

Transforming phylogenetic networks: Moving beyond tree space

Katharina T. Huber^a, Vincent Moulton^a, Taoyang Wu^{a,*}

^a *School of Computing Sciences, University of East Anglia, Norwich, NR4 7TJ, UK*

Abstract

Phylogenetic networks are a generalization of phylogenetic trees that are used to represent reticulate evolution. Unrooted phylogenetic networks form a special class of such networks, which naturally generalize unrooted phylogenetic trees. In this paper we define two operations on unrooted phylogenetic networks, one of which is a generalization of the well-known nearest-neighbor interchange (NNI) operation on phylogenetic trees. We show that any unrooted phylogenetic network can be transformed into any other such network using only these operations. This generalizes the well-known fact that any phylogenetic tree can be transformed into any other such tree using only NNI operations. It also allows us to define a generalization of tree space and to define some new metrics on unrooted phylogenetic networks. To prove our main results, we employ some fascinating new connections between phylogenetic networks and cubic graphs that we have recently discovered. Our results should be useful in developing new strategies to search for optimal phylogenetic networks, a topic that has recently generated some interest in the literature, as well as for providing new ways to compare networks.

Keywords:

phylogenetic network, local transformation, network space, NNI operation, network metric

*Corresponding author

Email addresses: katharina.huber@cmp.uea.ac.uk (Katharina T. Huber),
vincent.moulton@cmp.uea.ac.uk (Vincent Moulton), taoyang.wu@gmail.com (Taoyang Wu)

1. Introduction

Phylogenetic networks are a generalization of phylogenetic trees that are gaining growing acceptance by biologists due to their importance in representing reticulate evolution (Bapteste et al., 2013). Certain types of networks, such as neighbornets (see e.g. Bryant and Moulton, 2004; Huson and Bryant, 2006) and median networks (see e.g. Bandelt et al., 1995) are now commonly used in the literature. Moreover, there has recently been much focus on developing ways to construct special classes of networks to explicitly model evolution (see e.g. Huson et al., 2010; Nakhleh, 2011; Gusfield, 2014). Even so, there are still several aspects of phylogenetic network theory that remain to be more fully explored. One such aspect is how to transform one network into another one by using a collection of specified network operations (see e.g. Cardona et al., 2009; Yu et al., 2014; Huber et al., 2016, for some results in this direction) which we consider in this paper.

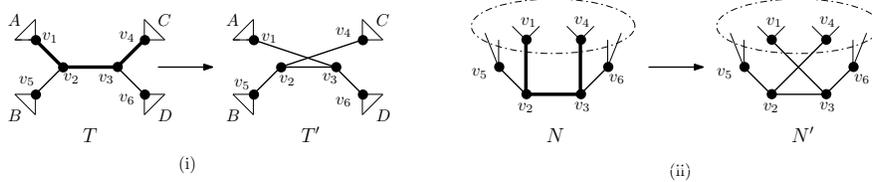


Figure 1: (i) An NNI operation on a phylogenetic tree T . The tree T' is obtained from T by performing an operation on the path highlighted in bold that results in the subtrees labelled A and C being swapped. (ii) An NNI operation on a network N . The network N' is obtained from N by one NNI operation that is performed on the path highlighted in bold, just like for phylogenetic trees. Note that vertices v_1, v_4, v_5 and v_6 in N could all have degree 1 or 3, and that N contains neither edge $\{v_1, v_3\}$ nor edge $\{v_2, v_4\}$.

Transformations of phylogenetic trees have been studied for several years, and have applications to tree search algorithms and comparing trees (cf. e.g. Felsenstein, 2004, Chapters 4 and 30). Probably the best known and simplest way to transform one phylogenetic tree into another is to use a *nearest-neighbor interchange (NNI) operation* which we now recall. For a set X of three or more

species or taxa, a *phylogenetic tree (on X)* is a tree in which every vertex has degree 1 or 3 with leaf set X . A pair of distinct trees differ by one NNI operation if one tree can be obtained from the other by swapping two of the four subtrees adjacent with an interior edge (Figure 1(i)). Note that the NNI operation is *reversible*, i.e. there is a unique NNI operation (or *reverse* operation) that can be applied to get back to the original tree. A well-known result concerning the NNI operation states that given any pair of phylogenetic trees T, T' on X , it is possible to transform T into T' by some sequence of NNI operations (Robinson, 1971). This implies that NNI operations can be used to explore all possible phylogenetic trees on a set X , a useful fact that underpins several algorithms for reconstructing phylogenetic trees (see, e.g. Felsenstein, 2004).

We now turn to the analogous problem for phylogenetic networks. More specifically, we shall consider the problem of transforming *unrooted phylogenetic binary networks (on X)*, or *networks* for short, into one another. Such networks are connected graphs in which every vertex has degree 1 or 3 and whose leaf set is X (Gambette et al., 2012). Here we concentrate on binary networks since these generalise binary trees which are commonly employed within tree-inference methods (see, e.g. Chapter 4, Felsenstein, 2004) and, moreover, multifurcations may be problematic for phylogenetic inference (Whelan and Money, 2010). In Figure 2, we present an example of such a network R that was referred to as the “ring of life” in a study concerning the genome fusion origin of eukaryotes (Figure 1 in Rivera and Lake, 2004). Based on that network, whose construction employed whole genomes spanning the diversity of life, the authors concluded that the eukaryotic nuclear genome has resulted from a fusion of a relative of a protobacterium (P_γ) and a relative of an archeal ecocyte (E). Properties of networks have been studied in Gambette et al. (2012), and they can be generated by software such as T-REX (cf. Makarenkov, 2001) and Splitstree (cf. Huson and Bryant, 2006; Huson et al., 2005). In this paper, we show that it is possible to extend results concerning NNI operations on phylogenetic trees in a natural way to networks, which we expect could lead to applications to network search algorithms and comparison of networks.

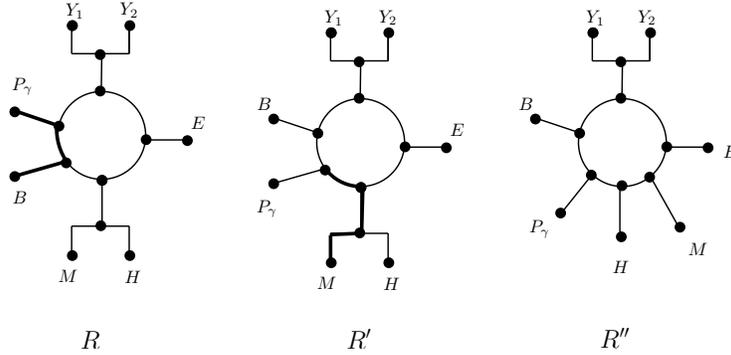


Figure 2: An example of transforming phylogenetic networks using NNI operations. Network R is from Rivera and Lake (2004, Figure 1), and is labelled by two yeasts (Y_1 , *Schizosaccharomyces pombe* and Y_2 , *Saccharomyces cerevisiae*), a γ -proteobacterium (P_γ , *Xylella fastidiosa*), a bacillus (B , *Staphylococcus aureus* MW2), a halobacterium (H , *Halobacterium* sp. NRC-1), an eocyte (E , *Sulfolobus tokodaii*) and a methanococcus (M , *Methanosarcina mazei* Goe1). The networks R' and R'' are obtained from R and R' , respectively, by applying an NNI operations to the path highlighted in bold.

Our starting point is to extend NNI operations to networks. This is based on the observation that the tree T' in Figure 1(i) can be obtained from T by replacing the length 3 path v_1, v_2, v_3, v_4 highlighted in bold with the path v_1, v_3, v_2, v_4 whilst preserving all other edges. This definition immediately extends to give what we shall call an NNI operation on networks. In particular, we start again with a path v_1, v_2, v_3, v_4 in a network N on X for which neither $\{v_1, v_3\}$ nor $\{v_2, v_4\}$ is an edge, and obtain a new network N' on X by replacing this path with the path v_1, v_3, v_2, v_4 (Figure 1(ii)). Note that N' has the same number of vertices as the original network, and that, just as with phylogenetic trees, the NNI operation is reversible.

In the first of our main results we show that, just as with phylogenetic trees, we can transform any network N on X to any other network N' on X with the same number of vertices as N by just using NNI operations (Theorem 3). We illustrate this in Figure 2: here network R is transformed into the network

R'' – which can be considered as an alternative ring of life hypothesis – by applying a sequence of two NNI operations. Even so, it is not possible to explore all alternative ways to potentially represent the ring of life using only NNI operations. To see this, note that we cannot transform the tree T in Figure 4 to the network R in Figure 2 using only NNI operations since, even though these networks are on the same leaf set, they have a different number of vertices and NNI operations must preserve this number.

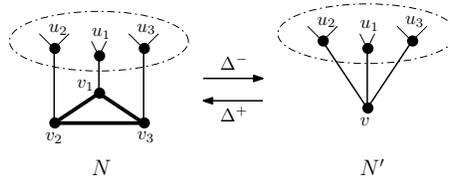


Figure 3: The two considered triangle operations: N' is obtained from N by one Δ^- operation that collapses the triangle v_1, v_2, v_3 highlighted in bold to the vertex v , preserving all other edges that are not contained in this triangle. Conversely, N is obtained from N' by one Δ^+ operation which replaces the vertex v by the triangle v_1, v_2, v_3 . Note that in both networks vertices u_1, u_2, u_3 are all distinct and each of them could have degree 1 or 3.

In our second main result (Theorem 4) we show that we need only one additional operation and its reverse to be able to transform any network on X into any other network on X . We call these operations *triangle operations* or Δ operations and picture them in Figure 3; they involve either inserting or removing a triangle (or, more technically, a length three cycle) from a network. We illustrate this result in Figure 4. The tree T in this figure is one of the phylogenetic trees presented in Rivera and Lake (2004, Figure 1). It is one of the five most probable trees computed using the method of conditioned reconstruction. The ring of life network R can be obtained from T by applying one Δ^+ operation and two NNI operations as illustrated in Figure 4.

We now summarize the rest of the paper. After presenting some preliminaries concerning graphs in the next section, we show in Theorems 1 and 2 that it is possible to construct any network starting either from a tree, or from a rather simple multigraph with just two vertices (called a generator) using a sequence

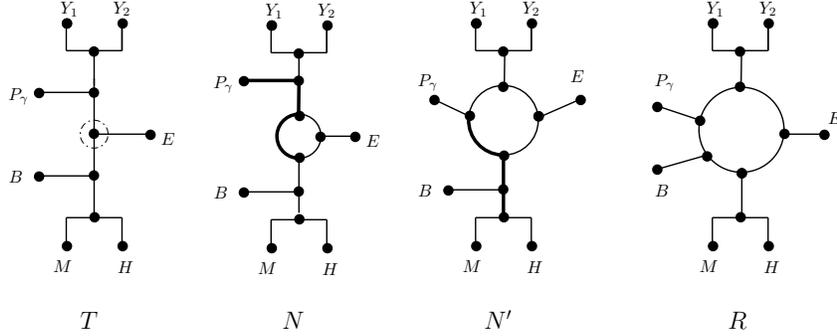


Figure 4: Transforming a phylogenetic tree into a network using NNI and Δ operations. The tree T appeared in Rivera and Lake (2004, Figure 1), and R is the same network as in Figure 2. Network N is obtained from T by a Δ^+ operation on the vertex highlighted in the dash circle. Networks N' and R are obtained from N and N' , respectively, by applying an NNI operation to the path highlighted in bold.

of operations. In Section 4, we then present the two main results stated above, which we prove using some new connections between networks and the theory of so-called cubic graphs that we discovered as part of this work. Using our results, in Section 5 we also define a generalization of tree space, and see how this can be used to give some new metrics that can be used to compare networks. We conclude in Section 6 with a brief discussion of future directions and open problems, including some potential consequences for network searching algorithms.

2. Preliminaries

We begin with some technical preliminaries concerning graphs, which give the underlying structure for phylogenetic networks. We assume throughout the paper that X is a set of species or taxa with three or more elements.

A *multigraph* G is an ordered pair $G = (V, E)$, with V a set of vertices and E a multiset of edges each being an unordered pair of elements in V . Note that we

assume that multigraphs do not contain loops (i.e. edges that connect vertices to themselves) unless otherwise stated, but they may contain *parallel edges* (i.e. edges with multiplicity greater than 1). If G contains no parallel edges we shall call it a *graph*. We also assume that all multigraphs are connected. For example, in Figure 8(i) we picture a connected multigraph with two vertices and three parallel edges.

Assume for the remainder of this section that G is a multigraph. An edge whose removal disconnects G is called a *bridge* of G ; if G contains no such edge, then G is called *bridgeless*. A *blob* in G is a maximal connected induced submultigraph of G that does not contain a bridge. For example, the cycle in the network R in Figure 2 is a blob and so is the subgraph of the network in Figure 7(i) with all labelled vertices and their adjacent edges removed. A triangle in G is called *isolated* if it does not share an edge with any other triangle in G , in which case the triangle is a blob of G . We say that G is *level- k* , $k \geq 0$ an integer, if a tree can be obtained from G by removing at most k edges from each blob of G . In this case, each blob of G is at most level- l , $0 \leq l \leq k$.

The proofs of our main results are closely related to some theorems in the theory of cubic graphs. We say that G is *cubic* if every vertex in G has degree three. Note that a cubic multigraph necessarily has an even number of vertices. Following Brinkmann et al. (2013), we call a multigraph G a *1-3 graph* if all vertices in G have degree either 1 or 3. These graphs naturally arise when studying phylogenetic trees as every vertex in such a tree has degree 1 or 3. Note that in a cubic or a 1-3 graph an isolated triangle cannot share a vertex with any other triangle in that graph and any parallel edges must have multiplicity two except in the (necessarily unique) level-2 cubic graph.

Finally we introduce the concept **of a generator** since, as we shall see later, they can be viewed as the building blocks of networks. A *generator* is either (i) an element of X , (ii) a loop with a single vertex, or (iii) a bridgeless, cubic multigraph. Note that this is an unrooted equivalent for the definition of a generator for directed networks presented in Gambette et al. (2009). In case of (i) and (ii), they are the (unique) generator with level-0 and level-1, respectively.

Also, for any integer $k \geq 2$, a generator is level- k if and only if it has $2(k - 1)$ vertices.

3. Constructing networks

Our first results concern the construction of networks by using simple operations. We shall first show that it is possible to construct any network starting from a phylogenetic tree using just NNI and Δ operations. This is a key step in proving that we can transform any network into any other using just NNI and Δ operations. Note that the T-REX algorithm mentioned in the introduction (Makarenkov, 2001) also builds networks from trees, but it does this somewhat differently by adding in new edges to a pre-computed tree that optimize **an objective** function.

Although somewhat technical, we include a proof for the first construction result since it nicely illustrates the type of considerations that are required to show that we can transform one network into another using only NNI and Δ operations. For an illustration of this theorem see Figure 4 (which we consider in reverse!).

Theorem 1. *Suppose that N is a graph such that X ($|X| \geq 3$) is contained in its vertex set. Then N is a network on X if and only if there is a phylogenetic tree T on X and a finite sequence of NNI and Δ^+ operations that can be applied one after the other, starting with T , to obtain N .*

Proof. If T is a phylogenetic tree on X and N is obtained from T by a sequence of NNI and Δ^+ operations, then it is straightforward to see that N is a network on X since applying each of these operations to a network results in a network.

Conversely, since each of the Δ^- and NNI operations preserve the leaf set of a network, it suffices to show that we can reduce any 1-3 graph to a binary tree by using a sequence of these two operations.

To prove this, we adapt the proof presented in Batagelj (1981, p. 94-95) for showing that every connected cubic graph different from a 4-clique (i.e., a

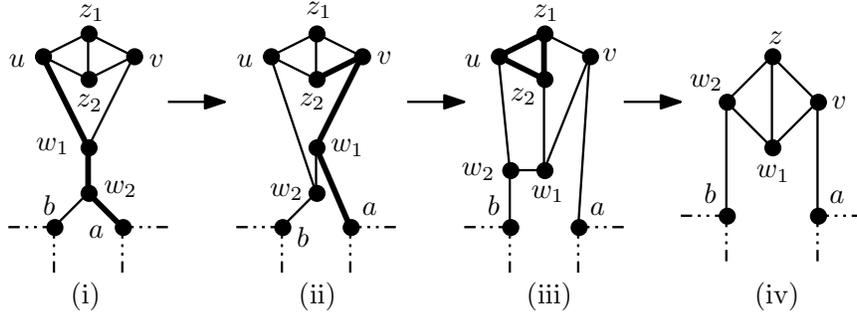


Figure 5: Case II, 2.1 in Theorem 1. (i) and (ii): The edges in bold indicate the path to which an NNI operation is applied. (iii) A Δ^- operation is applied to the vertices of the triangle whose edges are given in bold. (iv) Graph obtained from (iii); the vertex labeled z is the newly generated vertex.

graph with 4 vertices and all of the 6 possible edges) can be reduced (in terms of number of vertices) using the rules P1 and P7 considered in that paper.

Suppose that N is a 1-3 graph which is not a tree. Let C be a shortest cycle in N .

Case I: C is not a triangle.

In this case (and as exemplified by performing the NNI operation on network R'' that reverses the one that transforms R' to R'' in Figure 2), we can reduce the number of vertices of C by 1, by applying an NNI operation to a path v_1, v_2, v_3, v_4 where v_1 is not a vertex in C (and hence $\{v_1, v_2\}$ is not an edge in C) and $\{v_2, v_3\}$ and $\{v_3, v_4\}$ are two edges in C . Note that this is possible since $\{v_1, v_3\}$ and $\{v_2, v_4\}$ cannot be edges in N as this would contradict the minimality of C . We can therefore repeat this application of the NNI operation until we are in Case II.

Case II: C is a triangle.

Case II, 1: C is isolated, and so we can apply the Δ^- operation.

Case II, 2: C has an edge in common with another triangle.

Let u, v be the two (necessarily distinct) vertices in the two triangles that are not contained in the common edge. Let z_1 and z_2 denote the vertices adjacent with that edge.

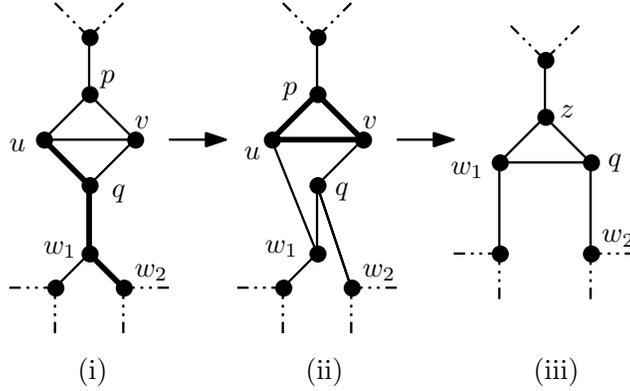


Figure 6: Case II, 2.2 in Theorem 1. (i) The edges in bold indicate the path to which an NNI operation is applied. (ii) A Δ^- operation is applied to the vertices of the triangle whose edges given in bold. (iii) The graph obtained from (ii); the vertex labeled z denotes again the newly generated vertex.

Note that $\{u, v\}$ is not an edge in N since otherwise N would be isomorphic to a 4-clique, which is not possible as G has at least three degree 1 vertices (since $|X| \geq 3$).

We distinguish now between two cases: u, v have a neighbor distinct from z_1 and z_2 in common, or not.

Case II, 2.1: Let $w_1 \notin \{z_1, z_2\}$ be the vertex which is a common neighbor of u and v . Note that the vertex w_2 adjacent to w_1 that is not u or v must have degree 3, otherwise N would have only one vertex with degree 1 as it is a 1-3 graph which is impossible. We can now perform the operations in Figure 5 (just as in Batagelj (1981, Case II/2.1)).

Case II, 2.2: Let p and q be the vertices that are adjacent to u and v , respectively. Then at least one of p or q must have degree 3, otherwise N would have only two vertices with degree 1 implying that X has two elements which is impossible. Without loss of generality, we may assume that the degree of q is 3. Now, we can reduce N as in Figure 6 (just as in Batagelj (1981, Case II/2.2)). \square

We now describe a second approach to constructing networks which is based

on generators. Although we do not require this result later on, it is of independent interest and it also illustrates how other operations could be potentially used for transforming networks.

Before proceeding, as an illustration of the concept of generators, in Figure 7(i) we picture a network that was referred to as the “rings of life” in (Lake and Sinsheimer, 2013). The authors interpreted the bottom right ring as a central system that contains the root of life and connects the two other rings. As we can see this network has a more complicated internal structure than the ring of life network R in Figure 2. This structure is further revealed if we remove all the pendant labelled edges and suppress any remaining degree 2 vertices. In this case we obtain the level-3 generator which is pictured in the far right of Figure 7.

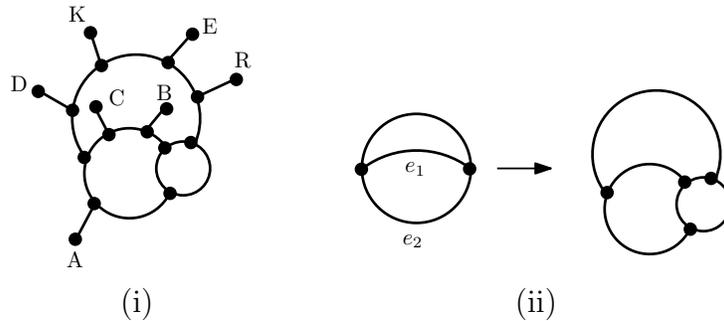


Figure 7: (i) The rings of life from Lake and Sinsheimer (2013) on eukaryotes (K), double-membrane (Gram-negative) prokaryotes (D), eocytes (E), the *Euryarchaeota* (R), *Actinobacteria* (A), and the *Firmicutes* (the *Clostridia*, C, and the *Bacilli*, B). Note that by removing all pendant edges in this network and suppressing degree 2 vertices we obtain a level-3 generator. (ii) The level-3 generator underlying the network in (i) can be constructed by applying one H operation to the edges e_1 and e_2 of the theta graph on the left of Figure 7(ii).

We now explain how level- k generators, $k \geq 2$, can be constructed starting with the multigraph in Figure 8(i) using just the two simple operations in Figure 8(ii) which we call the H and L operations, respectively. We illustrate the H operation in Figure 7(ii), where it is used to build the level-3 generator underlying the rings of life network pictured in Figure 7(i). This could be useful

in exhaustively exploring alternative network hypotheses such as the rings of life example in Figure 7. Note that the H operation, which connects two edges by adding an additional edge, was introduced in 1891 by J. de Vries (see, e.g. Brinkmann et al., 2011), and the L operation, which is applied to one edge, was considered in Brinkmann et al. (2013, Figure 4) where it was called operation 3.

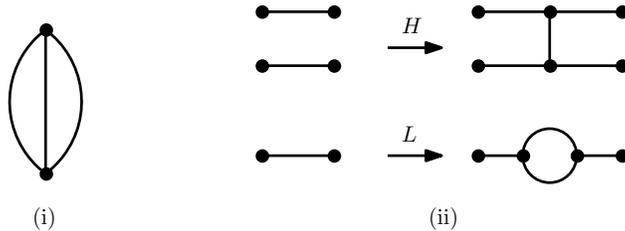


Figure 8: (i) The (necessarily unique) cubic multigraph with two vertices (called theta graph in (Brinkmann et al., 2013, Figure 5)). (ii) The H and L operations on cubic multigraphs. Note that the two edges on which H operates are also allowed to have one or two vertices in common (e.g. they could be parallel).

The following result implies that any level- k generator, $k \geq 2$, can be constructed using just H and L operations. It can be regarded as an analogue of the approach given in Gambette et al. (2009, Section 2) for constructing generators for rooted networks. As we will not require this result later, we give its proof in Appendix A.

Theorem 2. *Suppose that G is a multigraph and $k \geq 2$. Then G is a level- k generator if and only if G can be constructed from the theta graph in Figure 8(ii) via performing a sequence of $(k - 2)$ L or H operations.*

We remark that using this result it is in principle possible to explicitly construct (and therefore directly count) all level- k generators. For example, it can be checked in this way that there is one level-1 generator, two level-3 generators, and six level-4 generators. However, this could probably be done more efficiently using the results in Brinkmann et al. (2013).

We now briefly explain how generators can be used to construct networks; full details are given in Appendix B. The construction is based on the fact that if N is a network and G is a blob in N , then G is either a level-0 generator (i. e. a vertex), or the graph obtained from G by suppressing all degree 2 vertices is a level- k generator, $k \geq 1$ (see Lemma 6). It follows that we can build any network N from the collection of generators corresponding to the blobs of N by gluing together the generators one-by-one along new edges onto a growing network until we obtain N . We illustrate this process in Figure 9.

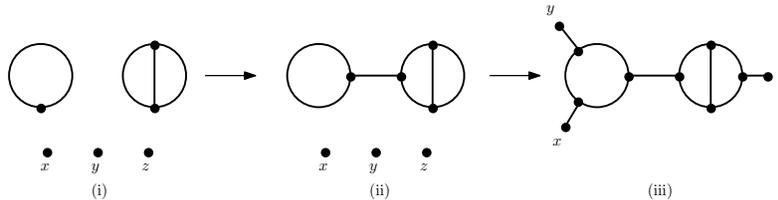


Figure 9: (i) A collection of generators. (ii) Top: a cubic multigraph obtained by joining the level-1 and level-2 generators in (i). (iii) A phylogenetic network obtained by attaching the level-0 generators to the cubic multigraph in (ii).

4. Transforming phylogenetic networks

We now turn our attention to the other main results of this paper. In the first result we see that somewhat surprisingly we can transform any network to any other network with the same number of vertices using only NNI operations. For an illustration of this result, see Figure 2.

Theorem 3. *Suppose that N, N' are two networks on X ($|X| \geq 3$) with the same number of vertices. Then N can be transformed into N' by a finite sequence of NNI operations.*

Proof. Suppose that N, N' are networks on X . We put $N \sim N'$ if, starting with N , there is a sequence of NNI operations (possibly of length 0) that can be

applied to obtain N' . Note that \sim is an equivalence relation on the set $\mathcal{N}(X)$ of all networks on X .

Now, suppose N, N' are networks on X with vertex sets $V(N)$ and $V(N')$, respectively, such that with $|V(N)| = |V(N')|$. The proof that $N \sim N'$ must hold is almost identical to that in Tsukui (1996, Theorem II). However, we shall present a sketch proof for the sake of completeness.

As in the proof of Tsukui (1996, Theorem II) we use three observations which can be proven in essentially the same way as the analogous results for cubic graphs as indicated in brackets.

Observation 1 (cf. Tsukui, 1996, Lemma 4.2): If N is a network on X that is not a phylogenetic tree, then $N \sim N'$ for N' a network on X which contains an isolated triangle.

It is straightforward to check that Observation 1 holds by employing an argument similar to that used in the proof of Theorem 1 Case I (note that as in that theorem we require $|X| \geq 3$ for this to hold). The argument presented in Tsukui (1996) is similar but no reference is given to Batagelj (1981).

Observation 2 (cf. Tsukui, 1996, Lemma 4.3): If N is a network on X , and $u, v \in (V(N) - X)$, then $N^u \sim N^v$, for the networks N^u and N^v obtained from N by performing a Δ^+ operation at u and v , respectively.

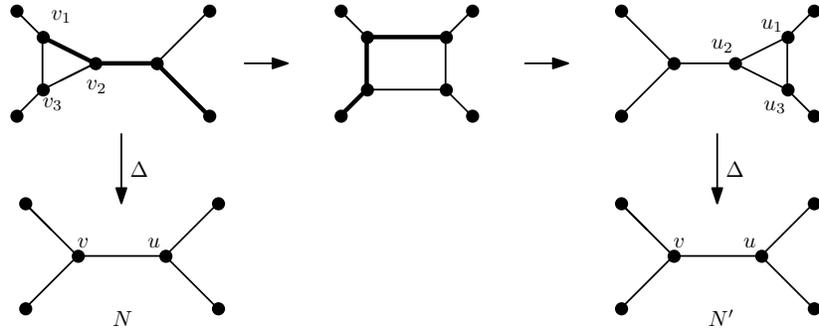


Figure 10: An illustration of the proof of Observation 2 in Theorem 3. The edges in bold indicate the paths to which NNI operations are applied. The Δ operations are applied to the triangles with vertices v_1, v_2, v_3 and u_1, u_2, u_3 , respectively.

To see that Observation 2 holds, it suffices to assume that u is adjacent to v ; in this situation the operation indicated in Figure 10 (adapted from Tsukui, 1996, Figure 12) can be applied to see that $N^u \sim N^v$ holds.

Observation 3 (cf. Tsukui, 1996, Lemma 4.4): Suppose that M is a network on X with an isolated triangle, and N is obtained by performing a Δ^- operation on the vertices of this triangle. If N' is obtained by performing one NNI operation on N , then there is a network M' on X with $M' \sim M$, such that N' can be obtained by performing a Δ^- operation on M' .

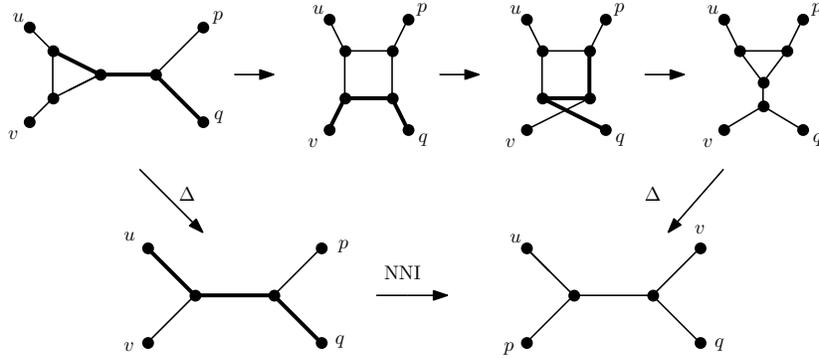


Figure 11: An illustration of the proof of Observation 3 in Theorem 3. The same conventions as for Figure 10 apply.

To see that Observation 3 holds, assume that the NNI operation performed on N to obtain N' is on the path v_1, v_2, v_3, v_4 , and that w is the vertex in N obtained by applying the Δ^- operation to the vertices of an isolated triangle in M . It is straightforward to see that Observation 3 holds in case $w \neq v_2, v_3$. In case $w = v_2$ (or $w = v_3$) the way to obtain the network M' is indicated in Figure 11 (adapted from Tsukui, 1996, Figure 14).

Using these three observations we now prove that $N \sim N'$, using induction on $m = |V(N)| = |V(N')|$. We shall use the fact that if N is a network on X , then $|V(N)| \geq 2|X| - 2$, with equality holding if and only if N is a phylogenetic tree on X .

First note that if $|V(N)| = 2|X| - 2$, then both N and N' are trees, and so

$N \sim N'$ (Robinson, 1971). So, suppose that $M \sim M'$ holds for all networks M and M' on X with $|V(M)| = |V(M')| = m \geq 2|X| - 2$ holding for some m . Assume that N, N' are two networks on X such that $|V(N)| = |V(N')| = m + 2$.

By Observation 1, we can assume that N, N' have isolated triangles Γ and Γ' , respectively. Let P, P' be the networks obtained by performing an Δ^- operation on N and N' on the vertices of the triangles Γ and Γ' , respectively.

By induction, there exists a sequence $P = P_0, P_1, \dots, P_q = P'$ with network P_{i+1} obtained from P_i by an NNI operation, $i = 0, 1, \dots, q - 1$.

Now, by Observation 3, it follows that there is a sequence N_0, N_1, \dots, N_q so that $N_0 = N$, N_i has an isolated triangle Γ_i with $\Gamma_0 = \Gamma$, $N_i \sim N_{i+1}$, and P_i is the network obtained by performing an Δ^- operation on the vertices of the triangle Γ_i in N_i . In particular, $N_q \sim N$ and the Δ^- operation applied to the vertices of the triangle Γ_q in N_q gives the network $P_q = P'$.

But then there must exist vertices u, v in $V(P_q) - X$ so that the networks N_q and N' can be obtained from P_q by performing a Δ^+ operation at u and v , respectively. By Observation 2, $N \sim N'$ follows, as required. \square

Using the last theorem we can now prove our second main result, namely that we can transform any network on X to any other network on X by just using NNI and Δ operations.

Theorem 4. *Suppose that N and N' are networks on X with $|X| \geq 3$. Then, starting with N , there exists a finite sequence of NNI and Δ operations that can be applied to obtain N' .*

Proof. By Theorem 1, there are phylogenetic trees T, T' on X that can be obtained from networks N, N' on X , respectively, by applying a sequence of NNI and Δ operations. Therefore the theorem follows from the fact that there is a sequence of NNI operations that can be applied, starting with T , to obtain T' , as shown in Robinson (1971). \square

5. Network spaces

For phylogenetic trees, NNI operations lead naturally to the concept of *tree space* (see e.g. Billera et al., 2001). In its simplest form, for X fixed, this space can be thought of as the set of all phylogenetic trees on X with two trees being neighbors if and only if they differ by one NNI operation. In other words, tree space is a graph with vertices being the phylogenetic trees on X and edges corresponding to pairs of such trees which differ by one NNI operation. As a consequence of the fact that any phylogenetic tree on X can be transformed into any other by a sequence of NNI operations, it follows that this graph is connected. This is important for optimization algorithms that search tree space using NNI operations since it implies that, in principle, every tree could be visited.

Using the operations introduced in this paper, we can define network space in a similar manner. Indeed, we can consider a graph with vertices consisting of all networks on X and edges corresponding to pairs of networks which differ by either one NNI operation or one Δ operation. Note that this space is well defined since NNI and Δ operations are both reversible. We illustrate part of this network space for $X = \{a, b, c, d\}$ in Figure 12.

By Theorem 4 it follows that network space is connected, just as with tree space. In fact, tree space is actually a subgraph of network space, as it consists of all of the level-0 networks (i.e. phylogenetic trees) and the NNI operation is just the usual NNI operation for trees. Moreover, Theorem 3 shows that network space has some additional structure since, for any non-negative integer i , the set $\mathcal{N}_i(X)$ of all networks on X with $2(|X| - 1 + i)$ vertices forms a “tier” in the space in the sense that any pair of networks in that tier can be joined by edges corresponding only to NNI operations (Figure 12). Moreover, two tiers $\mathcal{N}_i(X)$ and $\mathcal{N}_j(X)$ are adjacent if and only if $|j - i| = 1$, that is, the difference between the number of vertices for the networks in tier $\mathcal{N}_i(X)$ and that in tier $\mathcal{N}_j(X)$ is exactly two. In particular, there exists a network N in $\mathcal{N}_i(X)$ and a network N' in $\mathcal{N}_j(X)$ such that N and N' differ by precisely one Δ operation.

Note that partitioning network space into tiers is not equivalent to dividing up the set of networks according to their level. This alternative way to partition rooted networks has been considered in Yu et al. (2014) and also in Huber et al. (2016) where an alternative space of level-1 networks is defined.

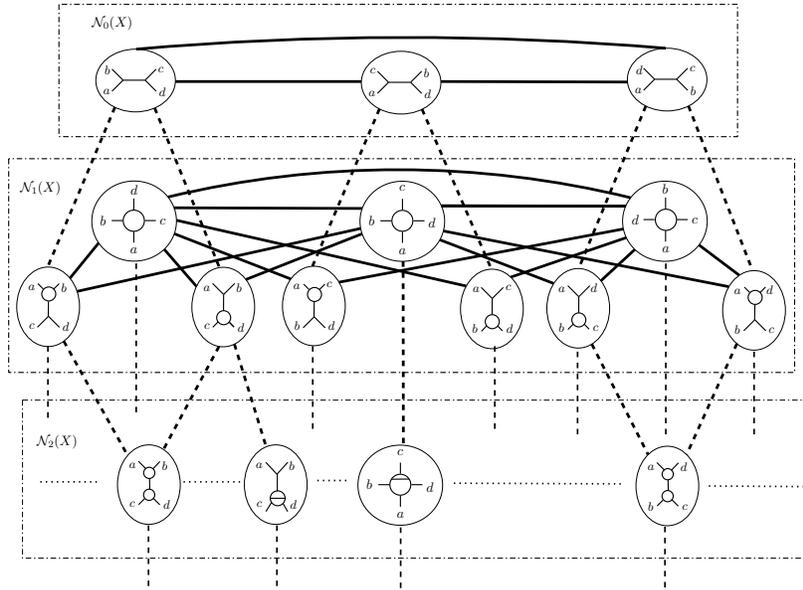


Figure 12: Part of the space of phylogenetic networks on the leaf set $\{a, b, c, d\}$. Two networks are connected by a solid line if one can be transformed into the other by one NNI operation, and a dashed line if this can be done by one Δ operation. This gives rise to the tiers $\mathcal{N}_i(X)$ described in the text.

Tree spaces have been extensively studied in the literature, and they have various important properties, many of which have direct bearing on the development and performance of tree search algorithms. As an example of this, tree space naturally gives rise to the well-known NNI metric on trees which is the shortest path between any pair of trees in tree space, or equivalently, the least number of NNI operations required to transform one tree into another (cf. DasGupta et al., 1997, and the references therein). As a consequence of our results we can define new metrics on the set $\mathcal{N}(X)$ of all networks on X which arise

from network space in a similar way.

More specifically, for $N, N' \in \mathcal{N}_i(X)$ we define the distance $d_{\text{NNI}}(N, N')$ between N and N' to be the minimal number of NNI operations required to transform N into N' , and for $N, N' \in \mathcal{N}(X)$ we let the distance $d(N, N')$ between N and N' be the minimal number of Δ and NNI operations needed to transform N into N' . By Theorems 3 and 4 we immediately obtain the following result.

Theorem 5. *Let X be a finite set with $|X| \geq 3$. The distance d is a metric on the set $\mathcal{N}(X)$. In addition, for any non-negative integer i , the distance d_{NNI} is a metric on the set $\mathcal{N}_i(X)$.*

It would be interesting to better understand properties of these metrics such as ways to compute or find bounds on the size of the distance between any pair of networks. Note that the NNI distance between a pair of phylogenetic trees is NP-hard to compute in general (DasGupta et al., 1997), and so as tree space is contained in network space it is unlikely that there will be a polynomial time algorithm to compute the metric d in general. However, understanding properties of d could yield clues as to how search algorithms on networks may work, as illustrated by the recent work on Ricci-Ollivier curvature of tree space based on the *subtree prune and regraft (SPR)* operation (Whidden et al., 2015).

6. Discussion

In this paper we have shown that it is possible to transform any network into any other network through using a sequence of two elementary operations, one of which is a generalized version of the NNI operation on trees. This allows us to consider spaces of networks which generalize spaces of trees, and which could serve as search spaces for network construction algorithms.

The NNI and Δ operations are well-known in the theory of cubic graphs, where the NNI operation has been called a *Whitehead move* (Cavendish, 2010) or a *slide transformation* (Tsukui, 1996), and the Δ operation is known as the *$P1^-$ rule* (Batagelj, 1981) and *Δ -reduction* (Tsukui, 1996). We believe that

links between the theories of networks and cubic graphs could be potentially of great interest since by using theoretical tools for cubic graphs and associated structures, it should be possible to obtain new approaches to understand and construct phylogenetic networks. For example, spaces of cubic graphs appear in algebraic geometry (Billera et al., 2001) and techniques from that area have proven useful in phylogenetics (e.g. in understanding phylogenetic invariants (Allman and Rhodes, 2003)). As another illustration of these possibilities, we note that properties of cubic graphs were recently used by McDiarmid et al. (2015) to help count rooted phylogenetic networks.

Although we have focused on unrooted networks (in which each vertex has degree 1 or 3), our results serve to illustrate the possibilities for developing similar results for spaces of more complicated networks. For example, it would be very interesting to see how our results could be extended to rooted networks. Some results in this direction are presented in (Radice, 2012; Yu et al., 2014). Our previous work on spaces of level-1 rooted networks (Huber et al., 2016) suggests that this will probably be somewhat more technical. However, as algorithms are now being developed for performing likelihood and Bayesian searches on rooted networks (Radice, 2012; Yu et al., 2014), it seems to be well worth while trying to understand general properties of rooted network spaces.

In light of our results, it is also of interest to consider how network optimization searches could be performed based on a combination of either fixing network size (e.g. for unrooted networks using only NNI operations) and/or level (cf. Yu et al. (2014, p. 16450) for a search strategy on rooted networks which fixes level using non-local operations). In regards to this, we could also consider other types of operations such as generalizations of the SPR and *tree bisection and reconnection (TBR)* operations on trees (Allen and Steel, 2001). Note that, in contrast to NNI operations, these generalizations will not necessarily be local, and that there are other operations already defined for cubic graphs (cf. e.g. Batagelj, 1981) which could also yield interesting network operations. Of course, there are pros and cons in deciding which operations to consider. For example, for tree space, NNI neighborhoods are smaller than

SPR or TBR neighborhoods (cf. e.g. Humphries and Wu, 2013) which, in practical terms, means that searches based on the latter operations can be more computationally expensive. However, SPR and TBR operations can potentially yield algorithms that are less likely to become trapped in local optima. Thus it would be interesting to better understand neighborhoods in network spaces, and to use their structure to develop new search algorithms.

Another technical issue in this respect is that network spaces are in general infinite, as it is always possible to increase the level of any network by adding in vertices and edges (e.g. using the Δ operation). Hence it could be useful to develop ways to determine bounds for the level of *any* network for any given data set (maybe using concepts such as the Bayesian information criterion (BIC) (Yu et al., 2014)), and then only search spaces of networks with bounded level for that data set. This would at least give a finite search space, such as the space of level-1 networks that we defined in Huber et al. (2016).

Appendix A. Proof of Theorem 2

Proof. Suppose that G is a multigraph and that $k \geq 2$. If, starting with the theta graph, G is constructed via a sequence of $(k - 2)$ L or H operations then it is straightforward to see that G is a level- k generator.

Conversely, suppose that G is a level- k generator. We proceed by using induction on the level of G . If $k = 2$, then G is the theta graph, and so the base case clearly holds. So, suppose that any level- l generator, $2 \leq l < k$, can be constructed using $(l - 2)$ H or L operations.

If G contains a parallel edge connecting two distinct vertices u and v , say, then it must contain the subgraph pictured in Figure A.13(i). But then by performing a reverse L operation on this pair of parallel edges (i.e., deleting one of the two parallel edges and suppressing the resulting degree two vertices) we obtain a bridgeless, cubic multigraph G' (as G is a level- k generator for some $k > 2$, and hence a bridgeless, cubic multigraph). Note that this operation decreases the level of G by one and so, by induction, G' can be constructed

using $k - 3$ H or L operations. By induction the theorem therefore holds in this case.

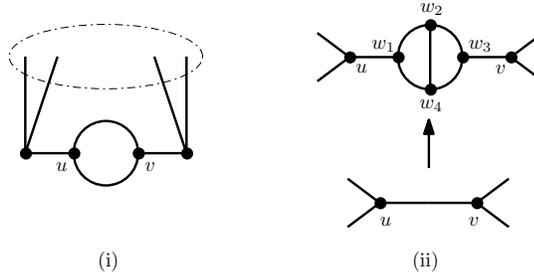


Figure A.13: Two configurations considered in the proof of Theorem 2. The vertices w_1, \dots, w_4 and their adjacent edges of the form $\{w_i, w_j\}$, $i, j \in \{1, \dots, 4\}$ are added by the operation mentioned in the text.

So, suppose G does not contain any parallel edges, i.e. G is a graph. Then by the main theorem in Batagelj (1981, part (T8)), it follows that either (a) G is a 4-clique, or (b) there is a bridgeless, cubic multigraph G' so that G can be obtained from G' by performing either an H operation or an operation of the form in Figure A.13(ii). If case (a) holds, then we can perform a single reverse H operation on G to obtain the theta graph, and the theorem is then easily seen to hold. In case (b) holds and G can be obtained from G' by performing an H operation, the theorem also follows by induction.

So, suppose that (b) holds and that G can be obtained from G' by using an operation as in Figure A.13(ii). But this operation is equivalent to first performing an L operation and then an H operation, starting with G' . Since the graph G'' obtained from G' by performing just the L operation is clearly a bridgeless cubic multigraph with level $k - 3$, and G is obtained by performing an H operation on G'' , this case again follows by induction. \square

Appendix B. Constructing networks from generators

We show how generators can be used to construct networks. We begin with a simple observation (see Figure 9).

Lemma 6. *Suppose that N is a network and G is a blob in N . Then G is either a level-0 generator, or the graph obtained from G by suppressing all degree 2 vertices (one-by-one until this is no longer possible) is a level- k generator, $k \geq 1$.*

Proof. Suppose that G is a blob in N and that G is not a level-0 generator (i.e. a vertex). Then, as G is a blob, every vertex in G must have degree 2 or 3.

If every vertex in G has degree 2, then G must be a cycle. Therefore if we suppress all but one degree 2 vertex of G we obtain a loop with a single vertex, i.e. the level-1 generator.

If G contains some vertex with degree 3, then the multigraph G' obtained by suppressing all vertices in G of degree 2 must be a bridgeless cubic multigraph (since G has no bridges), i.e. G' is a level- k generator, for some $k \geq 2$. \square

In light of Lemma 6 it follows that we can construct any network N from some collection of generators. To see this, let \mathcal{C}_N be the collection of generators (one for each blob in N) which is obtained by subjecting each blob in N to the vertex suppression process described in the statement of that lemma, if necessary. Then it is straightforward to check that to construct N from \mathcal{C}_N , we can start with some generator in \mathcal{C}_N and then continue to graft generators G in \mathcal{C}_N that have not been previously selected onto a growing multigraph M (which possibly contains loops) until we obtain N by applying one of the following operations. (a) If M is the level-0 generator u or the level-1 generator with vertex u , then to attach a new generator G to M add in the edge $\{u, v\}$, where v is the unique vertex in G if G is level-0 or level-1, or v is a new vertex which is added to G so as to subdivide some edge in G . (b) If M is neither the level-0 nor the level-1 generator, then subdivide some edge in M by adding in a new vertex u , and then, to attach a new generator G , add in the edge $\{u, v\}$

where v is as in (a). Note that this process can be regarded as an unrooted analogue of Theorem 3.1 in Gambette et al. (2009) which describes a similar way to construct directed phylogenetic networks from rooted generators.

Acknowledgments

We thank Simone Linz for interesting discussions on spaces of phylogenetic networks, and Lars Jermin and his colleagues for helpful comments on an earlier version of this manuscript. We also thank the Institute for Mathematical Sciences, Singapore, for inviting us to a workshop in July 2015, where some of the topics presented in this paper were discussed. We would also like to thank the two anonymous referees for their helpful and constructive comments on a previous version of this paper.

References

- Allen, B. and M. Steel. 2001. Subtree transfer operations and their induced metrics on evolutionary trees. *Annals of Combinatorics* 5:1–15.
- Allman, E. S. and J. A. Rhodes. 2003. Phylogenetic invariants for the general markov model of sequence mutation. *Mathematical Biosciences* 186:113–144.
- Bandelt, H.-J., P. Forster, B. C. Sykes, and M. B. Richards. 1995. Mitochondrial portraits of human populations using median networks. *Genetics* 141:743–753.
- Baptiste, E., L. van Iersel, A. Janke, S. Kelchner, S. Kelk, J. McInerney, D. Morrison, L. Nakhleh, M. Steel, L. Stougie, and J. Whitfield. 2013. Networks: expanding evolutionary thinking. *Trends in Genetics* 29(8):439–441.
- Batagelj, V. 1981. Inductive classes of cubic graphs. Pages 89–101 *in* *Colloquia Mathematica Societatis Janos Bolyai*, 37. Finite and Infinite Sets, Eger (Hungary). Elsevier Science.
- Billera, L. J., S. P. Holmes, and K. Vogtmann. 2001. Geometry of the space of phylogenetic trees. *Advances in Applied Mathematics* 27:733–767.

- Brinkmann, G., J. Goedgebeur, and B. D. McKay. 2011. Generation of cubic graphs. *Discrete Mathematics and Theoretical Computer Science* 13:69–79.
- Brinkmann, G., N. Van Cleemput, and T. Pisanski. 2013. Generation of various classes of trivalent graphs. *Theoretical Computer Science* 502:16–29.
- Bryant, D. and V. Moulton. 2004. Neighbor-net: an agglomerative method for the construction of phylogenetic networks. *Molecular Biology and Evolution* 21:255–265.
- Cardona, G., M. Llabres, F. Rossello, and G. Valiente. 2009. Metrics for phylogenetic networks ii: Nodal and triplets metrics. *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)* 6:454–469.
- Cavendish, W. 2010. Growth of the diameter of the pants graph modulo the mapping class group. Preprint, <https://web.math.princeton.edu/~wcavendi/PantsModMCG.pdf>.
- DasGupta, B., X. He, T. Jiang, M. Li, J. Tromp, and L. Zhang. 1997. On distances between phylogenetic trees. Pages 427–436 *in* Proceedings of the eighth annual ACM-SIAM symposium on Discrete algorithms Society for Industrial and Applied Mathematics.
- Felsenstein J. 2004. *Inferring Phylogenies*. Sinauer Associates Inc.
- Gambette, P., V. Berry, and C. Paul. 2009. The structure of level-k phylogenetic networks. Pages 289–300 *in* *Combinatorial Pattern Matching*. Springer.
- Gambette, P., V. Berry, and C. Paul. 2012. Quartets and unrooted phylogenetic networks. *Journal of Bioinformatics and Computational Biology* 10.
- Gusfield, D. 2014. *ReCombinatorics: The Algorithmics of Ancestral Recombination Graphs and Explicit Phylogenetic Networks*. MIT Press.
- Huber, K., S. Linz, V. Moulton, and T. Wu. 2015. Space of phylogenetic networks from generalized nearest-neighbor interchange operations. *Journal of Mathematical Biology*, 72: 699–725.

- Humphries, P. and T. Wu. 2013. The neighborhood of tree operations. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 10: 721–728.
- Huson, D. H. and D. Bryant. 2006. Application of phylogenetic networks in evolutionary studies. *Molecular Biology and Evolution* 23:254–267.
- Huson, D. H., T. Klöpper, P. J. Lockhart, and M. A. Steel. 2005. Reconstruction of reticulate networks from gene trees. Pages 233–249 *in* *Research in Computational Molecular Biology*. Springer.
- Huson, D. H., R. Rupp, and C. Scornavacca. 2010. *Phylogenetic Networks: Concepts, Algorithms and Applications*. Cambridge University Press.
- Lake, J. A. 2008. Reconstructing evolutionary graphs: 3D parsimony. *Molecular Biology and Evolution* 25:1677–1682.
- Lake, J. A. and J. S. Sinsheimer. 2013. The deep roots of the rings of life. *Genome Biology and Evolution* 5:2440–2448.
- Makarenkov, V. 2001. T-REX: reconstructing and visualizing phylogenetic trees and reticulation networks. *Bioinformatics* 17:664–668.
- McDiarmid, C., C. Semple, and D. Welsh. 2015. Counting phylogenetic networks. *Annals of Combinatorics* 19:205–224.
- McInerney, J. O. and M. Wilkinson. 2005. New methods ring changes for the tree of life. *Trends in Ecology & Evolution* 20:105–107.
- Nakhleh, L. 2011. Evolutionary phylogenetic networks: models and issues. Pages 125–158 *in* *Problem Solving Handbook in Computational Biology and Bioinformatics*. Springer.
- Radice, R. 2012. A Bayesian approach to modelling reticulation events with application to the ribosomal protein gene *rps11* of flowering plants. *Australian and New Zealand Journal of Statistics* 54:401–426.

- Rivera, M. C. and J. A. Lake. 2004. The ring of life provides evidence for a genome fusion origin of eukaryotes. *Nature* 431:152–155.
- Robinson, D. F. 1971. Comparison of labeled trees with valency three. *Journal of Combinatorial Theory, Series B* 11:105–119.
- Tsukui, Y. 1996. Transformations of cubic graphs. *Journal of the Franklin Institute* 333:565–575.
- Whelan, S. and D. Money. 2010. The prevalence of multifurcations in tree-space and their implications for tree-search. *Molecular biology and evolution* 27:2674–2677.
- Whidden, C., I. Matsen, and A. Frederick. 2015. Ricci-Ollivier curvature of the rooted phylogenetic subtree-prune-regraft graph. arXiv preprint arXiv:1504.00304 .
- Yu, Y., J. Dong, K. J. Liu, and L. Nakhleh. 2014. Maximum likelihood inference of reticulate evolutionary histories. *Proceedings of the National Academy of Sciences* 111:16448–16453.