diameter and modify a spanning tree in D , we need to replace P_y with a slightly longer path in G' . Moreover, such a path must be a Hamiltonian path of G' . This observation yields the following lemma and completes the proof of theorem 3.

Lemma 18. (G', s', t') is a YES-instance of HAMILTONIAN PATH if and only if (G, d, T_s, T_t) is a YES-instance of RST WITH LARGE DIAMETER.

Proof. We first prove the “if” direction. Suppose that (G, d, T_s, T_t) is a YES-instance. Then there is a reconfiguration sequence $\langle T_s = T_0, T_1, \dots, T_k = T_t \rangle$ between T_s and T_t in which all the spanning trees have diameter at least d . Let T_i be the first spanning tree in the sequence such that T_i is obtained from T_{i-1} by exchanging an edge in D , that is, $E(T_i) \cap E(D) = \{t't_1, t_1t_2, t_2t_3\}$ for all $j \in \{0, 1, \dots, i-1\}$ and




Figure 8: The graph G in the corresponding instance.

$E(T_i) \cap E(D) \neq \{t't_1, t_1t_2, t_2t_3\}$. Note that such i exists, because $E(T_k) \cap E(D) \neq \{t't_1, t_1t_2, t_2t_3\}$. Note also that $\bar{\ell}_{T_i}(t', t_3) = 2$ by the definition of T_i . Then, by Lemma 17, we obtain

$$\begin{aligned}
7n' + 1 &\leq \text{diam}(T_i) \\
&= \bar{\ell}_{T_i}(x_{3n'}, z_{3n'}) \\
&= \bar{\ell}_{T_i}(x_{3n'}, s') + \bar{\ell}_{T_i}(s', t') + \bar{\ell}_{T_i}(t', t_3) + \bar{\ell}_{T_i}(t_3, z_{3n'}) \\
&= 3n' + \bar{\ell}_{T_i}(s', t') + 2 + 3n',
\end{aligned}$$

and hence $\bar{\ell}_{T_i}(s', t') \geq n' - 1$. Since P_y contains only $n' - 2$ edges, all the edges in $T_i[s', t']$ are contained in G' . We thus conclude that $T_i[s', t']$ is a Hamiltonian path between s' and t' in G' , and hence the “if” direction follows.

We now prove the “only-if” direction. Suppose that (G', s', t') is a YES-instance, that is, G' contains a Hamiltonian path P^* between s' and t' . Let e^* be any edge in P^* and e_y be any edge in P_y . We define five spanning trees T_1, T_2, T_3, T_4 , and T_5 in G as follows:

$$\begin{aligned}
E(T_1) &= E(P_x) \cup E(P_y) \cup E(P_z) \cup E(P^* - e^*) \cup \{t't_1, t_1t_2, t_2t_3\}, \\
E(T_2) &= E(P_x) \cup E(P_y - e_y) \cup E(P_z) \cup E(P^*) \cup \{t't_1, t_1t_2, t_2t_3\}, \\
E(T_3) &= E(P_x) \cup E(P_y - e_y) \cup E(P_z) \cup E(P^*) \cup \{t't_2, t_1t_2, t_2t_3\}, \\
E(T_4) &= E(P_x) \cup E(P_y - e_y) \cup E(P_z) \cup E(P^*) \cup \{t't_2, t_1t_2, t_1t_3\}, \\
E(T_5) &= E(P_x) \cup E(P_y) \cup E(P_z) \cup E(P^* - e^*) \cup \{t't_2, t_1t_2, t_1t_3\}.
\end{aligned}$$

We observe that $\langle T_1, T_2, T_3, T_4, T_5 \rangle$ is a reconfiguration sequence from T_1 and T_5 in which all the spanning trees have diameter at least $d = 7n' + 1$. Thus, in order to show that T_s is reconfigurable to T_t , it suffices to show that T_s is reconfigurable to T_1 and T_5 is reconfigurable to T_t . Since $T_s[x_{3n'}, z_{3n'}] = T_1[x_{3n'}, z_{3n'}]$, lemma 5 shows that there is a reconfiguration sequence from T_s to T_1 in which all the spanning trees contain $E(T_s[x_{3n'}, z_{3n'}]) \subseteq E(T_s) \cap E(T_1)$. Therefore, every spanning tree in the sequence has diameter at least d , and hence T_s is reconfigurable to T_1 . Similarly, T_5 is reconfigurable to T_t . By combining them, we have that T_s is reconfigurable to T_t , which completes the proof of the “only-if” direction. \square

6 Small Diameter (Proof of Theorem 4)

In this section, we prove theorem 4, which we restate here.

Theorem 4. RST WITH SMALL DIAMETER is polynomial-time solvable.

After giving some preliminaries for the proof in Section 6.1, we describe a naive algorithm for the problem in Section 6.2, which does not necessarily run in polynomial time. Then, by modifying it, we give a polynomial-time algorithm in Section 6.3. Some technical parts in the proof are deferred to Sections 7 and 8.

Algorithm 2: First algorithm for RST WITH SMALL DIAMETER

Input: A graph G and two spanning trees T_s and T_t in G with diameter at most d .

Output: Is T_s reconfigurable to T_t ?

- 1 Compute $\text{center}(T_s)$ and $\text{center}(T_t)$, and construct \mathcal{G} ;
 - 2 if there is a path between $\text{center}(T_s)$ and $\text{center}(T_t)$ in \mathcal{G} then return YES;
 - 3 else return NO;
-

6.1 Preliminaries for the Proof

Throughout the proof of theorem 4, we fix a positive integer d . For each edge $e \in E$, we denote the middle point of e by p_e . We denote $R(H) := \{p_e \mid e \in E(H)\}$ for a subgraph H of G and let $R := R(G)$. Let \hat{G} be the graph on $V \cup R$ that is obtained from G by subdividing each edge. Then, since $\bar{\ell}_G(u, v) = \frac{1}{2}\bar{\ell}_{\hat{G}}(u, v)$ for $u, v \in V$, we can naturally extend the domain of the distance to $V \cup R$ by setting $\bar{\ell}_G(u, v) := \frac{1}{2}\bar{\ell}_{\hat{G}}(u, v)$ for $u, v \in V \cup R$. We also define $\epsilon_G(v) := \max\{\bar{\ell}_G(v, u) \mid u \in V\}$ for $v \in R$. If no confusion may arise, for $u, v \in V \cup R$, a u - v path in \hat{G} is sometimes called a u - v path in G . We can see that spanning trees with diameter at most d are characterized as follows (see also [8]).

Lemma 19. For any spanning tree T in $G = (V, E)$, $\text{diam}(T) \leq d$ if and only if there exists $r \in V \cup R(T)$ such that $\epsilon_T(r) \leq \frac{d}{2}$.

Proof. To show the “if” part, suppose that there exists $r \in V \cup R(T)$ such that $\epsilon_T(r) \leq \frac{d}{2}$. Then, for any $u, v \in V$, $\bar{\ell}_T(u, v) \leq \bar{\ell}_T(u, r) + \bar{\ell}_T(r, v) \leq 2\epsilon_T(r) \leq d$, which shows that $\text{diam}(T) \leq d$.

To show the “only-if” part, suppose that $\text{diam}(T) \leq d$. Let $d^* := \text{diam}(T)$ and let $u, v \in V$ be a pair of vertices such that $\bar{\ell}_T(u, v) = d^*$. Let $r \in V \cup R(T)$ be the middle point of u and v in T , that is, $\bar{\ell}_T(u, r) = \bar{\ell}_T(r, v) = \frac{d^*}{2}$. Since T is a spanning tree, for any $x \in V$, $\frac{d^*}{2} + \bar{\ell}_T(r, x) = \max\{\bar{\ell}_T(u, x), \bar{\ell}_T(v, x)\} \leq d^*$. This shows that $\bar{\ell}_T(r, x) \leq \frac{d^*}{2}$, that is, $\epsilon_T(r) \leq \frac{d}{2}$. \square

We say that a subgraph Q of G is a *spanning pseudotree* if it is a connected spanning subgraph containing at most one cycle. In other words, a spanning pseudotree is obtained from a spanning tree by adding at most one edge. For brevity, a spanning pseudotree is simply called a *pseudotree*. For a pseudotree Q , let C_Q denote the unique cycle in Q if it exists. We can easily see that, for two spanning trees T_1 and T_2 with diameter at most d , $T_1 \leftrightarrow T_2$ if and only if $T_1 \cup T_2$ forms a pseudotree. For a pseudotree Q , we refer a point $r \in V \cup R(Q)$ as a *center point* of Q if $\epsilon_Q(r) \leq \frac{d}{2}$. Note that a center point is not necessarily unique even if Q is a spanning tree. For a pseudotree Q , let $\text{center}(Q) \subseteq V \cup R(Q)$ be the set of all center points of Q .

6.2 Algorithm Using Center Points: First Attempt

In this subsection, as a first step, we give an algorithm for RST WITH SMALL DIAMETER whose running time is not necessarily polynomial. In the same way as RST WITH LARGE MAXIMUM DEGREE (Section 3), the proposed algorithm is based on testing the reachability in an auxiliary graph \mathcal{G} , which is defined as follows. The vertex set of \mathcal{G} is defined as $V \cup R$, where each vertex v in $V(\mathcal{G})$ corresponds to the set of all the spanning trees containing v as a center point. For any pair u, v of distinct vertices in $V(\mathcal{G})$, there is an edge $uv \in E(\mathcal{G})$ if and only if there is a pseudotree Q with $u, v \in \text{center}(Q)$. As we will see in proposition 22 later, for two spanning trees T_u and T_v having center points u and v , respectively, \mathcal{G} contains a u - v path if and only if T_u and T_v are reconfigurable to each other. Thus, to solve RST WITH SMALL DIAMETER, it is enough to find a path from a center point of T_s to a center point of T_t on \mathcal{G} . See Algorithm 2 for a pseudocode of our algorithm.

To show the correctness of Algorithm 2, we begin with easy but important lemmas.

Lemma 20. Let T_1 and T_2 be spanning trees in G with diameter at most d . If there exists a point $r \in \text{center}(T_1) \cap \text{center}(T_2)$, then T_1 is reconfigurable to T_2 .

Proof. Let T^* be the spanning tree that is obtained by applying the breadth first search from r in G . Here, if $r \in R$ is the middle point of $uv \in E$, then the breadth first search is started from $\{u, v\}$. Since $\bar{\ell}_{T^*}(r, v) \leq \bar{\ell}_{T_1}(r, v) \leq \frac{d}{2}$ for any $v \in V$, the diameter of T^* is at most d . For $v \in V$, let $P_{T^*}(v)$ (resp. $P_{T_1}(v)$) denote the unique path from r to v in T^* (resp. T_1). In T^* , we say that a vertex $u \in V$ is the *parent* of v if $uv \in E(T^*)$ and $\bar{\ell}_{T^*}(r, v) = \bar{\ell}_{T^*}(r, u) + 1$. The parent in T_1 is defined in the same way.

In order to show that T_1 is reconfigurable to T_2 , it suffices to show that T_i is reconfigurable to T^* for $i \in \{1, 2\}$. Suppose that $T_1 \neq T^*$ and let xy be an edge in $E(T^*) \setminus E(T_1)$ that minimizes $\min\{\bar{\ell}_{T^*}(r, x), \bar{\ell}_{T^*}(r, y)\}$. Without loss of generality, we assume that x is the parent of y in T^* . Let $w \in V$ be the parent of y in T_1 and define $T'_1 := T_1 + \{xy\} - \{wy\}$, which is a spanning tree in G . By the choice of xy , we obtain $P_{T_1}(x) = P_{T^*}(x)$, and hence $P_{T'_1}(y) = P_{T^*}(y)$ and $\bar{\ell}_{T'_1}(r, y) = \bar{\ell}_{T^*}(r, y) \leq \bar{\ell}_{T_1}(r, y)$. Since this shows that $\bar{\ell}_{T'_1}(r, v) \leq \bar{\ell}_{T_1}(r, v) \leq \frac{d}{2}$ for any $v \in V$, the diameter of T'_1 is at most d by lemma 19. We observe that replacing T_1 with T'_1 increases $|\{v \in V \mid P_{T_1}(v) = P_{T^*}(v)\}|$ by at least one, because $P_{T_1}(y) \neq P_{T'_1}(y) = P_{T^*}(y)$. Therefore, by applying this procedure at most $|V|$ times, we obtain a reconfiguration sequence from T_1 to T^* . We can also obtain a reconfiguration sequence from T_2 to T^* in the same way. Hence, the statement holds. \square

Lemma 21. *Let $r_1, r_2 \in V \cup R$ (possibly $r_1 = r_2$). There exists a pseudotree Q with $r_1, r_2 \in \text{center}(Q)$ if and only if there exist two spanning trees T_1 and T_2 such that $r_i \in \text{center}(T_i)$ for $i = 1, 2$ and $T_1 \leftrightarrow T_2$ (possibly $T_1 = T_2$).*

Proof. We first consider the “if” part. Suppose that there exist two spanning trees T_1 and T_2 such that $r_i \in \text{center}(T_i)$ for $i = 1, 2$ and $T_1 \leftrightarrow T_2$. Then $Q := T_1 \cup T_2$ is a desired pseudotree as $\epsilon_Q(r_i) \leq \epsilon_{T_i}(r_i) \leq \frac{d}{2}$ for $i = 1, 2$.

We next consider the “only-if” part. Suppose that Q is a pseudotree with $r_1, r_2 \in \text{center}(Q)$. For $i = 1, 2$, let T_i be the spanning tree that is obtained by applying the breadth first search from r_i in Q . Then, we obtain $\epsilon_{T_i}(r_i) = \epsilon_Q(r_i) \leq \frac{d}{2}$. Furthermore, since $|E(T_1) \setminus E(T_2)| \leq |E(Q) \setminus E(T_2)| = 1$, it holds that $T_1 \leftrightarrow T_2$. \square

By these lemmas, we can show the correctness of Algorithm 2.

Proposition 22. *Let T_s and T_t be spanning trees with diameter at most d . Then, T_s is reconfigurable to T_t if and only if \mathcal{G} contains a path from $\text{center}(T_s)$ to $\text{center}(T_t)$.*

Proof. We first show the “only if” part. Suppose that there exists a reconfiguration sequence $\langle T_s = T_0, T_1, \dots, T_k = T_t \rangle$ from T_s to T_t , where T_i is a spanning tree of diameter at most d for any $i \in \{0, 1, \dots, k\}$ and $T_i \leftrightarrow T_{i+1}$ for any $i \in \{0, 1, \dots, k-1\}$. Let r_i be a center point of T_i , where its existence is guaranteed by lemma 19. For $i \in \{0, 1, \dots, k-1\}$, by lemma 21, there exists a pseudotree Q_i having both r_i and r_{i+1} as center points. This means that either $r_i = r_{i+1}$ or \mathcal{G} contains an edge $r_i r_{i+1}$. Since $r_0 \in \text{center}(T_s)$ and $r_k \in \text{center}(T_t)$, \mathcal{G} contains a path from $\text{center}(T_s)$ to $\text{center}(T_t)$.

To show the “if” part, suppose that \mathcal{G} contains a path (r_0, r_1, \dots, r_k) from $\text{center}(T_s)$ to $\text{center}(T_t)$. For $i \in \{0, 1, \dots, k-1\}$, since $r_i r_{i+1} \in E(\mathcal{G})$ implies the existence of a pseudotree Q_i with $r_i, r_{i+1} \in \text{center}(Q_i)$, lemma 21 shows that there exist two spanning trees T_i^+ and T_{i+1}^- such that $r_i \in \text{center}(T_i^+)$, $r_{i+1} \in \text{center}(T_{i+1}^-)$, and $T_i^+ \leftrightarrow T_{i+1}^-$. Let $T_0^- := T_s$ and $T_k^+ := T_t$. Then, for $i \in \{0, 1, \dots, k\}$, since T_i^- and T_i^+ share r_i as a center point, T_i^- is reconfigurable to T_i^+ by lemma 20. This together with $T_i^+ \leftrightarrow T_{i+1}^-$ shows that T_s is reconfigurable to T_t . \square

Although this proposition shows the correctness of Algorithm 2, it does not imply a polynomial-time algorithm for RST WITH SMALL DIAMETER, because it is not easy to construct \mathcal{G} efficiently. Indeed, for $u, v \in V(\mathcal{G})$, we do not know how to decide whether $uv \in E(\mathcal{G})$ or not in polynomial time. To avoid this problem, we efficiently construct a subgraph \mathcal{G}' of \mathcal{G} such that the reachability of \mathcal{G}' is equal to that of \mathcal{G} , which is a key ingredient of our algorithm and discussed in the next subsection.

6.3 Modified Algorithm

In this subsection, we give a polynomial-time algorithm for RST WITH SMALL DIAMETER. In our algorithm, it is important to uniquely determine a shortest path between two points. To achieve this, we use a *perturbation technique* (see e.g., [3]).

For each edge e in G , we give a unique index $i(e) \in \{1, 2, \dots, |E|\}$ to e . For $j \in \{1, 2, \dots, |E|\}$, let $\chi_j \in \mathbb{R}^{|E|}$ be the unit vector such that the j th coordinate is one and the other coordinates are zero. For $e \in E$, define $\ell(e) := (1, \chi_{i(e)}) \in \mathbb{R} \times \mathbb{R}^{|E|}$. For two vectors $x, y \in \mathbb{R}^k$, we denote $x < y$ if x is lexicographically smaller than y . For two paths P_1 and P_2 in G , we say that P_1 is *shorter than* P_2 if $\ell(P_1) := \sum_{e \in E(P_1)} \ell(e)$ is lexicographically smaller than $\ell(P_2) := \sum_{e \in E(P_2)} \ell(e)$. Since the first coordinate of $\ell(P_i)$ is $|E(P_i)|$ for $i = 1, 2$, if $|E(P_1)| < |E(P_2)|$, then P_1 is shorter than P_2 . When $|E(P_1)| = |E(P_2)|$, we use the other coordinates to break ties. For $u, v \in V$, we define $\ell_G(u, v) := \min_P \sum_{e \in E(P)} \ell(e)$, where the minimum is taken over all the u - v paths P . Since $P_1 \neq P_2$ implies that $\ell(P_1) \neq \ell(P_2)$, the *shortest path* between two vertices is uniquely determined. We note that the unique shortest paths between two given vertices can be computed by using a standard shortest path algorithm. The running time is increased by the perturbation, but it is still polynomial.

For an edge $e = uv \in E$ of length $\ell(e) \in \mathbb{R} \times \mathbb{R}^{|E|}$, we regard e as a curve connecting u and v . An interior point p on e is represented by a triplet (u, v, α) with $\alpha \in \mathbb{R} \times \mathbb{R}^{|E|}$ such that $\mathbf{0} \leq \alpha \leq \ell(e)$, where \leq means the lexicographical order. Here, α represents the length between u and p , and hence (u, v, α) and $(v, u, \ell(e) - \alpha)$ represent the same point. For two points $p_1 = (u_1, v_1, \alpha_1)$ and $p_2 = (u_2, v_2, \alpha_2)$ in G , consider a curve C connecting p_1 and p_2 that consists of a u_1 - u_2 path P , a curve in u_1v_1 between u_1 and p_1 , and a curve in u_2v_2 between u_2 and p_2 . Such a curve C is called a p_1 - p_2 *path* in G , and its length is defined as $\ell(C) := \sum_{e \in E(P)} \ell(e) + \alpha_1 + \alpha_2$.

For a point $r \in V \cup R$, the *shortest path tree from* r is the spanning tree in G that contains the unique shortest r - v path for any $v \in V$. For a pseudotree Q and for two points x and y on Q , let $Q[x, y]$ denote the shortest x - y path in Q , where we use this notation only when the shortest x - y path is uniquely determined. For $\alpha \in \mathbb{R} \times \mathbb{R}^{|E|}$, let $\bar{\alpha}$ denote the first coordinate of α , that is, $\bar{\alpha}$ is the length before the perturbation.

We denote $r_1 \xrightarrow{Q} r_2$ if Q is a pseudotree and $r_1, r_2 \in \text{center}(Q)$ with $r_1 \neq r_2$. For any pseudotree Q and any points r_1 and r_2 with $r_1 \xrightarrow{Q} r_2$, we say that a triplet (r_1, r_2, Q) is *good* if

1. $\text{label}_{r_1, r_2, Q}(v) \leq \text{label}_{r_1, r_2, Q}(u) + \ell(uv)$ for any $uv \in E$, and
2. C_Q contains both r_1 and r_2 if C_Q exists,

where $\text{label}_{r_1, r_2, Q}(v) := \max\{\ell_Q(r_1, v), \ell_Q(r_2, v)\}$. Roughly speaking, the first condition means that $\text{label}_{r_1, r_2, Q}(v)$ can be seen as the distance from a certain point to v in an auxiliary graph. If r_1 and r_2 are clear from the context, $\text{label}_{r_1, r_2, Q}(v)$ is simply denoted by $\text{label}_Q(v)$. We define the graph \mathcal{G}' as follows: $V(\mathcal{G}') = V \cup R$ and \mathcal{G}' contains an edge r_1r_2 if and only if there is a pseudotree Q such that $r_1 \xrightarrow{Q} r_2$ and (r_1, r_2, Q) is good. Clearly, \mathcal{G}' is a subgraph of \mathcal{G} .

The following theorem shows that we can determine whether $r_1r_2 \in E(\mathcal{G}')$ or not in polynomial time, whose proof is given in Section 7.

Theorem 23. *Let r_1 and r_2 be points in $V \cup R$ with $r_1 \neq r_2$. We can find in polynomial time a pseudotree Q such that $r_1 \xrightarrow{Q} r_2$ and (r_1, r_2, Q) is good if it exists.*

The next theorem shows that the reachability of \mathcal{G}' is equal to that of \mathcal{G} , which is a key property of \mathcal{G}' . A proof is given in Section 8.

Theorem 24. *For any $r_1, r_2 \in V \cup R$ with $r_1r_2 \in E(\mathcal{G})$, \mathcal{G}' contains an r_1 - r_2 path.*

We are now ready to prove theorem 4. By proposition 22 and theorem 24, two spanning trees T_s and T_t are reconfigurable to each other if and only if \mathcal{G}' contains a path from $\text{center}(T_s)$ to $\text{center}(T_t)$. Since we can construct \mathcal{G}' in polynomial time by theorem 23, this can be tested in polynomial time. Therefore, RST WITH SMALL DIAMETER can be solved in polynomial time, which completes the proof of theorem 4. A pseudocode of our algorithm is given in Algorithm 3. Note that all the proofs are constructive, and hence we can find a desired reconfiguration sequence from T_s to T_t in polynomial time if it exists.

Algorithm 3: Modified algorithm for RST WITH SMALL DIAMETER

Input: A graph G and two spanning trees T_s and T_t in G with diameter at most d .
Output: Is T_s reconfigurable to T_t ?

- 1 Compute $\text{center}(T_s)$ and $\text{center}(T_t)$, and construct $\mathcal{G}' = (V \cup R, \emptyset)$;
- 2 **for** $r_1, r_2 \in V \cup R$ with $r_1 \neq r_2$ **do**
- 3 **if** there is a pseudotree Q such that $r_1 \xleftrightarrow{Q} r_2$ and (r_1, r_2, Q) is good **then**
- 4 Add an edge $r_1 r_2$ to \mathcal{G}' ;
- 5 **end**
- 6 **end**
- 7 **if** there is a path between $\text{center}(T_s)$ and $\text{center}(T_t)$ in \mathcal{G}' **then return YES**;
- 8 **else return NO**;




Figure 9: Construction of G^+ .

7 Proof of Theorem 23

In this section, we give an algorithm for finding a pseudotree Q such that $r_1 \xleftrightarrow{Q} r_2$ and (r_1, r_2, Q) is good (if it exists), and prove Theorem 23. Our algorithm consists of two procedures: one finds a desired spanning tree Q and the other finds a desired pseudotree Q with a cycle. We describe these two procedures in Sections 7.1 and 7.2, respectively. In what follows in this section, we fix $r_1, r_2 \in V \cup R$.

7.1 Finding a spanning tree Q

We first consider the case when Q is a spanning tree. Suppose that Q is a spanning tree such that $r_1 \xleftrightarrow{Q} r_2$ and (r_1, r_2, Q) is good. Let $e = v_1 v_2 \in E$ be the unique edge in Q such that $\ell_Q(r_1, v_1) < \ell_Q(r_2, v_1)$ and $\ell_Q(r_1, v_2) > \ell_Q(r_2, v_2)$. That is, e is the edge containing the middle point of r_1 and r_2 . Note that such e exists, because $\ell_Q(r_1, v) \neq \ell_Q(r_2, v)$ for each $v \in V$ as ℓ is the perturbed length. Define G^+ as the graph obtained from G by adding a new vertex r together with two new edges $e_1 = r v_1$ and $e_2 = r v_2$ (Figure 9). Set $\ell(e_1) := \ell_Q(r_2, v_1)$ and $\ell(e_2) := \ell_Q(r_1, v_2)$. We now show the following claim.

Claim 25. For each $v \in V$, it holds that $\text{label}_Q(v) = \ell_{G^+}(r, v)$. Furthermore, $Q - e + \{e_1, e_2\}$ is the shortest path tree starting from r in G^+ .

Proof. Let Q_1 and Q_2 be the connected components of $Q - e$ such that $v_i \in V(Q_i)$ for $i = 1, 2$. Then, $\text{label}_Q(v) = \ell_Q(r_2, v_1) + \ell_{Q-e}(v_1, v) = \ell(e_1) + \ell_{Q-e}(v_1, v)$ for $v \in V(Q_1)$ and $\text{label}_Q(v) = \ell_Q(r_1, v_2) + \ell_{Q-e}(v_2, v) = \ell(e_2) + \ell_{Q-e}(v_2, v)$ for $v \in V(Q_2)$. Therefore, for each $v \in V$, $Q - e + \{e_1, e_2\}$ contains an r - v path P_v whose length is $\text{label}_Q(v)$. Furthermore, since (r_1, r_2, Q) is good, $\text{label}_Q(v) \leq \text{label}_Q(u) + \ell(uv)$ holds for any $uv \in E$, by the correctness of the Bellman-Ford method (see e.g. [18, Section 8.3], [15, Section 7.1]), $\text{label}_Q(v)$ is equal to the shortest path length from r to v in G^+ . That is P_v is the unique shortest r - v path in G^+ , which shows the claim. \square

Since a subpath of a shortest path is also a shortest path, this claim implies that $Q[v_i, r_i]$ is the unique

shortest v_i - r_i path in G for $i = 1, 2$. Therefore, we obtain

$$\ell(e_1) = \ell_Q(r_2, v_1) = \begin{cases} \frac{1}{2}\ell(e) & \text{if } r_2 = p_e, \\ \ell_G(r_2, v_2) + \ell(e) & \text{otherwise,} \end{cases} \quad (7)$$

and

$$\ell(e_2) = \ell_Q(r_1, v_1) = \begin{cases} \frac{1}{2}\ell(e) & \text{if } r_1 = p_e, \\ \ell_G(r_1, v_1) + \ell(e) & \text{otherwise.} \end{cases} \quad (8)$$

Note that the right-hand sides of (7) and (8) can be computed without Q if we have v_1 and v_2 in hand.

Therefore, the following algorithm correctly finds a desired spanning tree Q : guess two vertices v_1 and v_2 in V , construct G^+ where $\ell(e_1)$ and $\ell(e_2)$ are defined as (7) and (8), compute the shortest path tree T from r in G^+ , define $Q := T - \{e_1, e_2\} + e$, and check whether Q satisfies the constraints or not. See Algorithm 4 for a pseudocode of our algorithm. Since the number of choices of v_1 and v_2 is $O(|E|)$, this algorithm runs in polynomial time.

Algorithm 4: Algorithm for finding a spanning tree Q

Input: A graph G with two distinct points r_1 and r_2 in $V \cup R$.

Output: A spanning tree Q such that $r_1 \xleftrightarrow{Q} r_2$ and (r_1, r_2, Q) is good.

```

1 for  $v_1, v_2 \in V$  with  $e = v_1v_2 \in E$  do
2   Construct  $G^+$ ;
3   Define  $\ell(e_1)$  and  $\ell(e_2)$  as (7) and (8);
4   Compute the shortest path tree  $T$  from  $r$  in  $G^+$ ;
5    $Q := T - \{e_1, e_2\} + e$ ;
6   if  $Q$  is a spanning tree such that  $r_1 \xleftrightarrow{Q} r_2$  and  $(r_1, r_2, Q)$  is good then return  $Q$ ;
7 end
8 return no solution exists;

```

7.2 Finding a pseudotree Q with a cycle

We next consider the case when Q contains a cycle. Suppose that Q is a pseudotree with a cycle C_Q such that $r_1 \xleftrightarrow{Q} r_2$ and (r_1, r_2, Q) is good. Since r_1 and r_2 are on C_Q , there exist two edges $e = v_1v_2$ and $e' = v'_1v'_2$ in $E(C_Q)$ such that $\ell_Q(r_1, v_1) < \ell_Q(r_2, v_1)$, $\ell_Q(r_1, v_2) > \ell_Q(r_2, v_2)$, $\ell_Q(r_1, v'_1) < \ell_Q(r_2, v'_1)$, and $\ell_Q(r_1, v'_2) > \ell_Q(r_2, v'_2)$. That is, each of e and e' is the edge containing the middle point of an r_1 - r_2 path, where we note that Q contains two r_1 - r_2 paths. Let Q_1 and Q_2 be the connected components of $Q - \{e, e'\}$ such that $v_i, v'_i \in V(Q_i)$ for $i = 1, 2$. Define G^+ as the graph obtained from G by adding a new vertex r together with four new edges $e_1 = rv_1$, $e_2 = rv_2$, $e'_1 = rv'_1$, and $e'_2 = rv'_2$ (Figure 10). Set

$$\ell(e_1) := \begin{cases} \frac{1}{2}\ell(e) & \text{if } r_2 = p_e, \\ \ell_{C_Q}(r_2, v_2) + \ell(e) & \text{otherwise,} \end{cases} \quad (9)$$

$$\ell(e_2) := \begin{cases} \frac{1}{2}\ell(e) & \text{if } r_1 = p_e, \\ \ell_{C_Q}(r_1, v_1) + \ell(e) & \text{otherwise,} \end{cases} \quad (10)$$

$$\ell(e'_1) := \begin{cases} \frac{1}{2}\ell(e') & \text{if } r_2 = p_{e'}, \\ \ell_{C_Q}(r_2, v'_2) + \ell(e') & \text{otherwise,} \end{cases} \quad (11)$$

$$\ell(e'_2) := \begin{cases} \frac{1}{2}\ell(e') & \text{if } r_1 = p_{e'}, \\ \ell_{C_Q}(r_1, v'_1) + \ell(e') & \text{otherwise.} \end{cases} \quad (12)$$




Figure 10: Construction of G^+ .




Figure 11: Paths $P_{u_1}, P_{u'_1}, P_{u_2},$ and $P_{u'_2}$.

Note that $\ell(e_1)$ (resp. $\ell(e'_1), \ell(e_2), \ell(e'_2)$) is the length of an r_2 - v_1 path (resp. an r_2 - v'_1 path, an r_1 - v_2 path, an r_1 - v'_2 path) in Q . We now show the following claim.

Claim 26. *For each $v \in V$, it holds that $\text{label}_Q(v) = \ell_{G^+}(r, v)$. Furthermore, $Q - \{e, e'\} + \{e_1, e_2, e'_1, e'_2\}$ contains the shortest path tree starting from r in G^+ .*

Proof. Since $\text{label}_Q(v) = \ell_Q(r_2, v) = \min\{\ell(e_1) + \ell_{Q_1}(v_1, v), \ell(e'_1) + \ell_{Q_1}(v'_1, v)\}$ for $v \in V(Q_1)$ and $\text{label}_Q(v) = \ell_Q(r_1, v) = \min\{\ell(e_2) + \ell_{Q_2}(v_2, v), \ell(e'_2) + \ell(e') + \ell_{Q_2}(v'_2, v)\}$ for $v \in V(Q_2)$, for each $v \in V$, $Q - \{e, e'\} + \{e_1, e_2, e'_1, e'_2\}$ contains an r - v path P_v whose length is $\text{label}_Q(v)$. Furthermore, since $\text{label}_Q(v) \leq \text{label}_Q(u) + \ell(uv)$ for any $uv \in E$, $\text{label}_Q(v)$ is equal to the shortest path length from r to v in G^+ , that is, P_v is the unique shortest r - v path in G^+ . Since the shortest path tree is the union of P_v 's, it is contained in $Q - \{e, e'\} + \{e_1, e_2, e'_1, e'_2\}$. \square

This claim shows that there exist two edges $f_1 = u_1u'_1$ and $f_2 = u_2u'_2$ in $E(Q) \cup \{e_1, e_2, e'_1, e'_2\}$ such that $Q + \{e_1, e_2, e'_1, e'_2\} - \{e, e', f_1, f_2\}$ is the shortest path tree starting from r in G^+ . Without loss of generality, we may assume that $f_i \in E(Q_i) \cup \{e_i, e'_i\}$ for $i = 1, 2$, because $E(Q_i) \cup \{e_i, e'_i\}$ contains a cycle. Furthermore, for $i = 1, 2$, we may assume that $v_i, u_i, u'_i,$ and v'_i appear in this order along C_Q if $f_i \in E(Q_i)$. For $i = 1, 2$, we observe the following.

- If $f_i \in \{e_i, e'_i\}$, then $Q_i[v_i, v'_i]$ is the shortest v_i - v'_i path in G , because it is a subpath of P_{v_i} or $P_{v'_i}$, where P_v is the unique shortest r - v path in G^+ for $v \in V$.
- If $f_i \in E(Q_i)$, then $Q_i[v_i, u_i]$ is the shortest v_i - u_i path in G because it is a subpath of P_{u_i} , and $Q_i[v'_i, u'_i]$ is the shortest v'_i - u'_i path in G because it is a subpath of $P_{u'_i}$ (Figure 11).

We are now ready to describe our algorithm. In our algorithm, we first guess four vertices $v_1, v_2, v'_1,$ and v'_2 in V with $e = v_1v_2 \in E$ and $e' = v'_1v'_2 \in E$, construct G^+ , and guess four vertices $u_1, u'_1, u_2,$ and u'_2 in V

with $f_1 = u_1 u'_1 \in E \cup \{e_1, e'_1\}$ and $f_2 = u_2 u'_2 \in E \cup \{e_2, e'_2\}$. For $i = 1, 2$, let J_i be the v_i - v'_i walk defined by

$$J_i := \begin{cases} P_{v_i, v'_i}^* & \text{if } f_i \in \{e_i, e'_i\}, \\ P_{v_i, u_i}^* \circ \{f_i\} \circ P_{u'_i, v'_i}^* & \text{if } f_i \in E, \end{cases} \quad (13)$$

where $P_{x,y}^*$ denotes the shortest x - y path in G for $x, y \in V$ and \circ denotes the concatenation of walks. Define a closed walk C as

$$C := J_1 \circ \{e'\} \circ \overline{J_2} \circ \{e\}, \quad (14)$$

where $\overline{J_2}$ is the reverse walk of J_2 . Note that if a desired pseudotree Q exists and if $v_1, v_2, v'_1, v'_2, u_1, u'_1, u_2$, and u'_2 are guessed correctly, then $C = C_Q$ holds by the above arguments. Therefore, it suffices to consider the case when C is a cycle containing r_1 and r_2 . We define $\ell(e_1), \ell(e_2), \ell(e'_1)$, and $\ell(e'_2)$ as (9)–(12) in which C_Q is replaced with C . Finally, we compute the shortest path tree T from r in G^+ , define $Q := (T + \{e, e', f_1, f_2\}) - \{e_1, e_2, e'_1, e'_2\}$, and check whether Q satisfies the constraints or not. See Algorithm 5 for a pseudocode of our algorithm.

By the above arguments, this algorithm finds a desired pseudotree Q if $v_1, v_2, v'_1, v'_2, u_1, u'_1, u_2$, and u'_2 are guessed correctly, which shows the correctness of the algorithm. Since the number of choices of these vertices is $O(|E|^4)$, this algorithm runs in polynomial time.

Algorithm 5: Algorithm for finding a pseudotree Q

Input: A graph G with two distinct points r_1 and r_2 in $V \cup R$.

Output: A pseudotree Q with a cycle such that $r_1 \xleftrightarrow{Q} r_2$ and (r_1, r_2, Q) is good.

```

1 for  $v_1, v_2, v'_1, v'_2 \in V$  with  $e = v_1 v_2 \in E$  and  $e' = v'_1 v'_2 \in E$  do
2   Construct  $G^+$ ;
3   for  $u_1, u'_1, u_2, u'_2 \in V$  with  $f_1 = u_1 u'_1 \in E \cup \{e_1, e'_1\}$  and  $f_2 = u_2 u'_2 \in E \cup \{e_2, e'_2\}$  do
4     Define  $J_1, J_2$ , and  $C$  as (13) and (14);
5     if  $C$  is a cycle containing  $r_1$  and  $r_2$  then
6       Define  $\ell(e_1), \ell(e_2), \ell(e'_1)$ , and  $\ell(e'_2)$  as (9)–(12) in which  $C_Q$  is replaced with  $C$ ;
7       Compute the shortest path tree  $T$  from  $r$  in  $G^+$ ;
8        $Q := (T + \{e, e', f_1, f_2\}) - \{e_1, e_2, e'_1, e'_2\}$ ;
9       if  $Q$  is a pseudotree with a cycle such that  $r_1 \xleftrightarrow{Q} r_2$  and  $(r_1, r_2, Q)$  is good then return  $Q$ ;
10    end
11  end
12 end
13 return no solution exists;
```

8 Proof of Theorem 24

For points $r_1, r_2 \in V \cup R$ and a pseudotree Q with $r_1 \xleftrightarrow{Q} r_2$, we define $f(r_1, r_2, Q) := (\ell(C_Q), L(r_1, r_2, Q))$, where $L(r_1, r_2, Q) = \sum_{v \in V} \text{label}_{r_1, r_2, Q}(v)$. Here, we denote $\ell(C_Q) := \sum_{e \in E(C_Q)} \ell(e)$ if Q contains a cycle and $\ell(C_Q) := \mathbf{0}$ otherwise. For two triplets (r_1, r_2, Q) and (r'_1, r'_2, Q') , we compare $f(r_1, r_2, Q)$ and $f(r'_1, r'_2, Q')$ lexicographically, that is, we denote $f(r_1, r_2, Q) < f(r'_1, r'_2, Q')$ if and only if either $\ell(C_Q) < \ell(C_{Q'})$ or $\ell(C_Q) = \ell(C_{Q'})$ and $L(r_1, r_2, Q) < L(r'_1, r'_2, Q')$.

To show Theorem 24, we apply the induction on $f(r_1, r_2, Q)$ by using the following proposition, which roughly says that we can reduce to the case with smaller $f(\cdot)$ if (r_1, r_2, Q) is not good.

Proposition 27. *Let $r_1, r_2 \in V \cup R$ be points and Q be a pseudotree with $r_1 \xleftrightarrow{Q} r_2$. If (r_1, r_2, Q) is not good, then one of the followings holds:*




Figure 12: Case when $r_1 \notin V(C_Q)$ and $r_2 \notin V(Q_{r_1})$.

- B1.** there exists a pseudotree Q' such that $r_1 \xleftrightarrow{Q'} r_2$ with $f(r_1, r_2, Q') < f(r_1, r_2, Q)$ or
B2. there exist a point $r_0 \in V \cup R$ and pseudotrees Q_1, Q_2 such that

- $r_1 \xleftrightarrow{Q_1} r_0$ with $f(r_1, r_0, Q_1) < f(r_1, r_2, Q)$ and
- $r_0 \xleftrightarrow{Q_2} r_2$ with $f(r_0, r_2, Q_2) < f(r_1, r_2, Q)$.

Proof. Suppose that (r_1, r_2, Q) is not good.

We first consider the case when Q is a spanning tree. Let r be the middle point of r_1 and r_2 with respect to $\ell(\cdot)$. Then, $\text{label}_Q(v) = \ell_Q(r_1, r) + \ell_Q(r, v)$ for $v \in V$. Since (r_1, r_2, Q) is not good, there exists an edge $uv \in E$ such that $\text{label}_Q(u) + \ell(uv) < \text{label}_Q(v)$. Let e be the unique edge on $Q[r, v]$ that is incident to v . Then, $Q' := Q - \{e\} + \{uv\}$ is a spanning tree such that $\ell_{Q'}(r, x) \leq \ell_Q(r, x)$ for any $x \in V$. Therefore,

$$\ell_{Q'}(r_i, x) \leq \ell_{Q'}(r_i, r) + \ell_{Q'}(r, x) \leq \ell_Q(r_i, r) + \ell_Q(r, x) = \max\{\ell_Q(r_1, x), \ell_Q(r_2, x)\}$$

for any $x \in V$ and for $i = 1, 2$. Since $r_1 \xleftrightarrow{Q} r_2$ implies that $\bar{\ell}_Q(r_i, x) \leq \frac{d}{2}$, we obtain $\bar{\ell}_{Q'}(r_i, x) \leq \frac{d}{2}$ for $x \in V$ and for $i = 1, 2$. This shows that $r_1 \xleftrightarrow{Q'} r_2$. Furthermore, since $\text{label}_{Q'}(x) \leq \text{label}_Q(x)$ for any $x \in V$ and $\text{label}_{Q'}(v) < \text{label}_Q(v)$, we have that $L(r_1, r_2, Q') < L(r_1, r_2, Q)$ and $f(r_1, r_2, Q') < f(r_1, r_2, Q)$. Thus, Q' satisfies the conditions in **B1**.

We next consider the case when Q contains a cycle C_Q and at least one of r_1 and r_2 is not contained in C_Q . Without loss of generality, we assume that $r_1 \notin V(C_Q)$. Define T as the shortest path tree in Q starting from r_1 . Let Q_{r_1} be the connected component of $Q - E(C_Q)$ that contains r_1 . If $r_2 \in V(Q_{r_1})$, then $\text{label}_T(x) = \text{label}_Q(x)$ for any $x \in V$ and $f(r_1, r_2, T) < f(r_1, r_2, Q)$, and hence $Q' := T$ satisfies the conditions in **B1**. Otherwise, let r_0 be the unique vertex in $V(Q_{r_1}) \cap V(C_Q)$, where we note that $r_0 \neq r_1$ and $r_0 \neq r_2$ (Figure 12). Since $\ell_Q(r_0, x) < \ell_Q(r_2, x)$ for $x \in V(Q_{r_1})$ and $\ell_Q(r_0, x) < \ell_Q(r_1, x)$ for $x \in V \setminus V(Q_{r_1})$, it holds that $\ell_Q(r_0, x) < \max\{\ell_Q(r_1, x), \ell_Q(r_2, x)\} = \text{label}_Q(x)$ for $x \in V(Q_{r_1})$. Therefore, $\bar{\ell}_Q(r_0, x) \leq \frac{d}{2}$ for any $x \in V$, and hence $r_1 \xleftrightarrow{Q} r_0$ and $r_0 \xleftrightarrow{Q} r_2$. Furthermore, since $L(r_1, r_0, Q) < L(r_1, r_2, Q)$ and $L(r_0, r_2, Q) < L(r_1, r_2, Q)$, we obtain $f(r_1, r_0, Q) < f(r_1, r_2, Q)$ and $f(r_0, r_2, Q) < f(r_1, r_2, Q)$. This shows that $Q_1 := Q_2 := Q$ and r_0 satisfy the conditions in **B2**.

In what follows, we consider the case when Q contains a cycle C_Q and both r_1 and r_2 are contained in C_Q . Since (r_1, r_2, Q) is not good, there exists an edge $uv \in E(Q)$ such that $\text{label}_Q(v) > \text{label}_Q(u) + \ell(uv)$. If $v \notin V(C_Q)$, then $Q' := Q - \{e\} + \{uv\}$ satisfies the conditions in **B1**, where e is the unique edge on $Q[r_1, v]$ that is incident to v . This is because $C_Q = C_{Q'}$ and $\text{label}_{Q'}(x) < \text{label}_Q(x)$ for any $x \in V$ with $v \in V(Q[r_1, x])$ (see Figure 13). Thus, it suffices to consider the case when $v \in V(C_Q)$. For $i = 1, 2$, let g_i be the farthest point on C_Q from r_i , that is, $\ell_Q(r_i, g_i) = \frac{1}{2}\ell(C_Q)$, and let $e_i \in E(Q)$ be the edge containing g_i . Then, $Q - e_i$ is the shortest path tree in Q starting from r_i . We may assume that $e_1 \neq e_2$, since otherwise $Q' := Q - e_1$ satisfies the conditions in **B1**. Let g_1 and g_2 be the points in C_Q such that $\ell_Q(g_i, r_1) = \ell_Q(g_i, r_2) =: \ell_i$ for $i = 1, 2$ (Figure 14). Then, we observe that

$$\text{label}_Q(x) = \max\{\ell_Q(r_1, x), \ell_Q(r_2, x)\} = \min\{\ell_1 + \ell_Q(g_1, x), \ell_2 + \ell_Q(g_2, x)\} \quad (15)$$




Figure 13: Case when $v \notin V(C_Q)$.




Figure 14: Positions of q_i , e_i , and g_i .




Figure 15: Proof of Claim 28.

for $x \in V$. Let w be the vertex in $V(Q_u) \cap V(C_Q)$, where Q_u is the connected component of $Q - E(C_Q)$ containing u . Note that if u is in $V(C_Q)$, then $w = u$. We now show the following claim, which is useful in our case analysis.

Claim 28. *Suppose that u, v, e_1 , and e_2 are defined as above. If $Q' := Q - \{e_1, e_2\} + \{uv\}$ is connected, then **B1** or **B2** holds.*

Proof.

Let $Q_1 := Q - \{e_1\} + \{uv\}$ and $Q_2 := Q - \{e_2\} + \{uv\}$. It is obvious that Q_1 and Q_2 are pseudotrees.

By changing the roles of g_1 and g_2 if necessary, we may assume that v is on the g_1 - g_2 path in Q that contains g_1 . Since $Q' = Q - \{e_1, e_2\} + \{uv\}$ is connected, w is on the g_1 - g_2 path in Q that contains g_2 (Figure 15).

We first consider the lengths of C_{Q_1} and C_{Q_2} . Since $\ell_1 + \ell_Q(g_1, v) = \text{label}_Q(v) \leq \frac{1}{2}\ell(C_Q) = \ell_1 + \ell_2$, it holds that $\ell_Q(g_1, v) \leq \ell_2$. Since $\text{label}_Q(u) = \ell_2 + \ell_Q(g_2, u)$, $\text{label}_Q(v) \leq \ell_1 + \ell_2$, and $\text{label}_Q(u) + \ell(uv) < \text{label}_Q(v)$, it holds that $\ell_Q(g_2, u) + \ell(uv) < \ell_1$. By combining these two inequalities, we obtain $\ell_Q(g_2, u) + \ell(uv) + \ell_Q(g_1, v) < \ell_1 + \ell_2$. Let P be the unique g_1 - g_2 path in $Q - \{e_1, e_2\} + \{uv\}$. Then, this inequality shows that $\ell(P) < \ell_1 + \ell_2$. Therefore, $\ell(C_{Q_i}) \leq \ell_1 + \ell_2 + \ell(P) < 2(\ell_1 + \ell_2) = \ell(C_Q)$ for $i = 1, 2$, that is, the lengths of C_{Q_1} and C_{Q_2} are shorter than that of C_Q .




Figure 16: w is on $Q[g_1, v]$.




Figure 17: w is on $Q[q_1, g_1]$.




Figure 18: w is on $Q[v, g_2]$.

Recall that $\bar{\ell}_1$ (resp. $\bar{\ell}_2$) denotes the first coordinate of ℓ_1 (resp. ℓ_2), which corresponds to the length before the perturbation. If $\ell(P) \geq \bar{\ell}_1$, then let r_0 be the point on P such that $\bar{\ell}_P(r_0, g_1) = \bar{\ell}_1$. Then, r_0 is in $V \cup R$, because $\bar{\ell}_P(r_0, v)$ is equal to $\bar{\ell}_Q(r_1, v)$ or $\bar{\ell}_Q(r_2, v)$, which is half-integral. If $\ell(P) < \bar{\ell}_1$, then let r_0 be the point on P that is closest to g_2 subject to $r_0 \in V \cup R$. This construction shows that there exists a point $r_0 \in V \cup R$ on P such that $\bar{\ell}_P(r_0, g_1) \leq \bar{\ell}_1$ and $\bar{\ell}_P(r_0, g_2) \leq \bar{\ell}_2$. Let $V_1 \subseteq V$ (resp. $V_2 \subseteq V$) be the vertex set of the connected component of $Q - \{e_1, e_2\}$ that contains v (resp. w). For $i = 1, 2$ and for $x \in V_1$, we obtain

$$\begin{aligned} \bar{\ell}_{Q_i}(r_0, x) &\leq \bar{\ell}_P(r_0, g_1) + \bar{\ell}_Q(g_1, x) \leq \bar{\ell}_1 + \bar{\ell}_Q(g_1, x) \\ &= \max\{\bar{\ell}_Q(r_1, x), \bar{\ell}_Q(r_2, x)\} \leq \frac{d}{2}, \end{aligned}$$

where we use (15) in the equality. Similarly, for $i = 1, 2$ and for $x \in V_2$, it holds that

$$\begin{aligned} \bar{\ell}_{Q_i}(r_0, x) &\leq \bar{\ell}_P(r_0, g_2) + \bar{\ell}_Q(g_2, x) \leq \bar{\ell}_2 + \bar{\ell}_Q(g_2, x) \\ &= \max\{\bar{\ell}_Q(r_1, x), \bar{\ell}_Q(r_2, x)\} \leq \frac{d}{2}. \end{aligned}$$

The above inequalities show that $r_0 \in \text{center}(Q_i)$ for $i = 1, 2$. Since $r_i \in \text{center}(Q_i)$ for $i = 1, 2$, we obtain $r_1 \xrightarrow{Q_1} r_0$ if $r_0 \neq r_1$ and $r_0 \xrightarrow{Q_2} r_2$ if $r_0 \neq r_2$.

Since the lengths of C_{Q_1} and C_{Q_2} are shorter than that of C_Q , we have that $f(r_1, r_0, Q_1) < f(r_1, r_2, Q)$ and $f(r_0, r_2, Q_2) < f(r_1, r_2, Q)$. Therefore, if $r_0 \notin \{r_1, r_2\}$, then r_0, Q_1 , and Q_2 satisfy the conditions in **B2**. If $r_0 = r_1$ (resp. $r_0 = r_2$), then Q_2 (resp. Q_1) satisfies the conditions in **B1**. This completes the proof. \diamond

In what follows, by changing the roles of g_1 and g_2 if necessary, we may assume that $\ell_Q(g_1, r_1) > \ell_Q(g_2, r_1)$. Then, g_1, q_2, r_1, g_2, r_2 , and q_1 appear in this order along C_Q . Let $V_2 \subseteq V$ be the vertex set of the connected component of $G - \{e_1, e_2\}$ that contains r_1 and r_2 . Let $V_1 := V \setminus V_2$. Then, we see that a vertex $x \in V$ is in V_1 if and only if $\ell_1 + \ell_Q(g_1, x) < \ell_2 + \ell_Q(g_2, x)$. By the symmetry of r_1 and r_2 , we may assume that v is contained in one of $Q[g_1, q_2]$, $Q[q_2, r_1]$, and $Q[r_1, g_2]$. In the remaining of the proof, we consider each case separately.

Case 1. Suppose that v appears on $Q[g_1, q_2]$. We consider the following cases separately.

- a. Suppose that w appears on $Q[g_1, v]$ or $Q[q_1, g_1]$ (Figures 16 and 17). Let e be the edge on $Q[g_1, v]$ that is incident to v and let $Q' := Q - \{e\} + \{uv\}$, which is a pseudotree. Since $\text{label}_Q(u) + \ell(uv) < \text{label}_Q(v)$, $\text{label}_Q(u) = \ell_1 + \ell_Q(g_1, u)$, and $\text{label}_Q(v) = \ell_1 + \ell_Q(g_1, v)$, it holds that $\ell_Q(g_1, u) + \ell(uv) < \ell_Q(g_1, v)$. Hence, $\ell(C_{Q'}) < \ell(C_Q)$ holds. For $i = 1, 2$ and for $x \in V_1$, we obtain

$$\begin{aligned} \ell_{Q'}(r_i, x) &\leq \ell_{Q'}(r_i, g_1) + \ell_{Q'}(g_1, x) \\ &\leq \ell_1 + \ell_Q(g_1, x) = \max\{\ell_Q(r_1, x), \ell_Q(r_2, x)\} \end{aligned}$$

by (15), which implies that $\bar{\ell}_{Q'}(r_i, x) \leq \frac{d}{2}$. Similarly, for $i = 1, 2$ and for $x \in V_2$, we obtain

$$\begin{aligned} \ell_{Q'}(r_i, x) &\leq \ell_{Q'}(r_i, g_2) + \ell_{Q'}(g_2, x) \\ &= \ell_2 + \ell_Q(g_2, x) = \max\{\ell_Q(r_1, x), \ell_Q(r_2, x)\}, \end{aligned}$$




Figure 19: w is on $Q[v, r_1]$




Figure 20: w is on $Q[r_1, g_2]$




Figure 21: w is on $Q[g_2, r_2]$




Figure 22: w is on $Q[r_2, q_1]$

which implies that $\bar{\ell}_{Q'}(r_i, x) \leq \frac{d}{2}$. This shows that $r_1 \xrightarrow{Q'} r_2$, and hence Q' satisfies the conditions in **B1**.

- b. Suppose that w appears on $Q[v, q_2]$ (Figure 18). Then, $\text{label}_Q(v) < \text{label}_Q(w) \leq \text{label}_Q(u)$, which contradicts the assumption that $\text{label}_Q(v) > \text{label}_Q(u) + \ell(uv)$.
- c. Suppose that w appears on $Q[q_2, r_1]$, $Q[r_1, g_2]$, $Q[g_2, r_2]$, or $Q[r_2, q_1]$. In this case, **B1** or **B2** holds by claim 28.

Case 2. Suppose that v appears on $Q[q_2, r_1]$. We consider the following cases separately.

- a. Suppose that w appears on $Q[q_1, g_1]$ or $Q[g_1, q_2]$. In this case, **B1** or **B2** holds by claim 28.
- b. Suppose that w appears on $Q[q_2, v]$. Then, $\text{label}_Q(v) < \text{label}_Q(w) \leq \text{label}_Q(u)$, which contradicts the assumption that $\text{label}_Q(v) > \text{label}_Q(u) + \ell(uv)$.
- c. Suppose that w appears on $Q[v, r_1]$, $Q[r_1, g_2]$, $Q[g_2, r_2]$, or $Q[r_2, q_1]$ (Figures 19–22). Let e be the edge on $Q[g_2, v]$ that is incident to v . Let $Q_0 := Q - \{e\} + \{uv\}$ and $Q_1 := Q - \{e_1\} + \{uw\}$. Since $\text{label}_Q(u) + \ell(uv) < \text{label}_Q(v)$, $\text{label}_Q(u) = \ell_2 + \ell_Q(g_2, u)$, and $\text{label}_Q(v) = \ell_2 + \ell_Q(g_2, v)$, it holds that $\ell_Q(g_2, u) + \ell(uv) < \ell_Q(g_2, v)$. Hence, $\ell(C_{Q_0}) < \ell(C_Q)$ holds. We also have that $\ell(C_{Q_1}) \leq \ell_Q(g_2, v) + \ell_Q(g_2, u) + \ell(uv) < 2\ell_Q(g_2, v) < \ell(C_Q)$. Therefore, the lengths of C_{Q_0} and C_{Q_1} are shorter than that of C_Q .

Let P be the unique g_2 - v path in $Q - \{e, e_1\} + \{uw\}$. Since $|E(P)| \leq \bar{\ell}_Q(g_2, v) = \bar{\ell}_2 + \bar{\ell}_Q(r_1, v)$, there exists a point $r_0 \in V \cup R$ on P such that $\bar{\ell}_P(r_0, g_2) \leq \bar{\ell}_2$ and $\bar{\ell}_P(r_0, v) \leq \bar{\ell}_Q(r_1, v)$. Note that such r_0 exists in $V \cup R$, because $\bar{\ell}_Q(r_1, v)$ is half-integral.

Then, we see that $r_0 \in \text{center}(Q_i)$ for $i = 0, 1$, $r_2 \in \text{center}(Q_0)$, and $r_1 \in \text{center}(Q_1)$ by the following observations.

- For $i = 0, 1$ and for $x \in V_1$,

$$\begin{aligned} \bar{\ell}_{Q_i}(r_0, x) &\leq \bar{\ell}_P(r_0, v) + \bar{\ell}_Q(v, g_1) + \bar{\ell}_Q(g_1, x) \\ &\leq \bar{\ell}_Q(r_1, v) + \bar{\ell}_Q(v, g_1) + \bar{\ell}_Q(g_1, x) \\ &= \bar{\ell}_1 + \bar{\ell}_Q(g_1, x) = \max\{\bar{\ell}_Q(r_1, x), \bar{\ell}_Q(r_2, x)\} \leq \frac{d}{2}. \end{aligned}$$




Figure 23: w is on $Q[v, g_2]$.




Figure 24: w is on $Q[g_2, r_2]$.

- For $i = 0, 1$ and for $x \in V_2$,

$$\begin{aligned} \bar{\ell}_{Q_i}(r_0, x) &\leq \bar{\ell}_P(r_0, g_2) + \bar{\ell}_{Q_i}(g_2, x) \\ &\leq \bar{\ell}_2 + \bar{\ell}_Q(g_2, x) = \max\{\bar{\ell}_Q(r_1, x), \bar{\ell}_Q(r_2, x)\} \leq \frac{d}{2}. \end{aligned}$$

- For $x \in V_j$ with $j \in \{1, 2\}$,

$$\begin{aligned} \bar{\ell}_{Q_0}(r_2, x) &\leq \bar{\ell}_{Q_0}(r_2, g_j) + \bar{\ell}_{Q_0}(g_j, x) \\ &\leq \bar{\ell}_j + \bar{\ell}_Q(g_j, x) = \max\{\bar{\ell}_Q(r_1, x), \bar{\ell}_Q(r_2, x)\} \leq \frac{d}{2}. \end{aligned}$$

- For $x \in V$, it holds that $\bar{\ell}_{Q_1}(r_1, x) \leq \bar{\ell}_{Q-e_1}(r_1, x) = \bar{\ell}_Q(r_1, x) \leq \frac{d}{2}$, because $Q - e_1$ is the shortest path tree starting from r_1 in Q .

Therefore, if $r_0 \notin \{r_1, r_2\}$, then we obtain $r_1 \xleftarrow{Q_1} r_0$ and $r_0 \xleftarrow{Q_0} r_2$, and hence r_0 , Q_1 , and Q_0 satisfy the conditions in **B2**. If $r_0 = r_1$ (resp. $r_0 = r_2$), then Q_0 (resp. Q_1) satisfies the conditions in **B1**.

Case 3. Suppose that v appears on $Q[r_1, g_2]$. We consider the following cases separately.

- Suppose that w appears on $Q[q_1, g_1]$ or $Q[g_1, q_2]$. In this case, **B1** or **B2** holds by claim 28.
- Suppose that w appears on $Q[q_2, v]$. Then, $\text{label}_Q(v) < \text{label}_Q(w) \leq \text{label}_Q(u)$, which contradicts the assumption that $\text{label}_Q(v) > \text{label}_Q(u) + \ell(uv)$.
- Suppose that w appears on $Q[v, g_2]$ or $Q[g_2, r_2]$ (Figures 23 and 24). Let e be the edge on $Q[v, g_2]$ that is incident to v and let $Q' := Q - \{e\} + \{uw\}$. Since $\text{label}_Q(u) + \ell(uw) < \text{label}_Q(v)$, $\text{label}_Q(u) = \ell_2 + \ell_Q(g_2, u)$, and $\text{label}_Q(v) = \ell_2 + \ell_Q(g_2, v)$, it holds that $\ell_Q(g_2, u) + \ell(uw) < \ell_Q(g_2, v)$. Hence, $\ell(C_{Q'}) < \ell(C_Q)$ holds. For $i = 1, 2$ and for $x \in V_1$, we obtain

$$\begin{aligned} \ell_{Q'}(r_i, x) &\leq \ell_{Q'}(r_i, g_1) + \ell_{Q'}(g_1, x) \\ &= \ell_1 + \ell_Q(g_1, x) = \max\{\ell_Q(r_1, x), \ell_Q(r_2, x)\} \end{aligned}$$

by (15), which implies that $\bar{\ell}_{Q'}(r_i, x) \leq \frac{d}{2}$. Similarly, for $i = 1, 2$ and for $x \in V_2$, we obtain

$$\begin{aligned} \ell_{Q'}(r_i, x) &\leq \ell_{Q'}(r_i, g_2) + \ell_{Q'}(g_2, x) \\ &\leq \ell_2 + \ell_Q(g_2, x) = \max\{\ell_Q(r_1, x), \ell_Q(r_2, x)\}, \end{aligned}$$

which implies that $\bar{\ell}_{Q'}(r_i, x) \leq \frac{d}{2}$. This shows that $r_1 \xleftarrow{Q'} r_2$, and hence Q' satisfies the conditions in **B1**.

- Suppose that w appears on $Q[r_2, q_1]$. Then, $\text{label}_Q(v) \leq \text{label}_Q(r_1) = \text{label}_Q(r_2) \leq \text{label}_Q(w) \leq \text{label}_Q(u)$, which contradicts the assumption that $\text{label}_Q(v) > \text{label}_Q(u) + \ell(uw)$.

By the above case analysis, the proposition holds. \square

By this proposition, we can derive Theorem 24.

Proof of Theorem 24. It suffices to prove the following statement by the induction on $f(r_1, r_2, Q)$: for a triplet (r_1, r_2, Q) with $r_1 \xleftrightarrow{Q} r_2$, \mathcal{G}' contains an r_1 - r_2 path.

Since $f(\cdot)$ can take only finitely many values, there exists a triplet (r_1^*, r_2^*, Q^*) with $r_1^* \xleftrightarrow{Q^*} r_2^*$ that minimizes $f(r_1^*, r_2^*, Q^*)$. If (r_1^*, r_2^*, Q^*) is not good, then we can apply proposition 27 to obtain a triple with smaller $f(\cdot)$, which contradicts the minimality of $f(r_1^*, r_2^*, Q^*)$. Therefore, (r_1^*, r_2^*, Q^*) is good, and hence $r_1^* r_2^* \in E(\mathcal{G}')$, which shows the base case of the statement.

To show the induction step, suppose that (r_1, r_2, Q) is a triplet with $r_1 \xleftrightarrow{Q} r_2$. If (r_1, r_2, Q) is good, then $r_1 r_2 \in E(\mathcal{G}')$, and hence we are done. Otherwise, we apply proposition 27 to obtain **B1** or **B2**. When we obtain **B1**, since $r_1 \xleftrightarrow{Q'} r_2$ and $f(r_1, r_2, Q') < f(r_1, r_2, Q)$, \mathcal{G}' contains an r_1 - r_2 path by the induction hypothesis. When we obtain **B2**, the induction hypothesis shows that \mathcal{G}' contains an r_1 - r_0 path and an r_0 - r_2 path, which shows the existence of an r_1 - r_2 path in \mathcal{G}' . \square

9 Concluding Remarks

In this paper, we have investigated the computational complexity of RST WITH SMALL (or LARGE) MAXIMUM DEGREE and RST WITH SMALL (or LARGE) DIAMETER.

We have proved in Theorem 2 that RST WITH SMALL MAXIMUM DEGREE is PSPACE-complete for $d \geq 3$. One can naturally ask what happens for the case of maximum degree at most 2. In this case, the problem becomes the HAMILTONIAN PATH RECONFIGURATION problem, in which a feasible solution is a Hamiltonian path. We were not able to determine the complexity of this problem and we left it as an open problem. Note that HAMILTONIAN PATH RECONFIGURATION problem can be also seen as a special case of RST WITH LARGE DIAMETER in which the lower bound on the diameter is $|V(G)| - 1$. Note also that, for the Hamiltonian *cycle* case, the HAMILTONIAN CYCLE RECONFIGURATION problem is known to be PSPACE-complete [21], in which two edge flips are executed in one step.

We have proved in Theorem 3 that RST WITH LARGE DIAMETER is NP-hard, but it is unclear whether this problem belongs to the class NP. We conjecture that the problem is PSPACE-complete, and left this question as another open problem.

References

- [1] Matthieu Barjon, Arnaud Casteigts, Serge Chaumette, Colette Johnen, and Yessin M. Neggaz. Maintaining a spanning forest in highly dynamic networks: The synchronous case. In Marcos K. Aguilera, Leonardo Querzoni, and Marc Shapiro, editors, *Principles of Distributed Systems*, pages 277–292, Cham, 2014. Springer International Publishing.
- [2] Nicolas Bousquet, Takehiro Ito, Yusuke Kobayashi, Haruka Mizuta, Paul Ouvrard, Akira Suzuki, and Kunihiro Wasa. Reconfiguration of spanning trees with many or few leaves. In *28th Annual European Symposium on Algorithms, ESA 2020, September 7-9, 2020, Pisa, Italy (Virtual Conference)*, pages 24:1–24:15, 2020.
- [3] Sergio Cabello, Erin W. Chambers, and Jeff Erickson. Multiple-source shortest paths in embedded graphs. *SIAM Journal on Computing*, 42(4):1542–1571, 2013.
- [4] Artur Czumaj and Willy-B. Strothmann. Bounded degree spanning trees. In Rainer Burkard and Gerhard Woeginger, editors, *Algorithms — ESA '97*, volume 1284 of LNCS, pages 104–117, Berlin, Heidelberg, 1997. Springer Berlin Heidelberg.

- [5] Martin Furer and Balaji Raghavachari. Approximating the minimum-degree steiner tree to within one of optimal. *Journal of Algorithms*, 17(3):409–423, 1994.
- [6] Michel X. Goemans. Minimum bounded degree spanning trees. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), 21-24 October 2006, Berkeley, California, USA, Proceedings*, pages 273–282. IEEE Computer Society, 2006.
- [7] Refael Hassin and Asaf Levin. Minimum restricted diameter spanning trees. *Discrete Applied Mathematics*, 137(3):343–357, 2004.
- [8] Refael Hassin and Arie Tamir. On the minimum diameter spanning tree problem. *Inf. Process. Lett.*, 53(2):109–111, January 1995.
- [9] Robert A. Hearn and Erik D. Demaine. PSPACE-completeness of sliding-block puzzles and other problems through the nondeterministic constraint logic model of computation. *Theoretical Computer Science*, 343(1-2):72–96, 2005.
- [10] Jan van den Heuvel. The complexity of change. In Simon R. Blackburn, Stefanie Gerke, and Mark Wildon, editors, *Surveys in Combinatorics*, volume 409 of *London Mathematical Society Lecture Note Series*, pages 127–160. Cambridge University Press, 2013.
- [11] Silu Huang, Ada Wai-Chee Fu, and Ruifeng Liu. Minimum spanning trees in temporal graphs. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, SIGMOD '15*, page 419–430, New York, NY, USA, 2015. Association for Computing Machinery.
- [12] Giuseppe F. Italiano and Rajiv Ramaswami. Maintaining spanning trees of small diameter. *Algorithmica*, 22(3):275–304, 1998.
- [13] Takehiro Ito, Erik D. Demaine, Nicholas J.A. Harvey, Christos H. Papadimitriou, Martha Sideri, Ryuhei Uehara, and Yushi Uno. On the complexity of reconfiguration problems. *Theoretical Computer Science*, 412(12):1054–1065, 2011.
- [14] Richard M. Karp. Reducibility among combinatorial problems. In R. Miller and J. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.
- [15] Bernhard H. Korte and Jens Vygen. *Combinatorial Optimization: Theory and Algorithms*. Springer-Verlag, New York, NY, 2012.
- [16] Haruka Mizuta, Tatsuhiko Hatanaka, Takehiro Ito, and Xiao Zhou. Reconfiguration of minimum steiner trees via vertex exchanges. In *44th International Symposium on Mathematical Foundations of Computer Science, MFCS 2019, August 26-30, 2019, Aachen, Germany.*, pages 79:1–79:11, 2019.
- [17] Naomi Nishimura. Introduction to reconfiguration. *Algorithms*, 11(4):52, 2018.
- [18] Alexander Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.
- [19] Mohit Singh and Lap Chi Lau. Approximating minimum bounded degree spanning trees to within one of optimal. *J. ACM*, 62(1), March 2015.
- [20] Michael J. Spriggs, J. Mark Keil, Sergei Bespamyatnikh, Michael Segal, and Jack Snoeyink. Computing a $(1 + \epsilon)$ -approximate geometric minimum-diameter spanning tree. *Algorithmica*, 38(4):577–589, 2004.
- [21] Asahi Takaoka. Complexity of Hamiltonian cycle reconfiguration. *Algorithms*, 11(9), 2018.