December 2019

Rearrangement operations on unrooted phylogenetic networks

Remie Janssen

Delft University of Technology, remiejanssen92@gmail.com

Jonathan Klawitter

University of Auckland, jo.klawitter@gmail.com

Follow this and additional works at: <https://digitalcommons.georgiasouthern.edu/tag>

Part of the [Discrete Mathematics and Combinatorics Commons](#)

Recommended Citation

Janssen, Remie and Klawitter, Jonathan (2019) "Rearrangement operations on unrooted phylogenetic networks," *Theory and Applications of Graphs*: Vol. 6 : Iss. 2 , Article 6.

DOI: 10.20429/tag.2019.060206

Available at: <https://digitalcommons.georgiasouthern.edu/tag/vol6/iss2/6>

This article is brought to you for free and open access by the Journals at Digital Commons@Georgia Southern. It has been accepted for inclusion in Theory and Applications of Graphs by an authorized administrator of Digital Commons@Georgia Southern. For more information, please contact digitalcommons@georgiasouthern.edu.

Rearrangement operations on unrooted phylogenetic networks

Cover Page Footnote

The first author was supported by the Netherlands Organization for Scientific Research (NWO) Vidi grant 639.072.602. The second author thanks the New Zealand Marsden Fund for their financial support.

Abstract

Rearrangement operations transform a phylogenetic tree into another one and hence induce a metric on the space of phylogenetic trees. Popular operations for unrooted phylogenetic trees are NNI (nearest neighbour interchange), SPR (subtree prune and regraft), and TBR (tree bisection and reconnection). Recently, these operations have been extended to unrooted phylogenetic networks—generalisations of phylogenetic trees that can model reticulated evolutionary relationships—where they are called NNI, PR, and TBR moves. Here, we study global and local properties of spaces of phylogenetic networks under these three operations. In particular, we prove connectedness and asymptotic bounds on the diameters of spaces of different classes of phylogenetic networks, including tree-based and level- k networks. We also examine the behaviour of shortest TBR-sequence between two phylogenetic networks in a class, and whether the TBR-distance changes if intermediate networks from other classes are allowed: for example, the space of phylogenetic trees is an isometric subgraph of the space of phylogenetic networks under TBR. Lastly, we show that computing the TBR-distance and the PR-distance of two phylogenetic networks is NP-hard.

1 Introduction

Phylogenetic trees and networks are leaf-labelled graphs that are used to visualise and study the evolutionary history of taxa like species, genes, or languages. While phylogenetic trees are used to model tree-like evolutionary histories, the more general phylogenetic networks can be used for taxa whose past includes reticulate events like hybridisation or horizontal gene transfer [SS03, HRS10, Ste16]. Such reticulate events arise in all domains of life [TN05, RW07, MMM⁺17, WWK⁺17]. In some cases, it can be useful to distinguish between rooted and unrooted phylogenetic networks. In a rooted phylogenetic network, the edges are directed from a designated root towards the leaves. Hence, it models evolution along the passing of time. An unrooted phylogenetic network, on the other hand, has undirected edges and thus represent evolutionary relatedness of the taxa. In some cases, unrooted phylogenetic networks can be thought of as rooted phylogenetic networks in which the orientation of the edges has been disregarded. Such unrooted phylogenetic networks are called proper [JJE⁺18, FHM18]. Here we focus on unrooted, binary, proper phylogenetic networks, where binary means that all vertices except for the leaves have degree three. The set of phylogenetic networks on the same taxa can be partitioned into tiers that contain all networks of the same size.

A rearrangement operation transforms a phylogenetic tree into another tree by making a small graph theoretical change. An operation that works locally within the tree is the NNI (nearest neighbour interchange) operation, which changes the order of the four edges incident to an edge e . See for example the NNI from T_1 to T_2 in Figure 1. Two further popular rearrangement operations are the SPR (subtree prune and regraft) operation, which as the name suggests prunes (cuts) an edge and then regrafts (attaches) the resulting half edge again, and the TBR (tree bisection and reconnection) operation, which first removes an edge and then adds a new one to reconnect the resulting two smaller trees. See, for example, the SPR from T_2 to T_3 and the TBR from T_3 to T_4 in Figure 1.

The set of phylogenetic trees on a fixed set of taxa together with a rearrangement operation yields a graph where the vertices are the trees and two trees are adjacent if

they can be transformed into each other with the operation. We call this a space of phylogenetic trees. This construction also induces a metric on phylogenetic trees as the distance of two trees is then given as the distance in this space, that is, the minimum number of applications of the operation that are necessary to transform one tree into the other [SOW96]. However, computing the distance of two trees under NNI, SPR, and TBR is NP-hard [DHJ⁺97, HDRCB08, AS01]. Nevertheless, both the space of phylogenetic trees and a metric on them are of importance for the many inference methods for phylogenetic trees that rely on local search strategies [Gus14, SJ17].

Figure 1: The three rearrangement operations on unrooted phylogenetic trees: The NNI from T_1 to T_2 changes the order of the four edges incident to e ; the SPR from T_2 to T_3 prunes the edge e' , and then regrafts it again; and the TBR from T_3 to T_4 first removes the edge e'' , and then reconnects the resulting two trees with a new edge. Note that every NNI is also an SPR and every SPR is also a TBR but not vice versa.

Recently, these rearrangement operations have been generalised to phylogenetic networks, both for unrooted networks [HLMW16, HMW16, FHMW18] and for rooted networks [BLS17, FHMW18, GvIJ⁺17, Kla19]. For unrooted networks, Huber et al. [HLMW16] first generalised NNI to level-1 networks, which are phylogenetic networks where all cycles are vertex disjoint. This generalisation includes a horizontal move that changes the topology of the network, like an NNI on a tree, and vertical moves that add or remove a triangle to change the size of the network. Among other results, they then showed that the space of level-1 networks and its tiers are connected under NNI [HLMW16, Theorem 2]. Note that connectedness implies that the distance between any two networks in such a space is finite and that NNI thus induces a metric. This NNI operation was then extended by Huber et al. [HMW16] to work for general unrooted phylogenetic networks. Again, connectedness of the space was proven. Later, Francis et al. [FHMW18] gave lower and upper bounds on the diameter (the maximum distance) of the space of unrooted phylogenetic network of a fixed size under NNI. They also showed that SPR and TBR can straightforwardly be generalised to phylogenetic networks, that the connectedness under NNI implies connectedness under SPR and TBR, and they gave bounds on the diameters. These bounds for SPR were made asymptotically tight by Janssen et al. [JJE⁺18]. Here, we improve these bounds on the diameter under TBR.

There are several generalisations of SPR on rooted phylogenetic trees to rooted phylogenetic networks for which connectedness and diameters have been obtained [BLS17, FHMW18, GvIJ⁺17, JJE⁺18, Jan18]. For example, Bordewich et al. [BLS17] introduced SNPR (subnet prune and regraft), a generalisation of SPR that includes vertical moves, which add or remove an edge. They then proved connectedness under SNPR for the space of rooted phylogenetic networks and for special classes of phylogenetic networks including tree-based networks. Roughly speaking, these are networks that have a spanning tree that is the subdivision of a phylogenetic tree on the same taxa [FS15, FHM18]. Furthermore, Bordewich et al. [BLS17] gave several bounds on the SNPR-distance of two phylogenetic

networks. Further bounds and a characterisation of the SNPR-distance of a tree and a network were recently proven by Klawitter and Linz [KL19]. Here, we show that these bounds and characterisation on the SNPR-distance of rooted phylogenetic networks are analogous to the TBR-distance of two unrooted phylogenetic networks.

In this paper, we study spaces of unrooted phylogenetic networks under NNI, PR (prune and regraft), and TBR. Here, the PR and the TBR operation are the generalisation of SPR and TBR on trees, respectively, where vertical moves add or remove an edge like the vertical moves of the SNPR operation in the rooted case. After the preliminary section, we examine the relation of NNI, PR, and TBR; in particular, how a sequence using one of these operations can be transformed into a sequence using another operation (Section 3). We then study properties of shortest paths under TBR in Section 4. This includes the translation of the results from Bordewich et al. [BLS17] and Klawitter and Linz [KL19] on the SNPR-distance of rooted phylogenetic networks to the TBR-distance of unrooted phylogenetic networks. Next, we consider the connectedness and diameters of spaces of phylogenetic networks for different classes of phylogenetic networks, including tree-based networks and level- k networks (Section 5). A subspace of phylogenetic networks (e.g., the space of tree-based networks) is an isometric subgraph of a larger space of phylogenetic networks if, roughly speaking, the distance of two networks is the same in the smaller and the larger space. In Section 6 we study such isometric relations and answer a question by Francis et al. [FHMW18] by showing that the space of phylogenetic trees is an isometric subgraph of the space of phylogenetic networks under TBR. We use this result in Section 7 to show that computing the TBR-distance is NP-hard. In the same section, we also show that computing the PR-distance is NP-hard.

2 Preliminaries

This section provides notation and terminology used in the remainder of the paper. In particular, we define phylogenetic networks and special classes thereof, and rearrangement operations and how they induce distances. Throughout this paper, $X = \{1, 2, \dots, n\}$ denotes a finite set of taxa.

Phylogenetic networks. An *unrooted, binary phylogenetic network* N on a set of taxa X is an undirected multigraph such that the leaves are bijectively labelled with X and all non-leaf vertices have degree three. It is called *proper* if every cut-edge separates two labelled leaves [FHM18], and *improper* otherwise. This property implies that every edge lies on a path that connects two leaves. More importantly, a network can be rooted at any leaf if and only if it is proper [JJE⁺18, Lemma 4.13]. If not mentioned otherwise, we assume that a phylogenetic network is proper. Furthermore, note that our definition of a phylogenetic network permits the existence of parallel edges in N , i.e., we allow that two distinct edges join the same pair of vertices. An *unrooted, binary phylogenetic tree* T on X is an unrooted, binary phylogenetic network on X that is a tree.

Let $u\mathcal{N}_n$ denote the set of all unrooted, binary proper phylogenetic networks on X and let $u\mathcal{T}_n$ denote the set of all unrooted, binary phylogenetic trees on X , where $X = \{1, 2, \dots, n\}$. To ease reading, we refer to an unrooted, binary proper phylogenetic network (resp. unrooted, binary phylogenetic tree) on X simply as phylogenetic network or network (resp. phylogenetic tree or tree). Figure 2 shows an example of a tree $T \in u\mathcal{T}_6$, a network in $N \in u\mathcal{N}_6$, and an improper network M .

Figure 2: An unrooted, binary phylogenetic tree $T \in u\mathcal{T}_6$ and an unrooted, binary proper phylogenetic network $N \in u\mathcal{N}_6$. The unrooted, binary phylogenetic network M is improper since the cut-edge e does not lie on a path that connects two leaves.

An edge of a network N is an *external* edge if it is incident to a leaf, and an *internal* edge otherwise. A *cherry* $\{a, b\}$ of N is a pair of leaves a and b in N that are adjacent to the same vertex. For example, each network in Figure 2 contains the cherry $\{1, 5\}$.

Tiers. We say a network $N = (V, E)$ has *reticulation number*¹ r for $r = |E| - (|V| - 1)$, that is, the number of edges that have to be deleted from N to obtain a spanning tree of N . For example, the network N in Figure 2 has reticulation number three. Note that a phylogenetic tree is a phylogenetic network with reticulation number zero. Let $u\mathcal{N}_{n,r}$ denote *tier* r of $u\mathcal{N}_n$, the set of networks in $u\mathcal{N}_n$ that have reticulation number r .

Embedding. Let G be an undirected graph. *Subdividing* an edge $\{u, v\}$ of G consists of replacing $\{u, v\}$ by a path from u to v that contains at least one edge. A *subdivision* G^* of G is a graph that can be obtained from G by subdividing edges of G . If G has no degree two vertices, there exists a canonical embedding of vertices of G to vertices of G^* and of edges of G to paths of G^* . Let $N \in u\mathcal{N}_n$. We say G has an *embedding* into N if there exists a subdivision G^* of G that is a subgraph of N such that the embedding maps each labelled vertex of G^* to a labelled vertex of N with the same label.

Displaying. Let $T \in u\mathcal{T}_n$ and $N \in u\mathcal{N}_n$. We say N *displays* T if T has an embedding into N . For example, in Figure 2 the tree T is displayed by both networks N and M . Let $D(N)$ be the set of trees in $u\mathcal{T}_n$ that are displayed by N . This notion can be extended to trees with fewer leaves, and to networks. For this, let M be a phylogenetic network on $Y \subseteq X = \{1, \dots, n\}$. We say N *displays* M if M has an embedding into N . Let $P = \{M_1, \dots, M_k\}$ be a set of phylogenetic networks M_i on $Y_i \subseteq X = \{1, \dots, n\}$. Then let $u\mathcal{N}_n(P)$ denote the subset of networks in $u\mathcal{N}_n$ that display each network in P .

Tree-based networks. A phylogenetic network $N \in u\mathcal{N}_n$ is a *tree-based* network if there is a tree $T \in u\mathcal{T}_n$ that has an embedding into N as a spanning tree. In other words, there exists a subdivision T^* of T that is a spanning tree of N . The tree T is then called a *base tree* of N . Let $u\mathcal{TB}_n$ denote the set of tree-based networks in $u\mathcal{N}_n$. For $T \in u\mathcal{T}_n$, let $u\mathcal{TB}_n(T)$ denote the set of tree-based networks in $u\mathcal{TB}_n$ with base tree T .

Level- k networks. A blob B of a network $N \in u\mathcal{N}_n$ is a nontrivial two-connected component of N . The *level* of B is the minimum number of edges that have to be removed from B to make it acyclic. The *level* of N is the maximum level of all blobs of

¹In graph theory the value $|E| - (|V| - 1)$ of a connected graph is also called the cyclomatic number of the graph [Die17].

N . If the level of N is at most k , then N is called a *level- k* network. Let $u\mathcal{LV}\text{-}k_n$ denote the set of level- k networks in $u\mathcal{N}_n$.

r -Burl. An r -burl is a specific type of blob that we define recursively: a 1-burl is the blob consisting of a pair of parallel edges; an r -burl is the blob obtained by placing a pair of parallel edges on one of the parallel edges of an $r - 1$ -burl for all $r > 1$. See for example the network M in Figure 3.

r -Handcuffed trees and caterpillars. Let $T \in u\mathcal{N}_n$ and let a and b be two leaves of T . Let e and f be the edges incident to a and b , respectively. Subdivide e and f with vertices $\{u_1, \dots, u_r\}$ and $\{v_1, \dots, v_r\}$, respectively, and add the edges $\{u_1, v_1\}, \dots, \{u_r, v_r\}$. The resulting network is an r -handcuffed tree $N \in u\mathcal{N}_n$ with base tree T on the *handcuffed* leaves $\{a, b\}$. Note that N has reticulation number r . If the tree T is a caterpillar and a and b form a cherry of T , then the resulting network N is an r -handcuffed caterpillar. Furthermore, we call an r -handcuffed caterpillar *sorted* if it is handcuffed on the leafs 1 and 2 and the leafs from 3 to n have a non-decreasing distance to leaf 1. See Figure 3 for an example.

Figure 3: A network M with a 3-burl and a sorted 3-handcuffed caterpillar N .

Suboperations. To define rearrangement operations on phylogenetic networks, we first define several suboperations. Let G be an undirected graph. A degree-two vertex v of G with adjacent vertices u and w gets *suppressed* by deleting v and its incident edges, and adding the edge $\{u, w\}$. The reverse of this suppression is the subdivision of $\{u, w\}$ with vertex v .

Let $N \in u\mathcal{N}_n$ be a network, and $\{u, v\}$ an edge of N . Then $\{u, v\}$ gets *removed* by deleting $\{u, v\}$ from N and suppressing any resulting degree-two vertices. We say $\{u, v\}$ gets *pruned* at u by transforming it into the half edge $\{\cdot, v\}$ and suppressing u if it becomes a degree-two vertex. Note that otherwise u is a leaf. In reverse, we say that a half edge $\{\cdot, v\}$ gets *regrafted* to an edge $\{x, y\}$ by transforming it into the edge $\{u, v\}$ where u is a new vertex subdividing $\{x, y\}$.

TBR. A TBR operation² is the rearrangement operation that transforms a network $N \in u\mathcal{N}_n$ into another network $N' \in u\mathcal{N}_n$ in one of the following four ways:

(TBR⁰) Remove an internal edge e of N , subdivide an edge of the resulting graph with a new vertex u , subdivide an edge of the resulting graph with a new vertex v , and add the edge $\{u, v\}$;

or, prune an external edge $e = \{u, v\}$ of N that is incident to leaf v at u , regraft $\{\cdot, v\}$ to an edge of the resulting graph.

(TBR⁺) Subdivide an edge of N with a new vertex u , subdivide an edge of the resulting graph with a new vertex v , and add the edge $e = \{u, v\}$.

(TBR⁻) Remove an edge e of N .

Note that a TBR⁰ can also be seen as the operation that prunes the edge $e = \{u, v\}$ at both u and v and then regrafts both ends. Hence, we say that a TBR⁰ *moves* the edge e . Furthermore, we say that a TBR⁺ *adds* the edge e and that a TBR⁻ *removes* the edge e . These operations are illustrated in Figure 4. Note that a TBR⁰ has an inverse TBR⁰ and that a TBR⁺ has an inverse TBR⁻, and that furthermore a TBR⁺ increases the reticulation number by one and a TBR⁻ decreases it by one.

Since a TBR operation has to yield a phylogenetic network, there are some restrictions on the edges that can be moved or removed. Firstly, if removing an edge by a TBR⁰ yields a disconnected graph, then in order to obtain a phylogenetic network an edge has to be added between the two connected components. Similarly, a TBR⁻ cannot remove a cut-edge. Secondly, the suppression of a vertex when removing an edge may not yield a loop $\{u, u\}$. Thirdly, removing or moving an edge cannot create a cut-edge that does not separate two leaves. Otherwise the network would not be proper.

Figure 4: Illustration of the TBR operation. The network N_2 can be obtained from N_1 by a TBR⁰ that moves the edge $\{u, v\}$ and the network N_3 can be obtained from N_2 by a TBR⁺ that adds the edge $\{u', v'\}$. Each operation has its corresponding TBR⁰ and TBR⁻ operation, respectively, that reverses the rearrangement.

The TBR⁰ operation equals the well known TBR (tree bisection and reconnection) operation on unrooted phylogenetic trees [AS01]. The TBR operation on trees has recently been generalised to TBR⁰ on improper unrooted phylogenetic networks by Francis et al. [FHMW18].

PR. A PR (*prune and regraft*) operation is the rearrangement operation that transforms a network $N \in \mathcal{UN}_n$ into another network $N' \in \mathcal{UN}_n$ with a $\text{PR}^+ = \text{TBR}^+$, a $\text{PR}^- = \text{TBR}^-$, or a PR^0 that prunes and regrafts an edge e only at one endpoint, instead of at both like a TBR⁰. Like for TBR, we say that the $\text{PR}^{0/+/-}$ *moves/adds/removes* the edge e in N . The PR operation is a generalisation of the well-known SPR (subtree prune and regraft) operation on unrooted phylogenetic trees [AS01]. Like for TBR, the generalisation of SPR to PR^0 for networks has been introduced by Francis et al. [FHMW18].

²The TBR operation is known on unrooted phylogenetic trees as *tree bisection and reconnection*. Since in general networks are not trees and a TBR on a network does not necessarily bisect it, we use TBR now as a word on its own. For the reader who would however like to have an expansion of TBR we suggest "total branch relocation". We welcome other suggestions.

NNI. An NNI (*nearest neighbour interchange*) operation is a rearrangement operation that transforms a network $N \in \mathcal{uN}_n$ into another network $N' \in \mathcal{uN}_n$ in one of the following three ways:

- (NNI⁰) Let $e = \{u, v\}$ be an internal edge of N . Prune an edge f ($f \neq e$) at u , and regraft it to an edge f' ($f' \neq e$) that is incident to v .
- (NNI⁺) Subdivide two adjacent edges with new vertices u' and v' , respectively, and add the edge $\{u', v'\}$.
- (NNI⁻) If N contains a triangle, remove an edge of the triangle.

These operations are illustrated in Figure 5. We say that an NNI⁰ *moves* the edge f . Alternatively, we call the edge e of an NNI⁰ the *axis* of the operation, as the operation can also be defined as pruning f at u , and $f'' \neq f'$ at v , and regrafting f at v and f'' at u . The NNI operation has been introduced on trees by Robinson [Rob71] and generalised to networks by Huber et al. [HLMW16, HMW16].

Figure 5: Illustration of the NNI operation. The network N_2 (resp. N_3) can be obtained from N_1 (resp. N_2) by an NNI⁰ with the axis $\{u, v\}$; alternatively, N_2 can be obtained from N_1 using the NNI⁰ of $\{1, u\}$ to the triangle, and N_3 from N_2 by moving $\{1, u\}$ to the bottom edge of the square. The labels are inherited naturally following the first interpretation of the NNI⁰ moves. The network N_4 can be obtained from N_3 by an NNI⁺ that extends x into a triangle. Each operation has its corresponding NNI⁰ and NNI⁻ operation, respectively, that reverses the transformation.

Sequences and distances. Let $N, N' \in \mathcal{uN}_n$ be two networks. A TBR-*sequence* from N to N' is a sequence

$$\sigma = (N = N_0, N_1, N_2, \dots, N_k = N')$$

of phylogenetic networks such that N_i can be obtained from N_{i-1} by a single TBR for each $i \in \{1, 2, \dots, k\}$. The *length* of σ is k . The TBR-*distance* $d_{\text{TBR}}(N, N')$ between N and N' is the length of a shortest TBR-sequence from N to N' , or infinite if no such sequence exists.

Let $\mathcal{C}_n$ be a class of phylogenetic networks. The TBR-distance on $\mathcal{C}_n$ is defined like on $\mathcal{uN}_n$ but with the restriction that every network in a shortest TBR-sequence has to be in $\mathcal{C}_n$. The class $\mathcal{C}_n$ is *connected* under TBR if, for all pairs $N, N' \in \mathcal{C}_n$, there exists a TBR-sequence σ from N to N' such that each network in σ is in $\mathcal{C}_n$. Hence, for the TBR-distance to be a metric on $\mathcal{C}_n$, the class has to be connected under TBR and the TBR operation has to be reversible. We already noted above that the latter holds for TBR (and NNI and PR). For a connected class $\mathcal{C}_n$, the *diameter* is the maximum distance between two of its networks under its metric. The definition for NNI and PR are analogous.

Let $\mathcal{C}'_n$ be a subclass of $\mathcal{C}_n$. Then $\mathcal{C}'_n$ is an *isometric subgraph* of a $\mathcal{C}_n$ under, say, TBR if for every $N, N' \in \mathcal{C}'_n$ the TBR-distance of N and N' in $\mathcal{C}'_n$ equals the TBR-distance of N and N' in $\mathcal{C}_n$.

3 Relations of rearrangement operations

On trees, it is well known that every NNI is also an SPR, which, in turn, is also a TBR. We observe that the same holds for the generalisations of these operations as defined above.

Observation 3.1.

Let $N \in u\mathcal{N}_n$. Then, on N , every NNI is a PR and every PR is a TBR.

For the reverse direction, we first show that every TBR can be mimicked by at most two PR like in $u\mathcal{T}_n$. Then we show how to substitute a PR with an NNI-sequence.

Lemma 3.2.

Let $N, N' \in u\mathcal{N}_n$ such that $d_{\text{TBR}}(N, N') = 1$. Then $1 \leq d_{\text{PR}}(N, N') \leq 2$, where a TBR^0 may be replaced by two PR^0 .

Proof. If N' can be obtained from N by a TBR^+ or TBR^- , then by the definition of PR^+ and PR^- it follows that $d_{\text{PR}}(N, N') = 1$. If N' can be obtained from N by a TBR^0 that is also a PR^0 , the statement follows. Assume therefore that N' can be obtained from N by a TBR^0 that moves the edge $e = \{u, v\}$ of N to $e' = \{x, y\}$ of N' . Let G be the graph obtained from N by removing e , or equivalently the graph obtained from N' by removing e' . If e is a cut-edge, then so is e' , and without loss of generality u and x as well as v and y subdivide an edge in the same connected components of G . Furthermore, if u subdivides an edge of a pendant blob in G , then so does x . Otherwise N' would not be proper. Therefore, the PR^0 that prunes e at u and regrafts it to obtain x yields a phylogenetic network N'' . The choices of u and x ensure that N'' is connected and proper. There is then a PR^0 from N'' to N' that prunes $\{x, v\}$ at v and regrafts it at y to obtain N' . Hence, $d_{\text{PR}}(N, N') \leq 2$. $\square$

Corollary 3.3.

Let $N, N' \in u\mathcal{N}_n$. Then $d_{\text{TBR}}(N, N') \leq d_{\text{PR}}(N, N') \leq 2 d_{\text{TBR}}(N, N')$.

Lemma 3.4.

Let $N, N' \in u\mathcal{N}_{n,r}$ such that there is a PR^0 that transforms N into N' . Let e be the edge of N pruned by this PR^0 .

Then there exists an NNI^0 -sequence from N to N' that only moves e and whose length is in $\mathcal{O}(n+r)$. Moreover, if neither N nor N' contains parallel edges, then neither does any intermediate networks in the NNI-sequence.

Proof. Assume that N can be transformed into N' by pruning the edge $e = \{u, v\}$ at u and regrafting it to $f = \{x, y\}$. Note that there is then a (shortest) path $P = (u = v_0, v_1, v_2, \dots, v_k = x)$ from u to x in $N \setminus \{e\}$, since otherwise N' would be disconnected. Without loss of generality, assume that P does not contain y . Furthermore, assume for now that P does not contain v . The idea is now to move e along P to f with NNI^0 . In particular, we show how to construct a sequence $\sigma = (N = N_0, N_1, \dots, N_k = N')$ such that either N_{i+1} can be obtained from N_i by an NNI^0 or $N_{i+1} = N_i$, and such that N_i

contains the edge $e_i = \{v_i, v\}$. This process is illustrated in Figure 6. Assume we have constructed the sequence up to N_i . Let $g = \{v_{i+1}, w\}$ with $w \neq v$ be the edge incident to v_{i+1} that is not on P . Obtain N_{i+1} from N_i by swapping e_i and g with an NNI^0 on the axis $\{v_i, v_{i+1}\}$. Note that this preserves the path P and that N_{i+1} may only contain a parallel edge if N or N' contains parallel edges. As a result, we get $N_k = N'$.

Figure 6: How to mimic the PR^0 that prunes the edge $\{u, v\}$ at u and regrafts to $\{x, y\}$ with NNI^0 operations that move u of $\{u, v\}$ along the path $P = (u = v_0, v_1, v_2 = x)$ (for the proof of Lemma 3.4). Labels follow the definition of NNI^0 along an axis.

It remains to show that every network in σ is proper. Assume otherwise and let N_{i+1} be the first improper network in σ . Then N_{i+1} contains a cut-edge e_c that separates a blob B from all leaves. We claim that e_c is part of P . Indeed, the pruning of the NNI^0 from N_i to N_{i+1} has to create B and the regrafting cannot be to B , so it has to pass along e_c (Figure 7). However, as P is a path, the moving edge cannot pass e_c again, so all networks N_j for $j > i$ including N' are improper; a contradiction. Hence, all intermediate networks N_i are proper and thus σ is an NNI^0 -sequence from N to N' .

Figure 7: How an NNI^0 in the proof of Lemma 3.4 may result an improper network where e_c separates a blob B from all leaves. The moving edge $\{v, v_i\}$ of N_i becomes the moving edge $\{v, v_{i+1}\}$ of N_{i+1} . Labels follow the definition of NNI^0 along an axis.

Next, assume that P contains $v_i = v$. Then first apply the process above to move v of $\{u, v\}$ along $P' = (v = v_i, v_{i+1}, \dots, v_k)$ to v_k . In the resulting network, apply the process above to move u of $\{u, v\} = \{u, v_k\}$ along $P'' = (u = v_0, v_1, \dots, v_i)$ to v_i . The process again avoids the creation of a network N_j with parallel edges, if neither N nor N' contains parallel edges. Furthermore, from Figure 7 we get that if σ would contain improper network then u would be contained in the blob B . However, then $\{u, v\}$ and e_c would be edges from B to the rest of the network; again a contradiction.

Lastly, note that the length of P is in $\mathcal{O}(n + r)$ since N contains only $2n + 3r - 1$ edges. Hence, the length of σ is also in $\mathcal{O}(n + r)$. $\square$

Lemma 3.5.

Let $n \geq 3$. Let $N, N' \in \mathcal{UN}_n$ such that there is a PR^- that transforms N into N' . Let e

be the edge of N removed by this PR^- . Let N have reticulation number r .

Then, there is an NNI^0 -sequence followed by one NNI^- that transforms N and N' by only moving and removing e and whose length is in $\mathcal{O}(n+r)$. Moreover, if neither N nor N' contains parallel edges, then neither do the intermediate networks in the NNI -sequence.

Proof. Assume the PR^- removes $e = \{u, v\}$ from N to obtain N' . If e is part of a triangle, the PR^- move is an NNI^- move. If e is a parallel edge, then move either u or v with an NNI^0 to obtain a network with a triangle that contains e . Then the previous case applies. So assume otherwise, namely that e is not part of a triangle or a pair of parallel edges. Then move u with an NNI^0 -sequence closer to v to form a triangle as follows.

Because removing e in N yields the proper network N' , it follows that $N \setminus \{e\}$ contains a shortest path P from u to v . Since e is not part of a triangle, this path must contain at least two nodes other than u and v . Let $\{x, y\}$ and $\{y, v\}$ be the last two edges on P . Consider the PR^0 that prunes $\{u, v\}$ at u and regrafts it to $\{x, y\}$. Note that this creates a triangle on the vertices y, u and v . By Lemma 3.4 we can replace this PR^0 with an NNI^0 -sequence. Lastly, we can remove $\{u, v\}$ with an NNI^- to obtain N' . The bound on the length of the NNI -sequence as well as the second statement follow from Lemma 3.4. $\square$

To conclude this section, we note that all previous results combined show that we can replace a TBR-sequence with a PR-sequence, which we can further replace with an NNI -sequence. For several connectedness results in Section 5 this allows us to focus on TBR and then derive results for NNI and PR.

4 Shortest paths

In this section, we focus on bounds on the distance between two specified networks. We restrict to the TBR-distance in $u\mathcal{N}_n$ and in $u\mathcal{N}_{n,r}$, and study the structure of shortest sequences of moves. We make several observations about these sequences in general, and some about shortest sequences between two networks that have certain structure in common, e.g., common displayed networks. Hence, we get bounds on the TBR-distance between two networks, and we uncover properties of the spaces of phylogenetic networks which allow for reductions of the search space. For example, if N and N' have reticulation number r , no shortest path from N to N' contains a network with a reticulation number less than r . The proof of this statement relies on the following observation about the order in which TBR^0 and TBR^+ operations can occur in a shortest path.

Observation 4.1.

Let $N, N' \in u\mathcal{N}_{n,r}$ such that there exists a TBR-sequence $\sigma_0 = (N, M, N')$ that uses a TBR^+ and a TBR^- . Then there is a TBR^0 that transforms N into N' .

Rephrasing Observation 4.1, a TBR^+ followed by a TBR^- , or vice versa, can be replaced by a TBR^0 . This case can thus not occur in a shortest TBR-sequence. Next, we look at a TBR^0 followed by a TBR^+ .

Lemma 4.2.

Let $N, N' \in u\mathcal{N}_n$ with reticulation number r and $r+1$ such that there exists a shortest TBR-sequence $\sigma_0 = (N, M, N')$ that starts with a TBR^0 .

Then there is a TBR-sequence $\sigma_+ = (N, M', N')$ that starts with a TBR^+ .

Proof. Note that the TBR^0 from N to M of σ_0 can be replaced with a sequence consisting of a TBR^+ followed by a TBR^- . This TBR^- and the TBR^+ from M to N' can now be combined to a TBR^0 , which gives us a sequence σ_+ . $\square$

Let $N, N' \in u\mathcal{N}_{n,r}$ and consider a shortest TBR-sequences from N to N' that contains TBR^+ and TBR^- operations. If the reverse statement of Lemma 4.2 would also hold, then we could shuffle the sequence such that consecutive TBR^+ and TBR^- can be replaced with a TBR^0 . This would imply that $u\mathcal{N}_{n,r}$ is an isometric subgraph of $u\mathcal{N}_n$ under TBR. However, we now show that the reverse statement of Lemma 4.2 does not hold in general, and, hence, adjacent operations of different types in a shortest TBR-sequence cannot always be swapped.

Lemma 4.3.

Let $n \geq 4$ and $r \geq 2$. Let $N, N' \in u\mathcal{N}_n$ with reticulation number r and $r + 1$ such that there exists a shortest TBR-sequence $\sigma_+ = (N, M', N')$ that starts with a TBR^+ . Then it is not guaranteed that there is a TBR-sequence $\sigma_0 = (N, M, N')$ that starts with a TBR^0 .

Proof. We claim that the networks N and N' in Figure 8 are a pair of networks for which no TBR-sequence $\sigma_0 = (N, M, N')$ exists that starts with a TBR^0 . The two networks M_1 and M_2 in Figure 8 are the only two TBR^- neighbours of N' . However, it is easy to check that the TBR^0 -distance of N and M_i , $i \in \{1, 2\}$, is at least two. Hence, a shortest TBR sequence from N to N' that starts with a TBR^0 has length three and so σ_0 cannot exist. Note that we can add an edge to each of the pair of parallel edges to obtain an example without parallel edges. Moreover, the example can be extended to higher n and r by adding extra leaves between leaf 3 and 4, and replacing a pair of parallel edges by a chain of parallel edges in each network. $\square$

Figure 8: Two networks $N, N' \in u\mathcal{N}_n$ with TBR-distance two such that there exist a shortest TBR-sequence from N to N' starting with a TBR^+ move (to M'). However, there is no shortest TBR-sequence starting with a TBR^0 , since the networks M_1 and M_2 , which are the TBR^- neighbours of N' , have TBR^0 -distance at least two to N .

Note that the TBR^0 used in Figure 8 to prove Lemma 4.3 is a PR^0 . Hence, the statement of Lemma 4.3 also holds for PR . On the positive side, if one of the two networks is a tree, then we can swap the TBR^+ with the TBR^0 .

Lemma 4.4.

Let $T \in u\mathcal{T}_n$ and $N \in u\mathcal{N}_n$ with reticulation number one such that there exists a shortest TBR-sequence $\sigma_+ = (T, N', N)$ that starts with a TBR^+ . Then there is a TBR-sequence $\sigma_0 = (T, T', N)$ that starts with a TBR^0 .

Proof. We show how to obtain σ_0 from σ_+ . Suppose that N' is obtained from T by adding the edge f and that N is obtained from N' by removing e' and adding e . Note that f is an edge of the cycle C in N' . Furthermore, e' and f are distinct. Indeed, otherwise there would be a shorter TBR-sequence from T to N that simply adds e to T .

Assume for now that e' is an edge of C in N' . Then, e' can be removed with a TBR^- from N' to obtain a tree T' . Hence, the TBR^+ from T to N' and the TBR^- from N' to T' can be merged into a TBR^0 from T to T' . Furthermore, the edge e can then be added to T' with a TBR^+ to obtain N . This yields the sequence σ_0 .

Next, assume that e' is not an edge of C in N' . Then, e' is a cut-edge in N' and e is a cut-edge in N . Let $\bar{e}$ be the edge of T that equals e' , if it exists, or the edge that gets subdivided by f into e' and another edge. Let $\bar{f}$ be the edge of N defined as follows: it is equal to f itself if f is not touched by the TBR^0 move from N' to N ; it is the extension of f if one of its endpoints is suppressed by this move; it is one of the two edges obtained by subdividing f . Now let T' be a tree obtained by removing $\bar{f}$ from N . Then, there is a TBR^0 from T to T' that moves $\bar{e}$ to $\bar{e}'$ and furthermore a TBR^+ that adds $\bar{f}$ to T' and yields N . We obtain again σ_0 . An example is given in Figure 9. $\square$

Figure 9: There is a TBR-sequence from T to N that first adds f with a TBR^+ and then moves e' to e with a TBR^0 . From this, a TBR-sequence can be derived that moves $\bar{e}$ to $\bar{e}'$ with a TBR^0 and then adds $\bar{f}$ with a TBR^+ .

Next, we look at shortest paths between a tree and a network. First, we show that if a network displays a tree, then there is a simple TBR^- -sequence from the network to the tree. Recall that $D(N)$ is the set of trees in $u\mathcal{T}_n$ displayed by $N \in u\mathcal{N}_n$. This result is the unrooted analogous to Lemma 7.4 by Bordewich et al. [BLS17] on rooted phylogenetic networks.

Lemma 4.5.

Let $N \in u\mathcal{N}_{n,r}$ and $T \in u\mathcal{T}_n$.

Then $T \in D(N)$ if and only if $d_{\text{TBR}}(T, N) = r$, that is, iff there exists a TBR^- -sequence of length r from N to T .

Proof. Note that $d_{\text{TBR}}(T, N) \geq r$, since a TBR can reduce the reticulation number by at most one. Furthermore, if we apply a sequence of r TBR^- moves on N , we arrive at a tree that is displayed by N . Hence, if $T \notin D(N)$, then $d_{\text{TBR}}(T, N) > r$.

We now use induction on r to show that $d_{\text{TBR}}(T, N) \leq r$ if $T \in D(N)$. If $r = 0$, then $T = N$ and the inequality holds. Now suppose that $r > 0$ and that the statement holds whenever a network with a reticulation number less than r displays T . Fix an embedding of T into N and colour all edges of N not covered by this embedding green. Note that removing a green edge with a TBR^- might result in an improper network or a loop. Therefore, we have to show that there is always at least one edge that can be removed such that the resulting graph is a phylogenetic network. For this, consider the subgraph H of N induced by the green edges. If H contains a component consisting of a single green edge e , then removing e from N with a TBR^- yields a network N' . If H contains a tree component S , then it is easy to see that removing an external edge of S from N with a TBR^- yields a network N' . Otherwise, as N is proper, a component S displays a tree T_S whose external edges cover exactly the external edges of S . We can then apply the same case distinction to the edges of S not covered by T_S and either directly find an edge to remove or find further trees that cover the smaller remaining components. Since S is finite, we eventually find an edge to remove. The induction hypothesis then applies to N' . This concludes the proof. $\square$

Note that the proof of Lemma 4.5 also works if T is a network displayed by N . Hence, we get the following corollary.

Corollary 4.6.

Let $N \in u\mathcal{N}_{n,r}$ and let $N' \in u\mathcal{N}_{n,r'}$ such that N' is displayed by N .

Then $d_{\text{TBR}}(N', N) = r - r'$, that is, there exists a TBR^- -sequence of length $r - r'$ from N to N' .

Lemma 4.5 and Corollary 4.6 now allow us to construct TBR-sequences between networks that go down tiers and then come up again. In fact, for rooted networks this can sometimes be necessary as Klawitter and Linz have shown [KL19, Lemma 13]. However, we now show that this is never necessary for TBR on unrooted networks.

Lemma 4.7.

Let $N, N' \in u\mathcal{N}_n$.

Then in no shortest TBR-sequence from N to N' does a TBR^- precede a TBR^+ .

Proof. Consider a minimal counterexample with $N, N' \in u\mathcal{N}_n$ such that there exists a shortest TBR-sequence σ from N to N' that uses exactly one TBR^- and TBR^+ and that starts with this TBR^- . If σ uses TBR^0 operations between the TBR^- and the TBR^+ , then, by Lemma 4.2, we can swap the TBR^+ forward until it directly follows the TBR^- . However, then we can obtain a TBR-sequence shorter than σ by combining the TBR^- and the TBR^+ into a TBR^0 by Observation 4.1; a contradiction. $\square$

Combining Lemmas 4.2 and 4.5 and Corollary 4.6, we easily derive the following two corollaries about short sequences that do not go down tiers before going back up again.

Corollary 4.8.

Let $N, N' \in u\mathcal{N}_n$ with reticulation number r and r' , with $r \geq r'$. Then

$$d_{\text{TBR}}(N, N') \leq \min\{d_{\text{TBR}}(T, T') : T \in D(N), T' \in D(N')\} + r.$$

Corollary 4.9.

Let $N, N' \in u\mathcal{N}_n$ with reticulation number r and r' , and $r \geq r'$. Let $T \in u\mathcal{T}_n$ such that $T \in D(N), D(N')$. Then

$$d_{\text{TBR}}(N, N') \leq r.$$

Both Corollaries 4.8 and 4.9 can easily be proven by first finding a sequence that goes down to tier 0 and back up to tier r , and then combining the r' TBR⁻ with r' TBR⁺ into r' TBR⁰ using Lemma 4.2.

The following lemma is the unrooted analogue to Proposition 7.7 by Bordewich et al. [BLS17]. We closely follow their proof.

Lemma 4.10.

Let $N, N' \in \mathcal{uN}_n$ such that $d_{\text{TBR}}(N, N') = k$. Let $T \in D(N)$. Then there exists a $T' \in D(N)$ such that

$$d_{\text{TBR}}(T, T') \leq k.$$

Proof. The proof is by induction on k . If $k = 0$, then the statement trivially holds. Suppose that $k = 1$. If $T \in D(N')$, then set $T' = T$, and we have $d_{\text{TBR}}(T, T') = 0 \leq 1$. So assume otherwise, namely that $T \notin D(N')$. Note that if N' has been obtained from N by a TBR⁺, then N' displays T . Therefore, distinguish whether N' has been obtained from N by a TBR⁰ or TBR⁻ σ .

Suppose that N' has been obtained from N by a TBR⁰ that moves the edge $e = \{u, v\}$ of N . Fix an embedding S of T into N . Since N' does not display T , the edge e is covered by S . Let $\bar{e}$ be the edge of T that gets mapped to the path of S that covers e . Let S_1 and S_2 be the subgraphs of $S \setminus \{e\}$. Note that S_1, S_2 have embeddings into N and N' . Now, if in N there exists a path P from the embedding of S_1 to the embedding of S_2 that avoids e , then the graph consisting of P, S_1 , and S_2 is a tree T' displayed by N' . Otherwise e is a cut-edge of N and the TBR⁰ moves e to an edge e' connecting the two components of $N \setminus \{e\}$. Then in N' there is path P from the embedding of S_1 to the embedding of S_2 in N' . Together they form an embedding of a tree T' displayed by N' . In both cases T' can also be obtained from T by moving $\bar{e}$ to where P attaches to S_1 and S_2 . If N' is obtained from N by a TBR⁻, then the first case has to apply.

Now suppose that $k \geq 2$ and that the hypothesis holds for any two networks with TBR-distance at most $k - 1$. Let $N'' \in \mathcal{uN}_n$ such that $d_{\text{TBR}}(N, N'') = k - 1$ and $d_{\text{TBR}}(N'', N') = 1$. Thus by induction there are trees T'' and T' such that $T'' \in D(N'')$ with $d_{\text{TBR}}(T, T'') \leq k - 1$ and $T' \in D(N')$ with $d_{\text{TBR}}(T'', T') \leq 1$. It follows that $d_{\text{TBR}}(T, T') \leq k$, thereby completing the proof of the lemma. $\square$

By setting one of the two networks in the previous lemma to be a phylogenetic tree and noting that the roles of N and N' are interchangeable, the next two corollaries are immediate consequences of Lemmas 4.5 and 4.10.

Corollary 4.11.

Let $T \in \mathcal{uT}_n$, $N \in \mathcal{uN}_{n,r}$ such that $d_{\text{TBR}}(T, N) = k$. Then for every $T' \in D(N)$

$$d_{\text{TBR}}(T, T') \leq k.$$

Corollary 4.12.

Let $N \in \mathcal{uN}_{n,r}$ and let $T, T' \in D(N)$. Then

$$d_{\text{TBR}}(T, T') \leq r.$$

The following theorem is the unrooted analogous of Theorem 7 by Klawitter and Linz [KL19] and their proof can be applied straightforward by swapping SNPR and rooted networks with TBR and unrooted networks, and by using Lemmas 4.5 and 4.10 and Theorem 6.1.

Theorem 4.13.

Let $T \in \mathcal{uT}_n$ and let $N \in \mathcal{uN}_{n,r}$. Then

$$d_{\text{TBR}}(T, N) = \min_{T' \in D(N)} d_{\text{TBR}}(T, T') + r.$$

5 Connectedness and diameters

Whereas in the previous section we studied the maximum distance between two given networks, here, we focus on global connectivity properties of several classes of phylogenetic networks under NNI, PR, and TBR. These results imply that these operations induce metrics on these spaces. For each connected metric space, we can ask about its diameter. Since a class of phylogenetic networks that contains networks with unbounded reticulation number naturally has an unbounded diameter, this questions is mainly of interest for the tiers of a class. First, we recall some known results from unrooted phylogenetic trees.

Theorem 5.1 (Li et al. [LTZ96], Ding et al. [DGH11]).

The space $\mathcal{uT}_n$ is connected under

- NNI⁰ with the diameter in $\Theta(n \log n)$,
- PR⁰ with the diameter in $n - \Theta(\sqrt{n})$, and
- TBR⁰ with the diameter in $n - \Theta(\sqrt{n})$.

5.1 Network space

Huber et al. [HMW16, Theorem 5] proved that the space of phylogenetic networks that includes improper networks is connected under NNI. We reprove this for our definition of $\mathcal{uN}_n$, but first look at the tiers of this space.

Theorem 5.2.

Let $n \geq 0$, $r \geq 0$, and $m = n + r$.

Then $\mathcal{uN}_{n,r}$ is connected under NNI with the diameter in $\Theta(m \log m)$.

Proof. Let $N \in \mathcal{uN}_{n,r}$ and let $T \in \mathcal{uT}_n$ be a tree displayed by N . We show that N can be transformed into a sorted r -handcuffed caterpillar N^* with $\mathcal{O}(m \log m)$ NNI. Our process is as follows and illustrated in Figure 10.

Step 1. Transform N into a network N_T that is tree-based on T .

Step 2. Transform N_T into handcuffed tree N_H on the leafs 1 and 2.

Step 3. Transform N_H into a sorted handcuffed caterpillar N^* .

We now describe this process in detail. For **Step 1**, we show how to construct an NNI⁰-sequence σ from N to N_T , and we give a bound on the length of σ . Let S be an embedding of T into N , that is, S is a subdivision of T and a subgraph of N . Colour all edges of N used by S black and all other edges green. Note that this yields green, connected subgraphs $G_1, \dots, G_l$ of N ; more precisely, the G_i are the connected components of the graph induced by the green edges of N . Note that each G_i has at least two vertices in S , since otherwise N would not be proper. Furthermore, if each G_i

Figure 10: The process used in the proof of Theorem 5.2. We transform a network N into a tree-based network N_T , then into a handcuffed tree N_H , and finally into a sorted handcuffed caterpillar N^* .

consists of a single edge, then N is tree-based on T . Assuming otherwise, we show how to break the G_i apart.

First, if there is a triangle on vertices v_1, u, v_2 where v_1 and v_2 are adjacent vertices in S and u is their neighbour in G_i , then change the embedding of S (and T) so that it takes the path v_1, u, v_2 instead of v_1, v_2 (see Figure 11a). Otherwise, there is an edge $\{v, u\}$ where v is in S and the other vertices adjacent to u are not adjacent to v . Let $\{u, w_1\}$ and $\{u, w_2\}$ be the other edges incident to u . Apply an NNI^0 to move $\{u, w_1\}$ to S as in Figure 11b. Note that each such NNI^0 decreases the number of vertices in green subgraphs and increases the number of vertices in S . Furthermore, the resulting networks is clearly proper. Therefore, repeat these cases until all G_i consist of single edges. Let the resulting graph be N_T . Since there are at most $2(r-1)$ vertices in all green subgraphs that are not in S , the number of required NNI^0 for Step 1 is at most

$$2(r-1). \quad (1)$$

Figure 11: Transformation and NNI^0 used in Step 1 to obtain a tree-based network N_T .

In **Step 2** we transform N_T into a handcuffed tree N_H on the leaves 1 and 2. Let $M = \{\{u_1, v_1\}, \{u_2, v_2\}, \dots, \{u_r, v_r\}\}$ be the set of green edges in N_T , that is, the edges that are not in the embedding S of T into N_T . Without loss of generality, assume that for $i \in \{1, \dots, r\}$ the distance between u_i and leaf 1 in S is at most the distance of v_i to leaf 1 in S . The idea is to sweep along the edges of S to move the u_i towards leaf 1 and then do the same for the v_i towards leaf 2.

For an edge e of T , let P_e be the path of S corresponding to e . Let e_1 be the edge of T incident to leaf 1. Impose directions on the edges of T towards leaf 1. Do the same for the edges of S accordingly. This gives a partial order $\preceq$ on the edges of T with e_1 as maximum. Let $\prec$ be a linear extension of $\preceq$ on the edges of T .

Let $e = (x, y)$ be the minimum of $\prec$. Let $P_e = (x, \dots, y)$ be the corresponding path in S . From x to y along P_e , proceed as follows.

- (i) If there is an edge (u_i, v_l) in P_e , then swap u_i and v_l with an NNI^0 .
- (ii) If there is an edge (u_i, u_j) in P_e then move the u_j endpoint of the green edge incident to u_j onto the green edge incident to u_i with an NNI^0 .

(iii) Otherwise, if there is an edge (u_i, y) in P_e , then move u_i beyond y .

This is illustrated in Figure 12. Informally speaking, we stack u_j onto u_i so they can move together towards e_1 . Repeat this process for each edge in the order given by $\prec$. For the last edge e_1 , ignore case (iii). Next “unpack” the stacked u_i ’s on e_1 .

We now count the number of NNI^0 needed. Firstly, each v_l is swapped at most once with a u_i . Secondly, each u_j is moving to and from a green edge at most once. Furthermore, each vertex of S corresponding to a vertex of T is swapped at most twice. Hence, the total number of NNI^0 required is at most

$$3r + 2n. \quad (2)$$

Figure 12: NNI^0 used in Step 2 to obtain a handcuffed tree N_H . The label of the moving endpoint follows this endpoint to its regrafting point.

Repeat this process for the v_i towards leaf 2. Since the v_i do not have to be swapped with u_j , the total number of NNI^0 required for this is at most

$$2r + 2n. \quad (3)$$

Note that the resulting network may not yet be a handcuffed tree as the order of the u_i and v_j may be different. Hence, lastly in Step 2, to obtain N_H sort the edges with the mergesort-like algorithm by Li et al. [LTZ96, Lemma 2]. They show that the required number of NNI^0 for this is at most

$$r(1 + \log r). \quad (4)$$

For **Step 3**, consider the path P in S from leaf 1 to 2. If P contains only one pendant subtree, then N_H is handcuffed on the cherry $\{1, 2\}$. Otherwise, use NNI^0 to reduce it to one pendant subtree. This takes at most n NNI^0 . Next, transform the pendant subtree of P into a caterpillar to obtain a handcuffed caterpillar, again with at most n NNI^0 . Lastly, sort the leaves with the algorithm from Li et al. [LTZ96, Lemma 2] to obtain the sorted handcuffed caterpillar N^* . The required number of NNI^0 to get from N_H to N^* is at most

$$2n + n \log n. \quad (5)$$

Since we can transform any network $N \in \mathcal{uN}_{n,r}$ into N^* , it follows that $\mathcal{uN}_{n,r}$ is connected under NNI . Furthermore, adding Equations (1) to (5) up and multiplying the result by two shows that the diameter of $\mathcal{uN}_{n,r}$ under NNI^0 is at most

$$2(6n + 8r + n \log n + r \log r) \in \mathcal{O}((n + r) \log(n + r)). \quad (6)$$

Francis et al. [FHMW18, Theorem 2] gave the lower bound $\Omega(m \log m)$ on the diameter of tier r of the space that allows improper networks under $\text{NNI}^0_{\text{improper}}$ (NNI^0 without the properness condition). Their proof consists of two parts: a lower bound on the total number of networks in a tier $|\mathcal{uN}_{n,r}|$, and upper bounds on the number of networks that

can be reached from one network for each fixed number of $\text{NNI}_{\text{improper}}^0$. The diameter of $u\mathcal{N}_{n,r}$ is at least the smallest number of moves needed for which previously mentioned upper bound is greater than the lower bound on $|u\mathcal{N}_{n,r}|$.

Our version of NNI^0 is stricter than theirs as we do not allow improper networks. Hence, the number of networks that can be reached with a fixed number of NNI^0 is at most the number of networks that can be reached with the same number of $\text{NNI}_{\text{improper}}^0$. Furthermore, their lower bound on $|u\mathcal{N}_{n,r}|$ is found by counting the number of *Echidna* networks, a class of networks only containing proper networks. Combining these two observations, we see that their lower bound for the diameter of $u\mathcal{N}_{n,r}$ under $\text{NNI}_{\text{improper}}^0$ is also a lower bound for $u\mathcal{N}_{n,r}$ under NNI^0 . $\square$

From Theorem 5.2 we get the following corollary.

Corollary 5.3.

The space $u\mathcal{N}_n$ is connected under NNI with unbounded diameter.

Since, by Observation 3.1, every NNI is also a PR and TBR, the statements in Theorem 5.2 and Corollary 5.3 also hold for PR and TBR. This observation has been made before by Francis et al. [FHMW18] for tiers of the space of networks that allow improper networks.

Corollary 5.4.

The spaces $u\mathcal{N}_n$ and $u\mathcal{N}_{n,r}$ are connected under the PR and TBR operation.

We now look at the diameters of $u\mathcal{N}_{n,r}$ under PR and TBR.

Theorem 5.5.

Let $n \geq 0$, $r \geq 0$.

Then the diameter of $u\mathcal{N}_{n,r}$ under PR^0 is in $\Theta(n+r)$ with the upper bound $n+2r$.

Proof. The asymptotic lower bound was proven by Francis et al. [FHMW18, Proposition 4]. Concerning an upper bound, Janssen et al. [JJE⁺18, Theorem 4.22] showed that the distance of two improper networks M and M' under PR is at most $n + \frac{8}{3}r$, of which $\frac{2}{3}r$ PR^0 moves are used to transform M and M' into proper networks N and N' . Hence, the PR-distance of N and N' is at most $n+2r$. $\square$

Theorem 5.6.

Let $n \geq 0$, $r \geq 0$.

Then the diameter of $u\mathcal{N}_{n,r}$ under TBR is in $\Theta(n+r)$ with the upper bound

$$n - 3 - \left\lfloor \frac{\sqrt{n-2} - 1}{2} \right\rfloor + r.$$

Proof. Like for PR, the lower bound was proven by Francis et al. [FHMW18, Proposition 4]. In Corollary 4.8 we show that the TBR-distance of two networks N and $N' \in u\mathcal{N}_{n,r}$ that display a tree T and $T' \in u\mathcal{T}_n$, respectively, is at most $d_{\text{TBR}}(T, T') + r$. Since $d_{\text{TBR}}(T, T') \leq n - 3 - \left\lfloor \frac{\sqrt{n-2} - 1}{2} \right\rfloor$ by Theorem 1.1 of Ding et al. [DGH11] it follows that $d_{\text{TBR}}(N, N') \leq n - 3 - \left\lfloor \frac{\sqrt{n-2} - 1}{2} \right\rfloor + r$. $\square$

5.2 Networks displaying networks

Bordewich [Bor03, Proposition 2.9] and Mark et al. [MMS16] showed that the space of rooted phylogenetic trees that display a set of triplets (trees on three leaves) is connected under NNI. Furthermore, Bordewich et al. [BLS17] showed that the space of rooted phylogenetic networks that display a set of rooted phylogenetic trees is connected. We give a general result for unrooted phylogenetic networks that display a set of networks. For this, we will use Lemma 4.5, which, as we recall, guarantees that if a network $N \in u\mathcal{N}_{n,r}$ displays a tree $T \in u\mathcal{T}_n$, then there is a sequence of r TBR^- from N to T .

Proposition 5.7.

Let $P = \{P_1, \dots, P_k\}$ be a set of k phylogenetic networks P_i on $Y_i \subseteq X = \{1, \dots, n\}$. Then $u\mathcal{N}_n(P)$ is connected under NNI, PR, and TBR.

Proof. Define the network $N_P \in u\mathcal{N}_n(P)$ as follows. Let $P_0 \in u\mathcal{T}_n$ be the caterpillar where the leaves are ordered from 1 to n ; that is, P_0 contains a path $(v_2, v_3, \dots, v_{n-1})$ such that leaf i is incident to v_i , leaf 1 is incident to v_2 , and leaf n is incident to v_{n-1} . Let e_i be the edge incident to leaf i in P_0 . Subdivide e_i with k vertices $u_i^1, \dots, u_i^k$. Now, for $P_j \in P$, $j \in \{1, \dots, k\}$, identify leaf i of P_j with u_i^j of P_0 and remove its label i . Finally, in the resulting network suppress any degree two vertex. This is necessary if one or more of the P_j have fewer than n leaves. The resulting network N_P now displays all networks in P . An example is given in Figure 13.

Figure 13: The canonical network $N_P \in u\mathcal{N}_5$ that displays the set of phylogenetic networks $P = (P_1, P_2)$ with the underlying caterpillar P_0 .

Let $N \in u\mathcal{N}_n(P)$. Construct a TBR-sequence from N to N_P by, roughly speaking, building a copy of N_P attached to N , and then removing the original parts of N . First, add P_0 to N by adding an edge $e = \{v_1, v_2\}$ from the edge incident to leaf 1 to the edge incident to leaf 2 with a TBR^+ . Then add another edge from e to the edge incident to leaf 3, and so on up to leaf n . Colour all newly added edges and the edges incident to the leaves blue, and all other edges red. Note that the blue edges now give an embedding of P_0 into the current network. Now, ignoring all red edges, it is straight forward to add the P_j , $j \in \{1, \dots, k\}$ one after the other with TBR^+ such that the resulting network displays N_P . For example, one could start by adding a tree displayed by P_j and then

adding any other edges. The first part works similar to the construction of P_0 and the second part is possible by Lemma 4.5. Lastly, remove all red edges with TBR^- such that every intermediate network is proper. This is again possible by Lemma 4.5 and yields the network N_P . Note that in the first two stages the red edges (plus external edges) display P and in the last phase the non-red edges display P .

Since we only used TBR^+ and TBR^- operations, the statement also holds for PR. For NNI, by Lemma 3.5 we can replace each of these operations that add or remove an edge e by NNI-sequences that only move and remove or add the edge e . Hence, the statement also holds for NNI. $\square$

For the following corollary, note that a quartet is an unrooted binary tree on four leaves and a quarternet is an unrooted binary, level-1 network on four leaves [HMSW18].

Corollary 5.8.

Let $X = \{1, \dots, n\}$. Let P be a set of phylogenetic trees on X , a set of quartets on X , or a set of quarternets on X . Then $\mathcal{u}\mathcal{N}_n(P)$ is connected under NNI, PR, and TBR.

5.3 Tree-based networks

A related but more restrictive concept to displaying a tree is being tree-based. So, next, we consider the class of tree-based networks. We start with the tiers of $\mathcal{u}\mathcal{TB}_{n,r}(T)$, which is the set of tree-based networks that have the tree T as base tree.

Theorem 5.9.

Let $T \in \mathcal{u}\mathcal{T}_n$. Then the space $\mathcal{u}\mathcal{TB}_{n,r}(T)$ is connected under

- TBR with the diameter being between $\lceil \frac{r}{3} \rceil$ and r ,
- PR with the diameter being between $\lceil \frac{r}{2} \rceil$ and $2r$, and
- NNI with the diameter being in $\mathcal{O}(r(n+r))$.

Proof. We start with the proof for TBR. Let $N, N' \in \mathcal{u}\mathcal{TB}_{n,r}(T)$. Consider embeddings of T into N and N' . Let $S = \{e_1, \dots, e_r\}$ and $S' = \{e'_1, \dots, e'_r\}$ be the set of all edges not covered by this embedding of T in N and in N' . Since N is tree-based, S and S' consist of vertex-disjoint edges. Following the embeddings of T into N and N' , it is straightforward to move each edge e_i with a TBR^0 from N to where e'_i is in N' . In total, this requires at most r TBR^0 . Since every intermediate network is clearly in $\mathcal{u}\mathcal{TB}_{n,r}(T)$, this gives connectedness of $\mathcal{u}\mathcal{TB}_{n,r}(T)$ and an upper bound of r on the diameter. For the lower bound, consider a network M with r pairs of parallel edges and M' without any. Observe that a TBR^0 can break at most three pairs of parallel edges and that only if a pair of parallel edges is removed and attached to two other pairs of parallel edge. Hence, for these particular N and N' we have that $d_{\text{TBR}}(N, N') \geq \lceil \frac{r}{3} \rceil$.

The constructed TBR^0 -sequence for N to N' above can be converted straightforwardly into a PR^0 -sequence from N to N' of length at most $2r$. For the lower bound, let M and M' be as above and note that a PR can break at most two pairs of parallel edges. Hence, $d_{\text{PR}}(M, M') \geq \lceil \frac{r}{2} \rceil$.

By Lemma 3.4, the PR-sequence can be used to construct an NNI-sequence from N to N' that only moves the edges e_i along paths of the embedding of T . Since the PR-sequence has length at most $2r$ and each PR can be replaced by an NNI sequence of length at most $\mathcal{O}(n+r)$, this gives the upper bound of $\mathcal{O}(r(n+r))$ on the diameter of $\mathcal{u}\mathcal{TB}_{n,r}(T)$ under NNI. $\square$

We use Theorem 5.9 to prove connectedness of other spaces of tree-based networks.

Theorem 5.10.

Let $T \in u\mathcal{T}_n$.

Then the spaces $u\mathcal{TB}_n(T)$, $u\mathcal{TB}_{n,r}$, and $u\mathcal{TB}_n$ are each connected under TBR, PR, and NNI. Moreover, the diameter of $u\mathcal{TB}_{n,r}$ is in $\Theta(n+r)$ under TBR and PR and in $\mathcal{O}(n \log n + r(n+r))$ under NNI.

Proof. Assume without loss of generality that T has the cherry $\{1, 2\}$. First, let N and N' be in tiers r and r' of $u\mathcal{TB}_n(T)$, respectively, such that they are r - and r' -handcuffed on the cherry $\{1, 2\}$. Then $d_{\text{NNI}}(N, N') = |r' - r|$, as we can decrease the number of handcuffs with NNI^- . Since, by Theorem 5.9, the tiers of $u\mathcal{TB}_{n,r}(T)$ are connected, the connectedness of $u\mathcal{TB}_n(T)$ follows.

Second, let $N, N' \in u\mathcal{TB}_{n,r}$ be tree-based networks on T and T' respectively, and with an r -burl on the edge incident to leaf 1. Ignoring the burls, by Theorem 5.1, N can be transformed into N' by transforming T into T' with $\mathcal{O}(n \log n)$ NNI^0 or with $\mathcal{O}(n)$ PR^0 or TBR^0 . With Theorem 5.9, the connectedness of $u\mathcal{TB}_{n,r}$ and the upper bounds on the diameter follow. The lower bound on the diameter under PR and TBR also follows from Theorem 5.1 and Theorem 5.9,

Lastly, the connectedness of $u\mathcal{TB}_n$ follows similarly from the connectedness of $u\mathcal{T}_n$ and $u\mathcal{TB}_{n,r}$. $\square$

5.4 Level- k networks

To conclude this section, we prove the connectedness of the space of level- k networks.

Theorem 5.11.

Let $n \geq 2$ and $k \geq 1$.

Then, the space $u\mathcal{LV}\text{-}k_n$ is connected under TBR and PR with unbounded diameter.

Proof. Let $N \in u\mathcal{LV}\text{-}k_n$ and $T \in u\mathcal{T}_n$. We show that N can be transformed into the network $M \in u\mathcal{LV}\text{-}k_n$ that can be obtained from T by adding a k -burl to the edge incident to leaf 1. First, create a k -burl in N on the edge incident to leaf 1. This can be done using k PR^+ . Next, using Lemma 4.5 remove all other blobs. This gives a network M' which consists of a tree T' with a k -burl at leaf 1. There is a PR^0 -sequence from T' to T , which is easily converted into a sequence from M' to M . This proves the connectedness of $u\mathcal{LV}\text{-}k_n$ under PR and also TBR. Lastly, note that the diameter is unbounded because the number of possible reticulations in a level- k network is unbounded. $\square$

Note that an NNI^+ cannot directly create a pair of parallel edges. We may instead add a triangle with an NNI^+ and then use an NNI^0 to transform it into a pair of parallel edges. However, adding the triangle within a level- k blob of a level- k network, then adding the triangle would increase the level. Therefore, to prove connectedness of level- k networks under NNI we use the same idea as for PR but are more careful to not increase the level.

Theorem 5.12.

Let $n \geq 3$ and $k \geq 1$.

Then, the space $u\mathcal{LV}\text{-}k_n$ is connected under NNI with unbounded diameter.

Proof. Let $N \in u\mathcal{LV}\text{-}k_n$ and let $T \in u\mathcal{T}_n$. Like in the proof of Theorem 5.11, we want to transform N into a network M obtained from T by adding a k -burl to the edge incident to leaf 1.

Let B be a level- k blob of N . Assume that N contains another blob B' . By Lemma 4.5 there is a PR^+ -sequence that can remove B' . Use Lemma 3.5 to substitute this sequence with an NNI-sequence that reduces B' to a level-1 blob. Note that this can be done locally within blob B' and its incident edges. Therefore, this process does not increase the level of a network along this sequence. If B' is now a cycle of size at least three, then we can shrink it to a triangle, if necessary, and remove it with an NNI^- . If B' is a pair of parallel edges and one of its vertices is incident to a degree three vertex v that is not part of a level- k blob, then use an NNI^0 to increase the size of B' into a triangle by including v or merge it with the blob containing v . Next, either remove the resulting triangle, or repeat the process above to remove the new blob. Otherwise, ignore B' for now and continue with another blob of the current network that is neither B' nor B . When this process terminates, we arrive at a network that has only blob B , and, potentially, pairs of parallel edges that are incident to both B and a leaf. That is the case since a pair of parallel edges incident to a degree three vertex not in B could be removed with an NNI^0 and an NNI^- .

If the edge incident to leaf 1 contains a pair of parallel edges or is incident to a degree three vertex not in B , then use $k - 1$ NNI^+ and NNI^0 (or k in the latter case) to create a k -burl next to leaf 1. Otherwise, if B is incident to three or more cut-edges, then one of them is not incident to leaf 1 and can be moved to the edge incident to leaf 1 with an NNI^0 -sequence. If B is incident to two or fewer cut-edges, there is a vertex incident to three cut edges (since $n \geq 3$) and one of them can be moved to the edge incident to leaf 1 with an NNI^0 -sequence. Then apply the first case again to create a k -burl. Finally, remove B and any remaining pair of parallel edges. This gives a network M' which consists of a tree T' with a k -burl at leaf 1. There is an NNI^0 -sequence from T' to T , which is easily converted into a sequence from M' to M . Lastly, note that the diameter is unbounded because for each $r \geq 0$, there is a level- k network with r reticulations. $\square$

6 Isometric relations between spaces

Recall that a space $\mathcal{C}_n$ is an isometric subgraph of $u\mathcal{N}_n$ under a rearrangement operation, say TBR, if the TBR-distance of two networks in $\mathcal{C}_n$ is the same as their TBR-distance in $u\mathcal{N}_n$. In this section, we investigate this question for $u\mathcal{T}_n$ under TBR, and for tree-based networks and level- k networks under TBR and PR.

We start with $u\mathcal{T}_n$. The proof of the following theorem follows the proof by Bordewich et al. [BLS17, Proposition 7.1] for their equivalent statement for SNPR on rooted phylogenetic trees and networks closely.

Theorem 6.1.

The space $u\mathcal{T}_n$ is an isometric subgraph of $u\mathcal{N}_n$ under TBR. Moreover, every shortest TBR-sequence from $T \in u\mathcal{T}_n$ to $T' \in u\mathcal{T}_n$ only uses TBR^0 .

Proof. Let $d_{\mathcal{T}}$ and $d_{\mathcal{N}}$ be the TBR-distance in $u\mathcal{T}_n$ and $u\mathcal{N}_n$ respectively. To prove the statement, it suffices to show that $d_{\mathcal{T}}(T, T') = d_{\mathcal{N}}(T, T')$ for every pair $T, T' \in u\mathcal{T}_n$. Note that $d_{\mathcal{T}}(T, T') \geq d_{\mathcal{N}}(T, T')$ holds by definition. To prove the converse, let $\sigma = (T = N_0, N_1, \dots, N_k = T')$ be a shortest TBR-sequence from T to T' . Consider the following

colouring of the edges of each N_i , for $i \in \{0, \dots, k\}$. Colour all edges of $T = N_0$ blue. For $i \in \{1, \dots, k\}$ preserve the colouring of N_{i-1} to a colouring of N_i for all edges except those affected by the TBR. In particular, an edge that gets added or moved is coloured red, an edge resulting from a vertex suppression is coloured blue if the two merged edges were blue and red otherwise, and the edges resulting from an edge subdivision are coloured like the subdivided edge.

Let F_i be the graph obtained from N_i by removing all red edges. We claim that F_i is a forest with at most $k + 1$ components. Since $F_0 = T$, the statement holds for $i = 0$. If N_i is obtained from N_{i-1} by a TBR^+ , then $F_i = F_{i-1}$. If N_i is obtained from N_{i-1} by a TBR^0 or TBR^- , then at most one component gets split. Note that F_k is a so-called agreement forest for T and T' and thus $d_{\mathcal{T}}(T, T') \leq k = d_{\mathcal{N}}(T, T')$ by Theorem 2.13 by Allen and Steel [AS01]. Furthermore, if σ would use a TBR^+ , then the forest F_k would contain at most k components. However, then $d_{\mathcal{T}}(T, T') < k$; a contradiction. $\square$

Francis et al. [FHMW18] gave the example in Figure 14 to show that the tiers $u\mathcal{N}_{n,r}$ for $n \geq 5$ and $r > 0$ are not isometric subgraphs of $u\mathcal{N}_n$ under NNI. Their question of whether tier zero, $u\mathcal{T}_n$, is an isometric subgraph of $u\mathcal{N}_n$ under NNI remains open.

Lemma 6.2.

Let $n \geq 5$ and $r \geq 0$. Then the space $u\mathcal{N}_{n,r}$ is not an isometric subgraph of $u\mathcal{N}_n$ under NNI.

Figure 14: An NNI-sequence from N to N' using an NNI^+ that adds f , an NNI^0 that moves e , and an NNI^- that removes e' . A shortest NNI^0 -sequence from N to N' has length three.

Lemma 6.3.

For $n = 4$ and $r = 13$ the space $u\mathcal{N}_{n,r}$ is not an isometric subgraph of $u\mathcal{N}_n$ under PR.

Proof. For the networks N and N' in $u\mathcal{N}_{n,r}$ shown in Figure 15 there is a length three PR-sequence that traverses tier $r + 1$, for example, like the depicted sequence $\sigma = (N = N_0, N_1, N_2, N_3 = N')$. To prove the statement we show that every PR^0 -sequence from N to N' has length at least four.

The networks N and N' contain the highlighted (sub)blobs B_1, B_2 , (resp. B'_1 and B'_2), B_3 , and B_4 . Observe that the edges between B_1 and B_2 and between B_3 and B_4 may only be pruned from a blob by a PR^0 if they get regrafted to the same blob again. Otherwise the resulting network is improper. Note that to derive B'_1 from B_1 an edge has to be regrafted to the “top” of B_1 and the edge to B_2 has to be pruned. By the first observation, combining these into one PR^0 cannot build the connection to B_3 . The same applies for the transformation of B_2 into B'_2 and its connection to B_4 . Therefore, we either need four PR^0 to derive B'_1 and B'_2 or two PR^0 plus two PR^0 to build the connections to B_3 and B_4 . In conclusion, at least four PR^0 are required to transform N into N' , which concludes this proof. $\square$

Figure 15: A length three PR-sequence from N to N' that uses a PR^+ , which adds f , a PR^0 , which moves e , and a PR^- , which removes e' . A PR^0 -sequence from N to N' has length at least four.

By replacing a leaf with a tree, and adding more pairs of parallel edges to edge leading to 4, this example can be made to work for $n \geq 4$ and $r \geq 13$.

Theorem 6.4.

For $n \geq 6$ the space $u\mathcal{TB}_n$ is not an isometric subgraph of $u\mathcal{N}_n$ under TBR and PR.

Proof. Let N be the network in Figure 16. Let N' be the network derived from N by swapping the labels 1 and 2. Note that $d_{TBR}(N, N') = d_{PR}(N, N') = 2$, since, from N to N' , we can move leaf 2 next to leaf 1 and then move leaf 1 to where leaf 2 was. However, then the network in the middle is not tree-based, since the blob derived from the Petersen graph has no Hamiltonian path if the two pendent edges of the blob are next to each other [FHM18]. We claim that there is no other length two TBR-sequence from N to N' . For this proof we call a blob derived from the Petersen graph a Petersen blob.

Figure 16: A tree-based network on the left and a Hamiltonian path through a blob derived from the Petersen graph on the right.

First, note that the TBR^0 -sequence of N and N' is at least two and there is thus no TBR -sequence that consists of a TBR^- and a TBR^+ . Otherwise, these two operations could be merged into a single TBR^0 by Observation 4.1. Note that we can only move leaf 1 or 2 by pruning an incident edge if we do not affect the split 1 versus 2, 3 or break the tree-based property. Therefore, they either have to be swapped using edges of the Petersen blobs or the $(4, 5, 6)$ -chain has to be reversed and leaf 3 moved to the other Petersen blob. However, it is straightforward to check that neither can be done with two TBR^0 . In particular, we can look at what edge the first TBR^0 might move and then check whether a second TBR^0 can arrive at N' . If the first TBR^0 breaks a Petersen blob, the problem is that the second TBR^0 has to restore it. We then find that this does not allow us to make the initially planned changes to arrive at N' . On the other hand, if we avoid breaking the Petersen blob and reverse the $(4, 5, 6)$ -chain, then leaf 3 is still on the wrong side; and if we move leaf 3 to the other Petersen blob, then not enough TBR^0 moves remain to reverse the chain.

Since there is no other length two TBR^0 -sequence there is also no other length two PR-sequence. $\square$

Theorem 6.5.

For $n \geq 5$ and large enough k , the space $u\mathcal{LV}\text{-}k_n$ is not an isometric subgraph of $u\mathcal{N}_n$ under TBR and PR.

Proof. For even k , the networks N and N' in Figure 17 have TBR- and PR-distance two via the network M . However, note that in M the blobs of size $\frac{k}{2} + 1$ and $\frac{k}{2}$ are merged into a blob of size $k + 1$. Therefore, M is not a level- k network. We claim that there is no TBR- or PR-sequence of length two that does not go through a level- $(k + 1)$ network like M . An example for odd k can be derived from this.

Figure 17: For even k , a PR^0 -sequence from a level- k network N to a level- k network N' (hidden reticulations of the blob-parts given inside, at least two leaves omitted: in B_1 and in B_3). However, the network M in the middle is a level- $(k + 1)$ but not a level- k network.

It is easy to see that the TBR-distance of N and N' is at least two and there is thus no TBR-sequence that consists of a TBR^- and a TBR^+ . Otherwise, these two operations could be merged into a single TBR^0 by Observation 4.1. We thus have to prove that there is no length two TBR^0 -sequence from N to N' that avoids a level- $(k + 1)$ network. Note that it requires two TBR^0 (or PR^0) to connect B_2 and B_3 into B'_2 . Similarly, it requires either two prunings from the upper five-cycle of B_2 to obtain the triangle B'_3 or one pruning within that cycle. However, in the latter option this would not contribute to connecting B_2 and B_3 and hence overall at least three operations would be needed. Therefore we have to combine the two operations necessary to create B'_2 and to create B'_3 , which however gives us a sequence like the one shown in Figure 17. $\square$

Note that the results of this section that show that the spaces of tree-based networks and level- k networks are not isometric subgraphs of the space of all networks also hold if we restrict these spaces to a particular tier r (for large enough r).

7 Computational complexity

In this section, we consider the computational complexity of computing the TBR-distance and the PR-distance. First, we recall the known results on phylogenetic trees.

Theorem 7.1 ([DHJ⁺97, HDRCB08, AS01]).

Computing the distance of two trees in $u\mathcal{T}_n$ is NP-hard for the NNI-distance, the SPR-distance, and the TBR-distance.

In Theorem 6.1, we have shown that $u\mathcal{T}_n$ is an isometric subgraph of $u\mathcal{N}_n$ under TBR. Hence, with Theorem 7.1, we get the following corollary.

Corollary 7.2.

Computing the TBR-distance of two arbitrary networks in $u\mathcal{N}_n$ is NP-hard.

We can use the same two theorems to prove that computing the TBR-distance in tiers is also hard.

Theorem 7.3.

Computing the TBR-distance of two arbitrary networks in $u\mathcal{N}_{n,r}$ is NP-hard.

Proof. We (linear-time) reduce the NP-hard problem of computing the TBR-distance of two trees in $u\mathcal{T}_n$ to computing the TBR-distance of two networks in $u\mathcal{N}_{n+1,r}$. For this, let $T, T' \in u\mathcal{T}_n$. Let e be the edge incident to leaf n of T . Obtain S from T by subdividing e with a new vertex u and adding the edge $\{u, v\}$ where v is a new vertex labelled $n+1$. Next, add r handcuffs to the cherry $\{n, n+1\}$ to obtain the network $N \in u\mathcal{N}_{n+1,r}$. Analogously obtain N' from T' .

The equality $d_{\text{TBR}}(T, T') = d_{\text{TBR}}(N, N')$ follows from Lemma 4.10, and the fact that networks handcuffed at a cherry display exactly one tree. More precisely, a TBR-sequence between T and T' induces a TBR-sequence of the same length between N and N' , hence $d_{\text{TBR}}(T, T') \geq d_{\text{TBR}}(N, N')$. Conversely, by Lemma 4.10 and the fact that $D(N) = \{T\}$ and $D(N') = \{T'\}$, it follows that $d_{\text{TBR}}(T, T') \leq d_{\text{TBR}}(N, N')$. Since computing the TBR-distance in $u\mathcal{T}_n$ is NP-hard, the statement follows. $\square$

To prove that computing the PR-distance is hard, we use a different reduction. Van Iersel et al. prove that deciding whether a tree is displayed by a (not necessarily proper) phylogenetic network (Unrooted Tree Containment; UTC) is NP-hard [VIKS⁺18]. Combining this with Lemma 4.5, we arrive at our result.

Theorem 7.4.

Computing the PR-distance of two arbitrary networks in $u\mathcal{N}_n$ is NP-hard.

Proof. We reduce from UTC to the problem of computing the PR-distance of two networks in $u\mathcal{N}_n$. Let (N, T) with N a (not necessarily proper) network and $T \in u\mathcal{T}_n$ be an arbitrary instance of UTC. We obtain an instance (N', T', r') of the PR-distance decision problem as follows: remove all cut-edges of N that do not separate two labelled leaves, and let N'' be the connected component containing all the leaves; now, let N' be the proper network obtained from N'' by suppressing all degree two nodes. The instance of the PR-distance decision problem consists of N' , $T' = T$, and the reticulation number r' of N' . As we can compute in polynomial time whether a cut edge separates two labelled leaves, the reduction is polynomial time. Because a displayed tree uses only cut-edges that separate two labelled leaves, T is displayed by N if and only if it is displayed by N' . By Lemma 4.5, T is a displayed tree of N , if and only if $d_{\text{PR}}(N', T') \leq r$, which concludes the proof. $\square$

Unlike for the hardness proof of TBR-distance, we cannot readily adapt this proof to the PR-distance in $u\mathcal{N}_{n,r}$. For this purpose, we need to learn more about the structure of PR-space.

8 Concluding remarks

In this paper, we investigated basic properties of spaces of unrooted phylogenetic networks and their metrics under the rearrangement operations NNI, PR, and TBR. We have proven connectedness and bounds on diameters for different classes of phylogenetic networks, including networks that display a particular set of trees, tree-based networks, and level- k networks. Although these parameters have been studied before for classes of rooted phylogenetic network [BLS17], this is the first paper that studies these properties for classes of unrooted phylogenetic networks besides the space of all networks. A summary of our results is shown in Table 1.

To see the improvements in diameter bounds, we compare our results to previously found bounds: For the space of phylogenetic trees $u\mathcal{T}_n$ it was known that the diameter is asymptotically linearithmic and linear in the size of the trees under NNI and SPR/TBR [LTZ96, DGH11], respectively. Here, we have shown that the diameter under NNI is also asymptotically linearithmic for higher tiers of phylogenetic networks. Whether this also holds in the rooted case is still open. We have further (re)proven the asymptotic linear diameter for PR and TBR of these tiers and, in particular, improved the upper bound on the diameter under TBR to $n - 3 - \lfloor \frac{\sqrt{n-2}-1}{2} \rfloor + r$ from the previously best bound $n + 2r$ [JJE⁺18].

class	NNI	PR	TBR
$u\mathcal{T}_n$	$\Theta(n \log n)$ [LTZ96]	$\Theta(n)$ [DGH11]	$\Theta(n)$ [DGH11]
$u\mathcal{N}_{n,r}$	$\Theta(m \log m)$ T. 5.2	$\Theta(m)$ [FHM18, JJE ⁺ 18]	$\Theta(m)$ T. 5.6
$u\mathcal{N}_n$	✓ Corollary 5.3	✓ Corollary 5.4	✓ Corollary 5.4
$u\mathcal{N}_n(P)$	✓ Proposition 5.7	✓ Proposition 5.7	✓ Proposition 5.7
$u\mathcal{TB}_{n,r}(T)$	$\mathcal{O}(rm)$ Theorem 5.9	$\Theta(r)$ Theorem 5.9	$\Theta(r)$ Theorem 5.9
$u\mathcal{TB}_{n,r}$	$\mathcal{O}(rm + n \log n)$ T. 5.10	$\Theta(m)$ Theorem 5.10	$\Theta(m)$ T. 5.10
$u\mathcal{TB}_n(T)$	✓ Theorem 5.10	✓ Theorem 5.10	✓ Theorem 5.10
$u\mathcal{TB}_n$	✓ Theorem 5.10	✓ Theorem 5.10	✓ Theorem 5.10
$u\mathcal{LV}\text{-}k_n$	✓ Theorem 5.12	✓ Theorem 5.11	✓ Theorem 5.11

Table 1: Connectedness and diameters, if bounded, for the various classes and rearrangement operations. Here $m = n + r$, P is a set of phylogenetic networks, and $T \in u\mathcal{T}_n$.

To uncover local structures of network spaces, we looked at properties of shortest sequences of moves between two networks. Here we found that shortest TBR-sequences between networks in the same tier never traverse lower tiers, and shortest TBR-sequences between trees also never traverse higher tiers. This implies that $u\mathcal{T}_n$ is an isometric subgraph of $u\mathcal{N}_n$, and that computing the TBR-distance between two networks in $u\mathcal{N}_n$ is NP-hard. This answers a question by Francis et al. [FHMW18]. We have attempted to prove similar results for other subspaces and rearrangement moves. However, for higher tiers, we have not been able to prove that shortest TBR-sequences never traverse higher tiers. To answer this question we may need to utilise agreement graphs such as frequently used for phylogenetic trees [AS01, BS05] and, more recently, also for rooted phylogenetic networks [KL19, Kla19]. Concerning NNI and PR we gave counterexamples to prove that higher tiers are not isometric subgraphs of $u\mathcal{N}_n$. The questions whether $u\mathcal{T}_n$ is isometrically embedded in $u\mathcal{N}_n$ under PR and NNI remains open. Answering these

questions positively would also provide an answer to the question whether computing the shortest NNI-distance between two networks is NP-hard, and clues toward proving whether the PR-distance between two networks in the same tier is NP-hard. Further negative results that we have shown are that the spaces of tree-based networks and level- k are not isometric subgraphs of the space of all phylogenetic networks.

Throughout this paper, we have restricted our attention to proper networks. We could also have chosen to use unrooted networks without the properness condition. This definition, which is mathematically more elegant, is used in most other papers, so it seems to be the obvious choice. However, it is not natural to have cut-edges that do not separate leaves: such networks carry no biological meaning. It is desirable that networks are rootable and thus have an evolutionary interpretation. Unrooted phylogenetic networks are rootable if they have at most one blob with one cut-edge. While using this in the definition of an unrooted phylogenetic network could therefore be sufficient, we go one step further, and ask that there is no such blob. This makes a network rootable at any leaf (i.e., with any taxon as out-group), which gives a stronger biological interpretation and usability.

The fact that our definition of unrooted phylogenetic networks is mathematically more restrictive, means that any positive result we have proven is likely also true when using a less restrictive definition. That is, connectedness for those definitions follows easily by finding sequences to proper networks, like done by Jansen et al. [JJE⁺18]. As we may be able to find short sequences for this purpose, the diameter results will likely also still hold. This means that whatever definitions may be used in practice, with minor additional arguments, our results provide the theoretical background necessary to justify local search operations.

References

- [AS01] B. L. Allen and M. Steel, “Subtree transfer operations and their induced metrics on evolutionary trees,” *Annals of Combinatorics*, vol. 5, no. 1, pp. 1–15, 2001. doi:10.1007/s00026-001-8006-8
- [BLS17] M. Bordewich, S. Linz, and C. Semple, “Lost in space? Generalising subtree prune and regraft to spaces of phylogenetic networks,” *Journal of Theoretical Biology*, vol. 423, pp. 1–12, 2017. doi:10.1016/j.jtbi.2017.03.032
- [Bor03] M. Bordewich, “The complexity of counting and randomised approximation,” Ph.D. dissertation, University of Oxford, 2003. <http://community.dur.ac.uk/m.j.r.bordewich/papers/Bordewich2003-a.pdf>
- [BS05] M. Bordewich and C. Semple, “On the computational complexity of the rooted subtree prune and regraft distance,” *Annals of Combinatorics*, vol. 8, no. 4, pp. 409–423, 2005. doi:10.1007/s00026-004-0229-z
- [DGH11] Y. Ding, S. Grünewald, and P. J. Humphries, “On agreement forests,” *Journal of Combinatorial Theory, Series A*, vol. 118, no. 7, pp. 2059–2065, 2011. doi:10.1016/j.jcta.2011.04.013
- [DHJ⁺97] B. DasGupta, X. He, T. Jiang, M. Li, J. Tromp, and L. Zhang, “On distances between phylogenetic trees,” in *Proceedings of the 8. annual ACM-SIAM Symposium on Discrete Algorithms*, 1997, pp. 427–436.

- [Die17] R. Diestel, *Graph Theory*, 5th ed. Springer Berlin Heidelberg, 2017. doi:10.1007/978-3-662-53622-3
- [FHM18] A. Francis, K. T. Huber, and V. Moulton, “Tree-based unrooted phylogenetic networks,” *Bulletin of Mathematical Biology*, vol. 80, no. 2, pp. 404–416, 2018. doi:10.1007/s11538-017-0381-3
- [FHMW18] A. Francis, K. T. Huber, V. Moulton, and T. Wu, “Bounds for phylogenetic network space metrics,” *Journal of Mathematical Biology*, vol. 76, no. 5, pp. 1229–1248, Apr 2018. doi:10.1007/s00285-017-1171-0
- [FS15] A. R. Francis and M. Steel, “Which phylogenetic networks are merely trees with additional arcs?” *Systematic Biology*, vol. 64, no. 5, pp. 768–777, 2015. doi:10.1093/sysbio/syv037
- [Gus14] D. Gusfield, *ReCombinatorics: the algorithmics of ancestral recombination graphs and explicit phylogenetic networks*. MIT Press, 2014.
- [GvIJ⁺17] P. Gambette, L. van Iersel, M. Jones, M. Lafond, F. Pardi, and C. Scornavacca, “Rearrangement moves on rooted phylogenetic networks,” *PLOS Computational Biology*, vol. 13, no. 8, pp. 1–21, 08 2017. doi:10.1371/journal.pcbi.1005611
- [HDRCB08] G. Hickey, F. Dehne, A. Rau-Chaplin, and C. Blouin, “SPR distance computation for unrooted trees,” *Evolutionary Bioinformatics*, vol. 4, p. EBO.S419, 2008. doi:10.4137/EB0.S419
- [HLMW16] K. T. Huber, S. Linz, V. Moulton, and T. Wu, “Spaces of phylogenetic networks from generalized nearest-neighbor interchange operations,” *Journal of Mathematical Biology*, vol. 72, no. 3, pp. 699–725, 2016. doi:10.1007/s00285-015-0899-7
- [HMSW18] K. T. Huber, V. Moulton, C. Semple, and T. Wu, “Quarnet inference rules for level-1 networks,” *Bulletin of Mathematical Biology*, vol. 80, no. 8, pp. 2137–2153, 2018. doi:10.1007/s11538-018-0450-2
- [HMW16] K. T. Huber, V. Moulton, and T. Wu, “Transforming phylogenetic networks: Moving beyond tree space,” *Journal of Theoretical Biology*, vol. 404, pp. 30–39, 2016. doi:10.1016/j.jtbi.2016.05.030
- [HRS10] D. H. Huson, R. Rupp, and C. Scornavacca, *Phylogenetic networks: concepts, algorithms and applications*. Cambridge University Press, 2010.
- [Jan18] R. Janssen, “Heading in the right direction? Using head moves to traverse phylogenetic network space,” *arXiv preprint arXiv:1810.09795*, 2018. arXiv:1810.09795
- [JJE⁺18] R. Janssen, M. Jones, P. L. Erdős, L. van Iersel, and C. Scornavacca, “Exploring the tiers of rooted phylogenetic network space using tail moves,” *Bulletin of Mathematical Biology*, vol. 80, no. 8, pp. 2177–2208, 2018. doi:10.1007/s11538-018-0452-0

- [KL19] J. Klawitter and S. Linz, “On the Subnet Prune and Regraft Distance,” *Electronic Journal of Combinatorics*, vol. 26, no. 2, pp. 329–355, 2019. <https://www.combinatorics.org/ojs/index.php/eljc/article/view/v26i2p3>
- [Kla19] J. Klawitter, “The agreement distance of rooted phylogenetic networks,” *Discrete Mathematics & Theoretical Computer Science*, vol. 21, no. 3, 2019. doi:10.23638/DMTCS-21-3-19
- [LTZ96] M. Li, J. Tromp, and L. Zhang, “Some notes on the nearest neighbour interchange distance,” in *Computing and Combinatorics*, J.-Y. Cai and C. K. Wong, Eds. Springer Berlin Heidelberg, 1996, pp. 343–351. doi:10.1007/3-540-61332-3_168
- [MMM⁺17] J. I. Meier, D. A. Marques, S. Mwaiko, C. E. Wagner, L. Excoffier, and O. Seehausen, “Ancient hybridization fuels rapid cichlid fish adaptive radiations,” *Nature Communications*, vol. 8, p. 14363, 2017. doi:10.1038/ncomms14363
- [MMS16] S. Mark, J. C. McLeod, and M. Steel, “A navigation system for tree space,” *Journal of Graph Algorithms and Applications*, vol. 20, no. 2, pp. 247–268, 2016. doi:10.7155/jgaa.00392
- [Rob71] D. F. Robinson, “Comparison of labeled trees with valency three,” *Journal of Combinatorial Theory, Series B*, vol. 11, no. 2, pp. 105–119, 1971. doi:10.1016/0095-8956(71)90020-7
- [RW07] L. H. Rieseberg and J. H. Willis, “Plant Speciation,” *Science*, vol. 317, no. 5840, pp. 910–914, 2007. doi:10.1126/science.1137729
- [SJ17] K. St. John, “Review paper: The shape of phylogenetic treespace,” *Systematic Biology*, vol. 66, no. 1, pp. e83–e94, 2017. doi:10.1093/sysbio/syw025
- [SOW96] D. L. Swofford, G. J. Olsen, and P. J. Waddell, “Phylogenetic inference,” in *Molecular Systematics*, D. M. Hillis, C. Moritz, and B. K. Mable, Eds. Sinauer Associates, Sunderland, MA, 1996, ch. 11, pp. 407–514.
- [SS03] C. Semple and M. Steel, *Phylogenetics*. Oxford University Press on Demand, 2003, vol. 24.
- [Ste16] M. Steel, *Phylogeny: discrete and random processes in evolution*. Society for Industrial and Applied Mathematics, 2016.
- [TN05] C. M. Thomas and K. M. Nielsen, “Mechanisms of, and Barriers to, Horizontal Gene Transfer between Bacteria,” *Nature Reviews Microbiology*, vol. 3, no. 9, pp. 711–721, 2005. doi:10.1038/nrmicro1234
- [VIKS⁺18] L. Van Iersel, S. Kelk, G. Stamoulis, L. Stougie, and O. Boes, “On unrooted and root-uncertain variants of several well-known phylogenetic network problems,” *Algorithmica*, vol. 80, no. 11, pp. 2993–3022, 2018. doi:10.1007/s00453-017-0366-5

- [WWK⁺17] A. Wagner, R. J. Whitaker, D. J. Krause, J.-H. Heilers, M. van Wolferen, C. van der Does, and S.-V. Albers, “Mechanisms of gene flow in archaea,” *Nature Reviews Microbiology*, vol. 15, no. 8, pp. 492–502, 2017. doi:10.1038/nrmicro.2017.41