

On the Elusiveness of Clusters

Steven M. Kelk, Celine Scornavacca, and Leo van Iersel

Abstract—Rooted phylogenetic networks are often used to represent conflicting phylogenetic signals. Given a set of clusters, a network is said to represent these clusters in the *softwired* sense if, for each cluster in the input set, at least one tree embedded in the network contains that cluster. Motivated by parsimony we might wish to construct such a network using as few reticulations as possible, or minimizing the *level* of the network, i.e., the maximum number of reticulations used in any “tangled” region of the network. Although these are NP-hard problems, here we prove that, for every fixed $k \geq 0$, it is polynomial-time solvable to construct a phylogenetic network with level equal to k representing a cluster set, or to determine that no such network exists. However, this algorithm does not lend itself to a practical implementation. We also prove that the comparatively efficient CASS algorithm correctly solves this problem (and also minimizes the reticulation number) when input clusters are obtained from two not necessarily binary gene trees on the same set of taxa but does not always minimize level for general cluster sets. Finally, we describe a new algorithm which generates in polynomial-time all binary phylogenetic networks with exactly r reticulations representing a set of input clusters (for every fixed $r \geq 0$).

Index Terms—Rooted phylogenetic networks, clusters, reticulate evolution, parsimony, computational complexity, polynomial-time algorithms.



1 INTRODUCTION

THE traditional abstraction for modeling evolution is the phylogenetic tree. The underlying principle of such a tree is that the observed diversity in a set of species (or, more abstractly, a set of *taxa*) can be explained by branching events that cause lineages to split into two or more sublineages [1], [2], [3]. However, there is increasing attention for the situation when observed data cannot satisfactorily be modeled by a tree. The field of phylogenetic networks has arisen with this challenge in mind. Phylogenetic networks generalize phylogenetic trees, but within this very general characterization there are many different definitions and models [4]. In this paper, we are concerned with *rooted* phylogenetic networks. Such networks assume that the observed data evolve from a unique starting point (the root) and that evolution is directed away from this root. The main way these networks differ from rooted phylogenetic trees is the presence of *reticulation* nodes: nodes with indegree 2 or higher. For the remainder of this paper, we will use the term phylogenetic network, or just network, to refer to rooted phylogenetic networks. We refer the reader to [4], [5], [6], [7], [8] for detailed background information.

Constructing a phylogenetic network that “explains” the observed data is a trivial problem if no optimality criteria are imposed upon the constructed network. One simple optimality criterion that has attracted a great deal of

attention in the literature, *reticulation minimization*, is to compute a phylogenetic network that explains the observed data but using as few reticulation events as possible. This is an algorithmically hard problem, irrespective of the exact construction technique that is being applied [8]. A related optimality criterion, *level minimization* [9], [10], [11], [12], [13] is motivated by the observation that a phylogenetic network can be regarded as some kind of tree backbone decorated with *tangles* of reticulate activity [4]. Here, the challenge is to construct a phylogenetic network that explains the observed data but such that the maximum number of reticulation events inside any biconnected component, the *level*, is as low as possible. Reticulation minimization is thus a global optimality criterion, and level minimization is in some sense a local optimality criterion; see Fig. 1. Both criteria will have an important role in this paper.

The question remains, when does a phylogenetic network explain the observed data? This depends very much on the exact construction technique being applied. The classical problem is motivated by the biological observation that, although the evolution of a set of organisms might best be explained by a phylogenetic network, the individual genes of the organisms will generally undergo treelike evolution [5]. In such a case, one can think of the gene trees as being *displayed by* (i.e., topologically embedded within) the species network. The reticulation nodes then have an explicit biological interpretation as (for example) hybridization, recombination or horizontal gene transfer events. Hence, the following problem: given a set of rooted phylogenetic trees, all on the same set of taxa, compute a phylogenetic network with a minimum number of reticulations that displays all the input trees. The problem is already NP-hard (and APX-hard) for the case of two input trees [14]. However, extensive research by different authors has shown that, by exploiting the fixed parameter tractability of the problem [15], [16], the two-tree problem can be solved to satisfaction for many instances [17], [18]. The case

• S.M. Kelk is with the Department of Knowledge Engineering (DKE), Maastricht University, PO Box 616, Maastricht 6200 MD, The Netherlands. E-mail: steven.kelk@maastrichtuniversity.nl.

• C. Scornavacca is with the Center for Bioinformatics (ZBIT), Tübingen University, Sand 14, Tübingen 72076, Germany. E-mail: scornava@informatik.uni-tuebingen.de.

• L. van Iersel is with the Department of Mathematics and Statistics, University of Canterbury, Private Bag 4800, Christchurch, New Zealand. E-mail: l.j.v.iersel@gmail.com.

Manuscript received 10 Mar. 2011; revised 22 July 2011; accepted 22 Aug. 2011; published online 27 Sept. 2011.

For information on obtaining reprints of this article, please send e-mail to: tcbb@computer.org, and reference IEEECS Log Number TCBB-2011-03-0059. Digital Object Identifier no. 10.1109/TCBB.2011.128.




Fig. 1. Example of a phylogenetic network with five reticulations. The encircled subgraphs form its biconnected components, also known as its “tangles.” This binary network has level equal to 2 since each biconnected component contains at most two reticulations.

of more than two input trees, or input trees that are not all binary (i.e., some nodes have outdegree three or higher), has been considerably less well studied [19], [20].

A parallel, and related, line of research concerns “piecewise” assembly of phylogenetic networks. Whereas the tree problem described above concerns the combination of a small number of large hypotheses (e.g., gene trees) into a phylogenetic network, an alternative strategy is to combine a large number of small hypotheses into a phylogenetic network. Examples of such small hypotheses include *rooted triplets* (phylogenetic trees defined on size-3 subsets of the taxa) [21], [12], (binary) characters (e.g., whether or not the taxon is vertebrate) [22], [23], [24], [25], and *clusters* (clades) [26], [7], [27]. Proponents of such piecewise assembly techniques argue that in this way it is easier (than with trees) to discard parts of the input that are not well supported. In this paper, we focus specifically on clusters, although the classical tree problem and all the other piecewise construction techniques do play a secondary role. This secondary role is linked to the fact, as observed in [8], that under certain circumstances all these different models behave in a unified way. We shall return to this point later.

Let us then say more about the cluster model. A cluster C is a subset of the taxa and we say that a phylogenetic network *represents* the cluster in the software sense if *some* tree embedded in the network contains a clade equal to that cluster [7]. In other words: *some* tree T embedded in the network has an edge such that C is exactly the set of all taxa reachable from the head of that edge by directed paths. The general problem is, given a set of clusters, to construct an optimal phylogenetic network that represents all the input clusters. The set of input clusters can be constructed in an ad hoc fashion, but often the set of clusters is generated by extracting the set of clusters induced by a set of rooted phylogenetic trees and then possibly excluding weakly supported clusters. This is the technique applied in the program DENDROSCOPE [28]. A disadvantage of this technique is that some of the topology of the original trees can be lost [8], but on the plus side it permits a focus on only well-supported clades, which is a major concern of practicing phylogeneticists.

In [7] it was shown, given a set of clusters, how to construct a *galled network* with a small number of reticulations that




Fig. 2. (a) The output of the galled network algorithm [7] for $C = \{\{a, b, f, g, i\}, \{a, b, c, f, g, i\}, \{a, b, f, i\}, \{b, c, f, i\}, \{c, d, e, h\}, \{d, e, h\}, \{b, c, f, h, i\}, \{b, c, d, f, h, i\}, \{b, c, i\}, \{a, g\}, \{b, i\}, \{c, i\}, \{d, h\}\}$ and (b) a (simple) network with two fewer reticulations that also represents this set of clusters.

represents the clusters. However, given that galled networks are a restricted subclass of phylogenetic networks, it was unclear how far that algorithm actually minimizes the number of reticulations (or the level) when ranging over the entire space of phylogenetic networks, see, for example, Fig. 2. This is the context in which the CASS algorithm was developed [27]. The CASS algorithm, in some sense a natural follow-up to the algorithm of [7], was formally shown to produce solutions of minimum *level* whenever the minimum level is at most two. However, the optimality of the CASS algorithm for “higher-level” inputs, and the performance of the algorithm in terms of minimizing number of reticulations, remained unclear. On the practical side the good news was that CASS produced solutions with fewer reticulations and lower level than the algorithm from [7]. Intriguingly it was also observed in [27] that, for several sets of input clusters induced by two binary trees, the networks produced by CASS had an identical number of reticulations to networks generated by algorithms that aim to display the trees themselves. This observation was the inspiration behind [8] in which it was proven that, in the case of two binary trees, the choice of construction technique (tree, triplets, characters, clusters) does not affect the number of reticulations required. However, this unification was shown to break down for data obtained from three or more binary trees.

After [8] several important questions about clusters remained open. Does CASS always minimize level? If not, can we find a different algorithm that efficiently minimizes level? Under which circumstances does CASS also minimize the number of reticulations? In how far do the unification results of [8] hold for two nonbinary trees?

The results in this paper settle many of these open questions. First, we show that in the case of clusters obtained from two not necessarily binary trees, a divide and conquer algorithm using CASS as subroutine, called here $CASS^{DC}$ and implemented in DENDROSCOPE [28], does minimize both the number of reticulations, and the level. Spin-off results from this include nonbinary versions of several unification results from [8], culminating in the observation that, in the case of clusters obtained from two trees, $CASS^{DC}$ also computes the minimum number of reticulations required to display the trees themselves (in the sense of [29]), rather than just the clusters from the trees. We also obtain deeper insights into why the two-tree case is so special and not representative for the problem on three or

more trees. In particular, the two-tree case seems to be best understood as the only point at which a very natural lower bound is guaranteed to be tight.

Second, we show that CASS does not, unfortunately, always minimize level when the input data require solutions of level 3 or higher. We give an explicit counterexample and explain what goes wrong with the CASS algorithm in this case.

To offset this negative result, we describe a polynomial-time algorithm that shows, for every fixed natural number k , how to determine whether a set of clusters can be represented by a network with level k . This algorithm, which is very different to CASS, is purely theoretical but does give important insights into the underlying structure of the cluster model. Also on the positive side we show that, for sets of clusters induced by arbitrarily large sets of *binary* trees, a simple polynomial-time algorithm can construct *all* binary phylogenetic networks with r reticulations that represent the clusters, for every fixed $r \geq 0$. To demonstrate an important design principle first observed in [8] we give a practical implementation of this algorithm, CLUSTISTIC, which elegantly “bootstraps” an existing software package for merging rooted triplets into a phylogenetic network.

To summarize, the results in this paper help advance our understanding of the cluster model considerably. Nevertheless, many questions remain, and in the final section of this paper we discuss a number of them. Perhaps the biggest question, which is the motivation for the title of this paper, concerns the fact that the cluster model so far has not enjoyed the same kind of steady algorithmic improvements witnessed in the tree literature. Why does it seem harder to work with the clusters inside the trees than the trees themselves? To what, exactly, can the elusiveness of clusters be attributed?

2 PRELIMINARIES

Consider a set \mathcal{X} of taxa. A *rooted phylogenetic network* (on \mathcal{X}), henceforth *network*, is a directed acyclic graph with a single node with indegree zero (the *root*), no nodes with both indegree and outdegree equal to 1, and leaves bijectively labeled by \mathcal{X} . In this paper, we identify the leaves with \mathcal{X} . The indegree of a node v is denoted $\delta^-(v)$ and v is called a *reticulation* if $\delta^-(v) \geq 2$. An edge (u, v) is called a *reticulation edge* if its target node v is a reticulation and is called a *tree edge* otherwise. When counting reticulations in a network, we count reticulations with more than two incoming edges more than once because, biologically, these reticulations represent several reticulate evolutionary events. Therefore, we formally define the *reticulation number* of a network $N = (V, E)$ as

$$r(N) = \sum_{v \in V: \delta^-(v) > 0} (\delta^-(v) - 1) = |E| - |V| + 1.$$

A *rooted phylogenetic tree* on \mathcal{X} , henceforth *tree*, is simply a network that has reticulation number zero. We say that a network N on \mathcal{X} *displays* a tree T if T can be obtained from N by performing a series of node and edge deletions and eventually by suppressing nodes with both indegree and outdegree equal to 1. We assume without loss of generality

that each reticulation has outdegree at least one. Consequently, each leaf has indegree one. We say that a network is *binary* if every reticulation node has indegree 2 and outdegree 1 and every tree node that is not a leaf has outdegree 2.

Proper subsets of \mathcal{X} are called *clusters*, and a cluster C is a *singleton* if $|C| = 1$. We say that an edge (u, v) of a tree *represents* a cluster $C \subset \mathcal{X}$ if C is the set of leaf descendants of v . A tree T represents a cluster C if it contains an edge that represents C . It is well known that the set of clusters represented by a tree is a laminar family, often called a *hierarchy* in the phylogenetics literature, and uniquely defines that tree. We say that a network N represents a cluster $C \subset \mathcal{X}$ “in the hardwired sense” if there exists a tree edge (u, v) of N such that C is the set of leaf descendants of v . Alternatively, we say that N represents C “in the softwired sense” if N displays some tree T on \mathcal{X} such that T represents C . In this paper, we only consider the softwired notion of cluster representation and henceforth assume this implicitly. A network represents a set of clusters \mathcal{C} if it represents every cluster in \mathcal{C} (and possibly more). The set of softwired clusters of a network can be obtained as follows: For a network N , we say that a *switching* of N is obtained by, for each reticulation node, deleting all but one of its incoming edges. Given a network N and a switching T_N of N , we say that an edge (u, v) of N represents a cluster C w.r.t. T_N if (u, v) is an edge of T_N and C is the set of leaf descendants of v in T_N . The set of softwired clusters of N is the set of clusters represented by all edges of N w.r.t. T_N , where T_N ranges over all possible switchings [4]. It is also natural to define that an edge (u, v) of N represents a cluster C if there exists some switching T_N of N such that (u, v) represents C w.r.t. T_N . Note that, in general, an edge of N might represent multiple clusters, and a cluster might be represented by multiple edges of N .

Given a set of clusters \mathcal{C} on \mathcal{X} , throughout the paper we assume that, for any taxon $x \in \mathcal{X}$, \mathcal{C} contains at least one cluster C containing x . For a set \mathcal{C} of clusters on \mathcal{X} we define $r(\mathcal{C})$ as $\min\{r(N) \mid N \text{ represents } \mathcal{C}\}$, we sometimes refer to this as the *reticulation number* of \mathcal{C} . The related concept of *level* requires some more background. A directed acyclic graph is *connected* (also called “weakly connected”) if there is an undirected path (ignoring edge orientations) between each pair of nodes. A node (edge) of a directed graph is called a *cut-node* (*cut-edge*) if its removal disconnects the graph. A directed graph is *biconnected* if it contains no cut-nodes. A biconnected subgraph B of a directed graph G is said to be a *biconnected component* if there is no biconnected subgraph $B' \neq B$ of G that contains B . A phylogenetic network is said to be a *level- $\leq k$ network* if each biconnected component has reticulation number less than or equal to k .¹ A level- $\leq k$ network is called a *simple level- $\leq k$ network* if the removal of a cut-node or a cut-edge creates two or more connected components of which at most one is nontrivial (i.e., contains at least one edge). A (simple) level- $\leq k$ network N is called a (simple) *level- k network* if the maximum reticulation number among the biconnected

1. Note that to determine the reticulation number of a biconnected component, the indegree of each node is computed using only edges belonging to this biconnected component.

components of N is precisely k . For example, the network in Fig. 1 is a level-2 network while the one in Fig. 2a is a simple level-4 network. Note that a tree is a level-0 network. For a set \mathcal{C} of clusters on \mathcal{X} we define $\ell(\mathcal{C})$, the *level* of \mathcal{C} , as the smallest $k \geq 0$ such that there exists a level- k network that represents \mathcal{C} . It is immediate that for every cluster set \mathcal{C} $r(\mathcal{C}) \geq \ell(\mathcal{C})$, because a level- k network always contains at least one biconnected component containing k reticulations.

We say that two clusters $C_1, C_2 \subset \mathcal{X}$ are *compatible* if either $C_1 \cap C_2 = \emptyset$ or $C_1 \subseteq C_2$ or $C_2 \subseteq C_1$, and *incompatible* otherwise. Consider a set of clusters \mathcal{C} . The *incompatibility graph* $IG(\mathcal{C})$ of \mathcal{C} is the undirected graph (V, E) that has node set $V = \mathcal{C}$ and edge set $E = \{\{C_1, C_2\} | C_1 \text{ and } C_2 \text{ are incompatible clusters in } \mathcal{C}\}$. We say that a set of taxa $\mathcal{X}' \subseteq \mathcal{X}$ is *compatible* with \mathcal{C} if every cluster $C \in \mathcal{C}$ is compatible with \mathcal{X}' , and *incompatible* otherwise.

We say that a set of clusters \mathcal{C} on \mathcal{X} is *separating* if it is incompatible with all sets of taxa \mathcal{X}' such that $\mathcal{X}' \subset \mathcal{X}$ and $|\mathcal{X}'| \geq 2$. We say that a set of clusters \mathcal{C} on \mathcal{X} is *tangled* [4] if:

1. $IG(\mathcal{C})$ is connected and has more than one node;
2. every size-2 subset of \mathcal{X} is incompatible with \mathcal{C} .

Remember that here we assume that any taxon of \mathcal{X} is contained in at least one cluster $C \in \mathcal{C}$. Then, it can easily be verified that, given a tangled set of clusters \mathcal{C} on \mathcal{X} , \mathcal{C} is separating.

The incompatibility graph and the concept of tangled cluster sets are important because they highlight an important difference between (the computation of) $r(\mathcal{C})$ and $\ell(\mathcal{C})$. In [27] the authors show that, if $\ell(\mathcal{C}) = k$, then a level- k network that represents \mathcal{C} can be constructed by combining in polynomial time simple level- $\leq k$ networks constructed independently for each connected component of $IG(\mathcal{C})$. The actual procedure is slightly more involved but it shows in any case that a polynomial-time algorithm for constructing simple level- $\leq k$ networks can easily be extended to a polynomial-time algorithm for constructing level- $\leq k$ networks. We will make use of this fact in Section 3.

Unfortunately, as has been observed by several authors, the same procedure does *not* necessarily lead to networks that have reticulation number $r(\mathcal{C})$ (see, for example, [7, Fig. 3], which is based on [22]). In other words, computation of $r(\mathcal{C})$ requires something more complicated than independently optimizing each connected component of $IG(\mathcal{C})$. An important special case, however, is when \mathcal{C} is separating; in this case any network N that represents \mathcal{C} is simple (or can be trivially modified to become simple) and $r(\mathcal{C}) = \ell(\mathcal{C})$. We will formalize this in due course.

To conclude the preliminaries we note that, throughout the paper, we often write that an algorithm is “polynomial time” without formally specifying what the input size is. Unless otherwise specified the input is a set of clusters \mathcal{C} on taxa set \mathcal{X} . It is sufficient to take $|\mathcal{C}| + |\mathcal{X}|$ as a lower bound on the size of the input. In some cases (such as Lemma 3 in Section 3) $|\mathcal{C}|$ is at most a constant factor larger than $|\mathcal{X}|$ and then it is sufficient to prove a running time polynomial in $|\mathcal{X}|$. In other cases a running time of the form $O(|\mathcal{C}|^a |\mathcal{X}|^b)$ is obtained, for constants a and b , and this is clearly polynomial in $|\mathcal{C}| + |\mathcal{X}|$ because $(|\mathcal{C}| + |\mathcal{X}|)^2 \geq |\mathcal{C}| |\mathcal{X}|$.

2.1 Structure of the Paper

To facilitate the mathematical exposition we build the results of this paper up in a specific order, which differs from the order presented in the introduction. We begin with Section 3, *A theoretical polynomial-time algorithm for constructing level- k networks*, where we prove that, for every fixed $k \geq 0$, the problem of determining whether a level- k network that represents \mathcal{C} exists (and if so to construct such a network) is solvable in polynomial time. This section is, compared to the rest of the paper, comparatively self-contained. In Section 4, *From theory to practice: the importance of ST-sets*, we describe several fundamental properties of the ST-set, a special structure that plays a central role throughout the rest of the paper. In Section 5, *Clusters obtained from sets of binary trees on \mathcal{X}* , we show how, given a set \mathcal{T} of binary trees on \mathcal{X} , and for each fixed $r \geq 0$, it is possible to construct in polynomial time all binary phylogenetic networks with reticulation number r that represent all the clusters in the input trees. We also describe CLUSTISTIC, which is our implementation of this algorithm built on top of already-existing software. In Section 6, *Witnesses and a natural lower bound*, we further develop the theory surrounding the ST-set, explicitly relating it to the computation of reticulation number. This is used extensively in Section 7, *The optimality and nonoptimality of CASS*, where we give both positive and negative results for the CASS algorithm, and in the process develop a number of powerful generalizations of unification results from [8].

3 A THEORETICAL POLYNOMIAL-TIME ALGORITHM FOR CONSTRUCTING LEVEL- k NETWORKS

In this section, we prove that, for every fixed $k \geq 0$, the problem of determining whether a level- k network exists that represents \mathcal{C} , and if so to construct such a network, can be solved in polynomial time.

We first require some auxiliary lemmas and definitions. The claims are intuitive but, because of subtleties that can arise when dealing with heavily “redundant” networks, the proofs are rather technical and have been deferred to the Appendix, which can be found on the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TCBB.2011.128>.

For a node v let $\mathcal{X}(v) \subseteq \mathcal{X}$ be the set of all taxa reachable from v by directed paths. For an edge $e = (u, v)$ we define $\mathcal{X}(e)$ to be equal to $\mathcal{X}(v)$.

Observation 1. *Let \mathcal{C} be a separating set of clusters on \mathcal{X} . Let N be any network that represents \mathcal{C} . Then, each node of N has at most one leaf child and for each cut-edge (u, v) in N , $|\mathcal{X}(v)| = 1$ or $\mathcal{X}(v) = \mathcal{X}$.*

Proof. Deferred to the Appendix, which can be found on the Computer Society Digital Library. \square

Lemma 1. *Let \mathcal{C} be a separating set of clusters on \mathcal{X} . Let N be any network that represents \mathcal{C} . Then there exists a simple network N^* with at most one leaf-child per node such that $\ell(N^*) \leq \ell(N)$.*

Proof. Deferred to the Appendix, which can be found on the Computer Society Digital Library. \square




Fig. 3. The single level-1 generator and the four level-2 generators. Here, the sides have been labeled with capital letters.

Observation 1 and Lemma 1 formalize the idea that any network that represents a separating set of clusters is simple or can easily be made simple by deleting certain “redundant” parts of it. The following lemma shows that, in terms of minimizing reticulation number or level, we can assume without loss of generality that networks are binary.

Lemma 2. *Let N be a phylogenetic network on \mathcal{X} . Then, we can transform N into a binary phylogenetic network N' such that N' has the same reticulation number and level as N and all clusters represented by N are also represented by N' .*

Proof. Deferred to the Appendix, which can be found on the Computer Society Digital Library. \square

Note that, given a simple binary network N , each node of N has at most one leaf child. Armed with these technical results, we are ready to prove the main result of this section.

Lemma 3. *Let \mathcal{C} be separating set of clusters on \mathcal{X} . Then, for every fixed $k \geq 0$, it is possible to determine in polynomial time whether a level- k network exists that represents \mathcal{C} , and if so to construct such a network.*

Proof. From Lemmas 1 and 2 it is sufficient to focus on simple binary networks. We assume then that, for fixed k , there exists a binary simple level- k network N that represents \mathcal{C} . Let $|\mathcal{X}| = n$. Then, \mathcal{C} will contain at most $2^{k+1}(n - 1)$ clusters, because there are at most 2^k trees displayed by a simple level- k network, and each tree represents at most $2(n - 1)$ clusters. Thus, for fixed k , the size of the input is polynomial in n . It follows from these observations that whether a set of clusters is represented by a given simple level- k network can be checked in polynomial time.

It is known that, if the leaves of N are removed and all nodes with both indegree and outdegree equal to 1 are suppressed, the resulting structure will be a level- k generator, defined in [11]. See also Fig. 3. For fixed k , there are only a constant number of level- k generators [30, Proposition 2.5]. Recall that the *sides* of a level- k generator are defined as the union of its edges and its nodes of indegree-2 and outdegree-0. For fixed k the maximum number of sides ranging over all level- k generators, is a constant.

For a cluster set \mathcal{C} on \mathcal{X} , we write $x \rightarrow y$ if and only if every nonsingleton cluster in \mathcal{C} that contains x , also contains y . For example, for the cluster set of Fig. 4, we have $e \rightarrow d$ and $f \rightarrow e$.

In the remainder of the proof, we illustrate a simple algorithm for determining whether a binary simple level- k network that represents \mathcal{C} exists by attempting to reconstruct such a network. Let g be the generator underlying N . We only require polynomially many tries to compute g , because there are only a constant number of generators. So assume we know g . For each side of g , we guess whether there are 0, 1, 2 or more than two leaves on that side. For each side containing exactly one leaf, we guess what that is. For each side s of g containing two or more leaves, we guess the leaf s^+ that is nearest to the root on that side, and the leaf s^- that is furthest from the root on that side. For an example see Fig. 4. Note that, since for each side we have only four options (0, 1, 2, or more than 2 leaves), and in the latter case only s^+ and s^- have to be chosen, it follows that we have a polynomial number of guesses to try.

We will now show how to add the remaining leaves. Note that we may fail to insert all leaves in the network. This means that we made the wrong guess and that another set of guesses has to be checked. We say that a side s is *lowest* if it does not yet have all its leaves, and there is no other such side s' reachable from s . By reachable we mean that in the underlying generator g , there is a directed path from the head of side s to the tail of side s' . Since N is a directed acyclic graph, until all leaves in \mathcal{X} have been added, there will always be a lowest side. For example, the side F in Fig. 4d is lowest. The idea is to add leaves to the lowest side s , until all its leaves have been added. We then continue with remaining lowest sides until we have reconstructed N .




Fig. 4. An example of the execution of the algorithm outlined in Lemma 3 for the separating set of clusters $\mathcal{C} = \{\{a, b\}, \{a, c\}, \{c, d\}, \{d, e\}, \{a, b, c\}, \{c, d, e\}, \{d, e, f\}, \{c, d, e, f\}, \{b, d, e, f\}, \{b, c, d, e, f\}, \{a\}, \{b\}, \{c\}, \{d\}, \{e\}, \{f\}\}$ on $\mathcal{X} = \{a, b, c, d, e, f\}$. The value of k is fixed to 2. (a) The chosen generator g is the generator 2c in Fig. 3. (b) We guess that the sides E , H , and G contain one leaf while the side F contains more than two leaves and all other sides contain zero leaves. (c) We guess the single leaves on sides E , H , and G . (d) We guess the leaves s^+ and s^- on side F . (e) We deduce the last leaf in F and we obtain a simple level-2 network on \mathcal{X} . By a stroke of luck, our first guess represents \mathcal{C} .

Given a lowest side s , with s^+ and s^- fixed, it is possible to tell in polynomial time what the correct remaining leaves for s are, as follows: Observe that a leaf x that is on side s in N and which has not yet been added has the property $s^+ \rightarrow x \rightarrow s^-$. Furthermore, there is at least one cluster $C \in \mathcal{C}$ such that $\{x, s^+, s^-\} \cap C = \{x, s^-\}$. There exists at least one such cluster because otherwise $\{x, s^+\}$ would be compatible with \mathcal{C} , a contradiction since \mathcal{C} is separating. We call such a cluster a *split cluster for side s* . Now, observe that for every split cluster C for side s , and for every side $t \neq s$ that contains two or more leaves in N , either $\{t^+, t^-\} \cap C = \{t^+, t^-\}$ or $\{t^+, t^-\} \cap C = \emptyset$. This follows because the only edges in N that represent C lie on side s . If this is not the case, our set of guesses was incorrect and a new one has to be checked.

Now, consider any leaf y that has not yet been added to the network. Assume that this leaf belongs to side t for some t . We want to have a simple test to avoid wrongly placing it on side s , with $s \neq t$. Side t will contain three or more leaves in N , so we can assume that t^+ and t^- exist. If $s^+ \rightarrow y \rightarrow s^-$ does not hold then it is immediately clear that y cannot be put on side s . So assume (conversely) that this condition does hold, and for the same reason assume there is a split cluster C for side s that contains y . In other words, there is a cluster C such that $\{y, s^+, s^-\} \cap C = \{y, s^-\}$. Since $t^+ \rightarrow y \rightarrow t^-$ holds, it follows that C also contains t^+ and t^- , because any cluster that contains y also contains t^- , and we know that C contains either both of t^+ and t^- , or neither of them. However, there is no edge in N that can represent C : the only edges that represent C lie on side s , but the fact that s is the lowest side means that no cluster beginning on side s can contain any leaves on side t . To summarize, we have a simple test for determining whether a leaf should be placed on side s . Once we have determined the set of leaves that should be placed on side s , it is easy to determine the correct order of those leaves by inspecting \rightarrow relationships. Indeed, if q and p are two leaves that belong to side s , and q is nearer to s^- in the network we aim to reconstruct, then obviously $p \rightarrow q$. Since \mathcal{C} is separating there will be (by separation) some cluster that contains q but not p , so $q \not\rightarrow p$. If all leaves can be added in such a way, we obtain a simple level- k network on \mathcal{X} and we can check in polynomial time if it represents \mathcal{C} . If it is not the case or we fail to insert at least one leaf in the network, another set of guesses can be checked until a simple level- k network representing \mathcal{C} (if any exist) is found. This concludes the proof. \square

The following corollary follows automatically from the generality of the proof of Lemma 3.

Corollary 1. *Let \mathcal{C} be a separating set of clusters on \mathcal{X} . Then, for every fixed $k \geq 0$, it is possible to construct in polynomial time all binary simple level- k networks that represent \mathcal{C} .*

Using Lemma 3, we can prove the following result:

Theorem 1. *Let \mathcal{C} be a (not necessarily separating) set of clusters on \mathcal{X} . Then, for every fixed $k \geq 0$, it is possible to determine in polynomial time whether a level- k network exists that represents \mathcal{C} , and if so to construct such a network.*

Proof. Recall that all tangled cluster sets are separating. It was shown in [27] that the existence of a polynomial-time algorithm for constructing a level- $\leq k$ network from a tangled cluster set, is sufficient to give a polynomial-time algorithm for constructing level- k networks from general cluster sets. (Specifically, several tangled cluster sets are obtained by processing each nontrivial connected component of the incompatibility graph of the original cluster set [27], [4]). Hence, we can assume without loss of generality that \mathcal{C} is tangled. Lemma 3 is thus sufficient, and we are done. \square

Note that the running time of this theoretical algorithm is extremely high, and is already impractical for $k = 2$ and small values of n . This is because

1. the number of level- k generators grows rapidly, lying between 2^{k-1} and $k!2^{50k}$ [30], and for $k = 1, 2, 3, 4, 5$ is 1, 4, 65, 1993, 91454, respectively, [12], [30], [31];
2. the number of edges (and thus also sides) in a generator grows linearly in k [12];
3. the need to guess whether there are 0, 1, 2 or more than two leaves on each side;
4. the need to make up to $O(n^2)$ guesses per side (i.e., to guess s^- and s^+ for each side).

3.1 Rooted Triplets

It is interesting to note that the proof technique used in Lemma 3 leads to a simplified proof, presented in the following corollary, of a complexity result that was first proven in [21]. (The algorithm in [21] yielded a much faster running time, however). Let us first recall several definitions related to rooted triplets. A (*rooted*) *triplet* on \mathcal{X} is a binary phylogenetic tree on a size-3 subset of \mathcal{X} . We use $xy|z$ to denote the triplet with taxa x, y on one side of the root and z on the other side of the root. For triplets, the notion of “represent” can be formalized by the notion of “display” introduced above. However, for triplets “consistent with” is often used instead of “displayed by.” A triplet $xy|z$ is *consistent* with a phylogenetic network N (and N is *consistent* with $xy|z$) if $xy|z$ is displayed by N . See Fig. 5 for an example. Given a phylogenetic tree T on \mathcal{X} , we let $Tr(T)$ denote the set of all rooted triplets on \mathcal{X} that are consistent with T . For a set of phylogenetic trees \mathcal{T} , we let $Tr(\mathcal{T})$ denote the set of all rooted triplets that are consistent with some tree in \mathcal{T} , i.e., $Tr(\mathcal{T}) = \bigcup_{T \in \mathcal{T}} Tr(T)$. A set of triplets on \mathcal{X} is *dense* if, for every size-3 subset $\{x, y, z\} \subseteq \mathcal{X}$, at least one of $xy|z, xz|y, yz|x$ is in the triplet set.

Corollary 2. *Let R be a dense set of triplets on \mathcal{X} . Then, for every fixed $k \geq 0$, it is possible to determine in polynomial time whether a binary simple level- k network exists that is consistent with R , and if so to construct such a network.*

Proof. As pointed out in [21] it is possible to determine in polynomial time whether a given network is indeed consistent with a set of input triplets. Then, the proof of Lemma 3 holds here almost entirely. The only significant difference concerns the adding of leaves to the lowest side: a not yet allocated leaf x belongs on lowest side s if and only if the triplet $s^-x|s^+$ is in the input. \square

We shall return to rooted triplets again later in the paper.




Fig. 5. A phylogenetic tree T (a) and a phylogenetic network N (b,c,d); (b) illustrates in red that N displays T (deleted edges are dashed); (c) illustrates that N is consistent with (among others) the triplet $cd|f$ (deleted edges are again dashed); (d) illustrates that N represents (among others) cluster $\{c, d, e\}$ in the softwired sense (dashed reticulation edges are “switched off”).

4 FROM THEORY TO PRACTICE: THE IMPORTANCE OF ST-SETS

The algorithm described in Section 3 is polynomial time but only of theoretical interest because its running time is too high to be useful in practice. In the rest of this paper, we will focus on practical polynomial-time algorithms. In all these algorithms the ST-set, which can informally be thought of as treelike subsets of \mathcal{X} , has a central role. We begin by formally defining ST-sets and describing their basic properties. We will expand upon these basic properties in subsequent sections of the paper.

4.1 Definition and Basic Properties of ST-Sets

Given a set $S \subseteq \mathcal{X}$ of taxa, we use $\mathcal{C} \setminus S$ to denote the result of removing all elements of S from each cluster in \mathcal{C} and we use $\mathcal{C}|S$ to denote $\mathcal{C} \setminus (\mathcal{X} \setminus S)$ (the restriction of \mathcal{C} to S). We say that a set $S \subseteq \mathcal{X}$ is an ST-set with respect to \mathcal{C} , if S is compatible with \mathcal{C} and any two clusters $C_1, C_2 \in \mathcal{C}|S$ are compatible. Note that, unlike in [27], we allow the possibility that $S = \emptyset$ or $S = \mathcal{X}$. (We say that an ST-set S is *trivial* if $S = \emptyset$ or $S = \mathcal{X}$). An ST-set S is *maximal* if there is no ST-set T with $S \subset T$, see Fig. 6 for an example.

Informally, the maximal ST-sets of \mathcal{C} are the result of repeatedly collapsing pairs of taxa compatible with \mathcal{C} for as long as possible; we can think of them as “islands of laminarity” within the cluster set.

ST-sets first explicitly appeared in [27] but, as we shall see in due course, they implicitly arose earlier in the *recombination network* literature. An important feature of ST-sets is that there can in general be very many of them. For example, suppose \mathcal{C} contains only $|\mathcal{X}| = n$ singleton clusters; then \mathcal{C} has 2^n ST-sets. However, as the following technical results show, \mathcal{C} will have at most n maximal ST-sets, and they will partition \mathcal{X} , i.e., they are mutually disjoint and entirely cover \mathcal{X} . Several of the proofs have been deferred to

the Appendix, which can be found on the Computer Society Digital Library.

Lemma 4. *Let \mathcal{C} be a set of clusters on \mathcal{X} and let $S_1 \neq S_2$ be two ST-sets of \mathcal{C} . If $S_1 \cap S_2 \neq \emptyset$ then $S_1 \cup S_2$ is an ST-set.*

Proof. Deferred to the Appendix, which can be found on the Computer Society Digital Library. □

Corollary 3. *Let \mathcal{C} be a set of clusters on \mathcal{X} and let $S_1 \neq S_2$ be two maximal ST-sets of \mathcal{C} . Then, $S_1 \cap S_2 = \emptyset$.*

Corollary 4. *Let \mathcal{C} be a set of clusters on \mathcal{X} . Then, there are at most n maximal ST-sets with respect to \mathcal{C} , they are uniquely defined and they partition \mathcal{X} .*

Proof. Deferred to the Appendix, which can be found on the Computer Society Digital Library. □

Lemma 5. *The maximal ST-sets of a set of clusters \mathcal{C} on \mathcal{X} can be computed in polynomial time.*

Proof. Deferred to the Appendix, which can be found on the Computer Society Digital Library. □

Corollary 4 and Lemma 5 are perhaps not so surprising, but we have nevertheless proven them rigorously to highlight the fact that *computing* maximal ST-sets is not a complexity bottleneck. In later sections, we shall see that there is a link between NP-hardness and maximal ST-sets, but that the hardness lies in selecting certain maximal ST-sets with special properties, not in the computation of the maximal ST-sets *per se*.

Let \mathcal{T} be a set of trees, where each $T \in \mathcal{T}$ is a tree on \mathcal{X} . For a tree T , we write $Cl(T)$ to denote the set of clusters induced by edges of T , i.e., $C \in Cl(T)$ if and only if some edge of T represents C . We let $Cl(\mathcal{T}) = \cup_{T \in \mathcal{T}} Cl(T)$. Whenever, we (reasonably) assume that all singleton clusters are present in the input,² it is easy to see that every cluster set \mathcal{C} on \mathcal{X} can be written as $Cl(\mathcal{T})$ for some \mathcal{T} as follows: We take any proper coloring of $IG(\mathcal{C})$ (i.e., map the nodes of $IG(\mathcal{C})$ to colors such that no two adjacent nodes have the same color) and use the resulting colors to partition \mathcal{C} . Clusters that have been colored the same are all mutually compatible, so can be represented by a single tree corresponding to that color. Finally, whenever a subset of clusters pertaining to a color does not cover all elements of \mathcal{X} , the missing taxa can be attached to the root. An

2. The presence or absence of the singleton clusters in the input does not change the complexity of the problems we study because it is trivial to modify a network without raising its reticulation number or level such that it also represents all the singleton clusters.




Fig. 6. Consider the set of clusters \mathcal{C} shown above. The maximal ST-sets are $\{a\}, \{b\}, \{c\}, \{d\}, \{e\}, \{f\}, \{g\}, \{h, i, j\}$. $\{h, i, j\}$ is the only nonsingleton maximal ST-set, while $\{h, i\}$ is the only nonsingleton, nonmaximal ST-set. Note that, if cluster $\{h, i\}$ was removed, $\{h, i, j\}$ would remain the only nonsingleton maximal ST-set, but $\{h, i\}$ and $\{i, j\}$ would both be nonsingleton, nonmaximal ST-sets.

obvious corollary of this is that the chromatic number of $IG(\mathcal{C})$ is a lower bound on the cardinality of \mathcal{T} .

Whenever $\mathcal{C} = Cl(\mathcal{T})$ there is an important relationship between the nodes and edges of trees in \mathcal{T} , and the (maximal) ST-sets of \mathcal{C} . Let T be a (not necessarily binary) tree on \mathcal{X} . In Section 2, we defined when an edge of a tree represents a cluster. Here, we extend this definition to *nodes* of trees. We say that a node v of T represents C if C is equal to the union of the clusters represented by some (not necessarily strict) subset of its outgoing edges. Note that if an edge (u, v) represents a cluster C then so does v .

Lemma 6. *Let $\mathcal{T} = \{T_1, \dots, T_m\}$ be a set of trees on \mathcal{X} . Let $\emptyset \subset \mathcal{X}' \subset \mathcal{X}$ be compatible with $Cl(\mathcal{T})$. Then, for each $T_i \in \mathcal{T}$ there exists an edge e_i or a node v_i in T_i such that e_i or v_i represents \mathcal{X}' .*

Proof. Consider an arbitrary tree $T_i \in \mathcal{T}$. Every cluster in $Cl(T_i)$ is either disjoint from \mathcal{X}' , a superset of it, or a subset of it, otherwise \mathcal{X}' would be incompatible with $Cl(\mathcal{T})$. If some edge of T_i represents \mathcal{X}' then we are done. Otherwise, consider some node v furthest from the root which represents a cluster C such that $\mathcal{X}' \subset C$. Such a v must exist because if necessary we can take the root as v . C is equal to the union of the clusters represented by some subset of the edges outgoing from v . Each cluster represented by an outgoing edge of v is either disjoint from \mathcal{X}' or a subset of it, because of the assumption on the distance of v from the root. For the same reason, \mathcal{X}' intersects with at least two such outgoing edge clusters of v . But when \mathcal{X}' intersects with such a cluster it must contain it entirely, so \mathcal{X}' is equal to the union of some subset of the clusters represented by edges outgoing from v . Hence, v represents \mathcal{X}' . \square

The following corollary is automatic.

Corollary 5. *Let $\mathcal{T} = \{T_1, \dots, T_m\}$ be a set of binary trees on \mathcal{X} . Let $\emptyset \subset \mathcal{X}' \subset \mathcal{X}$ be a set compatible with $Cl(\mathcal{T})$. Then, for each $T_i \in \mathcal{T}$ there exists an edge e_i such that e_i represents \mathcal{X}' .*

Note that Lemma 6 and Corollary 5 hold in particular for (maximal) ST-sets of $Cl(\mathcal{T})$, because all ST-sets are compatible with $Cl(\mathcal{T})$. Hence, the two following straightforward extensions to ST-sets, which we will use extensively in the next section. Recall that an ST-set S is trivial if $S = \emptyset$ or $S = \mathcal{X}$.

Corollary 6. *Let $\mathcal{T} = \{T_1, \dots, T_m\}$ be a set of binary trees on \mathcal{X} . Let S be a nontrivial ST-set with respect to $Cl(\mathcal{T})$. Then, for each $T_i \in \mathcal{T}$ there exists an edge e_i in T_i such that e_i represents S .*

Corollary 7. *Let $\mathcal{T} = \{T_1, \dots, T_m\}$ be a set of binary trees on \mathcal{X} . Then, $Cl(\mathcal{T})$ contains at most $2(n-1)$ nontrivial ST-sets, and for every such ST-set S of $Cl(\mathcal{T})$ and every tree $T_i \in \mathcal{T}$ there exists a unique edge e_i of T_i such that $\mathcal{X}(e_i) = S$ and such that the subtree rooted at the head of e_i is the unique tree that represents exactly the cluster set $Cl(\mathcal{T})|S$.*

Proof. Given any tree $T_i \in \mathcal{T}$, $Cl(T_i)$ will contain exactly $2(n-1)$ edges and consequently $2(n-1)$ clusters. Since an ST-set of $Cl(T_i)$ is by definition compatible with $Cl(T_i)$, each ST-set of $Cl(\mathcal{T})$ is a subset of the cluster set $Cl(T_i)$ (for any i) and there are thus at most $2(n-1)$

ST-sets. Moreover, by definition, for any ST-set S of $Cl(\mathcal{T})$ we have that the set $Cl(\mathcal{T})|S$ is compatible. Since two nonisomorphic binary trees on the same taxa set induce at least two incompatible clusters, this concludes the proof. \square

Informally, Corollary 7 states that when \mathcal{T} consists solely of binary trees, each nontrivial ST-set corresponds to some subtree that is common to all trees in \mathcal{T} .

5 CLUSTERS OBTAINED FROM SETS OF BINARY TREES ON \mathcal{X}

Let \mathcal{T} be a set of binary trees on \mathcal{X} . In this section, we prove that, for fixed $r \geq 0$, it is possible to construct in polynomial time all binary phylogenetic networks with reticulation number r that represent $Cl(\mathcal{T})$. We also describe our new program CLUSTISTIC which implements this algorithm. CLUSTISTIC is in itself an important “proof of concept:” it has been rapidly prototyped by, building on insights from [8], slightly modifying existing software that was originally conceived to reconstruct binary level- k networks not from clusters but rooted triplets.

Let N be a network on \mathcal{X} and let T' be some tree on $\mathcal{X}' \subset \mathcal{X}$. We say that T' is a *Subtree Below a Reticulation* (SBR) of N if there is a reticulation node v in N such that no reticulation nodes $v' \neq v$ are reachable from v by a directed path, and that the subnetwork rooted at v (or, when v has outdegree exactly 1, the child of v) is exactly equal to T' . It is easy to show that (by virtue of its acyclicity) every network contains at least one SBR [21]. A simple though critical observation is

Observation 2. *If T' is an SBR of a network N on \mathcal{X} that represents a cluster set \mathcal{C} , and T' has taxa set \mathcal{X}' , then \mathcal{X}' is an ST-set with respect to \mathcal{C} .*

We will make repeated use of this throughout the rest of the paper. (Note however that, in general, multiple SBRs might correspond to a single ST-set).

Given a network N with an SBR T we denote by $N \setminus T$ the network obtained from N by deleting T and for as long as necessary applying the following tidying-up operations until they are no longer needed: deleting any node with outdegree zero that is not labeled by an element of \mathcal{X} ; suppressing all nodes with indegree and outdegree both equal to 1; replacing multiedges with single edges; deleting nodes with indegree-0 and outdegree-1, see Fig. 7.

We call (S_1, S_2, \dots, S_p) ($p \geq 0$) a (maximal) ST-set sequence of \mathcal{C} if S_1 is a (maximal) ST-set of \mathcal{C} , S_2 is a (maximal) ST-set of $\mathcal{C} \setminus S_1$, S_3 is a (maximal) ST-set of $\mathcal{C} \setminus S_1 \setminus S_2$ and so on. Such a sequence is additionally a *tree* sequence if all the clusters in $\mathcal{C} \setminus S_1 \setminus \dots \setminus S_p$ are mutually compatible, i.e., can be represented by a tree. We denote by p be the *length* of the sequence; $p = 0$ denotes the empty tree sequence.

Lemma 7. *Let N be a network that represents some cluster set \mathcal{C} on \mathcal{X} . Then, there exists a sequence of SBRs that need to be removed to prune N into a tree, and this corresponds to an ST-set tree sequence of \mathcal{C} of length $r(N)$.*




Fig. 7. Let \mathcal{C} be some set of clusters represented by the network on the right. If we guess that ST-set $S = \{h, i, j\}$ corresponds to an SBR and remove it, we obtain the tree on the left. To reverse the process we add (a tree corresponding to) S below a reticulation node whose incoming edges subdivide edges x and y .

Proof. If N is a tree then there is an empty tree sequence. Otherwise, N has an SBR T' with taxa set \mathcal{X}' and the network $N \setminus_{T'}$ represents the cluster set $\mathcal{C} \setminus \mathcal{X}'$. Clearly $r(N') < r(N)$. By observation 2, we let S_1 equal \mathcal{X}' and $d = r(N) - r(N')$. If $d > 1$ then we let S_2, \dots, S_d all equal \emptyset , the empty ST-set. We use these empty ST-sets to model the situation when removing the SBR would cause multiple reticulation nodes to disappear simultaneously. Now, N' also has at least one SBR, so we can iterate the whole process. We can repeat this until we obtain a tree: at this point we will have an ST-set tree sequence of length exactly $r(N)$. \square

In the following observation, we make use of the fact that the operation $N \setminus_T$ is also meaningfully defined when the tree T is exactly the subtree rooted at some node of N .

Observation 3. Let $\mathcal{T} = \{T_1, \dots, T_m\}$ be a set of binary trees on \mathcal{X} . Let S be a nontrivial ST-set of $Cl(\mathcal{T})$. Then, $Cl(\mathcal{T}) \setminus S = Cl(\mathcal{T}')$ where \mathcal{T}' is a set of at most m binary trees $\{T'_1, \dots, T'_m\}$ on $\mathcal{X} \setminus S$ with $T'_i = T_i \setminus T_v$, where $e_i = (u, v)$ is the edge of T_i that represents S (which exists by Corollary 7) and T_v is the subtree rooted at v .

Theorem 2. Let $\mathcal{T} = \{T_1, \dots, T_m\}$ be a set of binary trees on \mathcal{X} . Then, for a constant $r \geq 0$ it is possible to construct in polynomial time all binary networks with reticulation number at most r that represent $Cl(\mathcal{T})$ (if any exist).

Proof. Without loss of generality assume we wish to construct all such networks with reticulation number exactly r . Let us suppose that at least one such network N exists. By Lemma 7, there is an ST-set tree sequence $\mathcal{S} = (S_1, \dots, S_r)$ for $Cl(\mathcal{T})$ and this corresponds to a sequence of SBRs that, when removed, will prune N into a tree. Let $|X| = n$. Now, note that by Observation 3 and Corollary 7 there are at most $O(n^r)$ ST-set tree sequences, which is polynomial in n for fixed constant r . \mathcal{S} will be one of these, so we can find \mathcal{S} in polynomial time. It remains to show

how, assuming we have found \mathcal{S} , we can reconstruct N . First, note that for a binary tree T on \mathcal{X} , T is the *unique* tree on \mathcal{X} that represents $Cl(T)$. The clusters that still remain after removing S_r can be represented by a tree, and in particular (by Observation 3) by a unique binary tree. We call this tree N_r . We want to obtain N_{r-1} by inserting (a tree corresponding to) S_r into N_r . In particular, we wish to introduce a new reticulation node into N_r below which a tree T (itself binary and unique by Corollary 7) that represents S_r will be attached. We have two possibilities to do this: we subdivide two (not necessarily distinct) edges and use these as the tails of the new reticulation edges, see Figs. 8a and 8b, or we subdivide one edge *once* and use two *identical* reticulation edges, see Fig. 8c. Note that, if we only subdivide one edge once, we actually create a multiedge and by our definition of phylogenetic network such edges are not allowed. However, it might be necessary to create such multiedges during *intermediate* iterations to ensure that all phylogenetic networks, including “redundant” ones, are constructed.

Unfortunately, we do not know in general which edge(s) of N_r to subdivide to create the new connection point(s). However, there are only polynomially many edges in N_r , so we simply try them all. Then we repeat the process, inserting S_{r-1} into N_{r-1} to obtain N_{r-2} , and so on until we have obtained N_0 . Given that r is a constant we can test in polynomial time whether N_0 represents all the clusters in \mathcal{C} (see proof of Lemma 3).

Just as in the similar algorithm described in [21] there are several slight technicalities that should be noted. Whenever some $S_i = \emptyset$ we use a “dummy” taxon (i.e., some taxon not in \mathcal{X}) as the tree that we attach below a reticulation. The function of this is to ensure that subsequent iterations can subdivide the edge leaving the reticulation, i.e., it is a placeholder. (This will be necessary when the removal of a single SBR caused the disappearance of two or more reticulations). These dummy taxa can be removed just before N_0 is inspected to check whether it represents \mathcal{C} . Any N_0 that at this point still contains a dummy taxon whose parent is a reticulation, should be rejected, because it means at least one reticulation was not used. Second, when we construct the tree N_r we actually add a “dummy root” which is simply a new node ρ' and a single edge from (ρ', ρ) where ρ is the root of N_r . This deals with the situation when the removal of some SBR caused the current root to disappear and a new root to take its place. At the end, ρ' and the edge leaving it should be removed. Finally, note that any multiedges




Fig. 8. The different ways of adding a reticulation back into a network, as discussed in the proof of Theorem 2. (a) Two different edges are subdivided; (b) one edge is subdivided twice; (c) one edge is subdivided once, under which a multiedge is placed.

created by intermediate iterations of the algorithm should all have disappeared (i.e., have been subdivided by reticulation edges) by the time N_0 has been reached; for this reason we reject any N_0 that still contains multiedges.

The network N we would like to reconstruct will eventually be found as some N_0 , and given that we made no assumptions about N this shows that the algorithm constructs all possible N . \square

Corollary 8. Let $\mathcal{T} = \{T_1, \dots, T_m\}$ be a set of binary trees on \mathcal{X} . Then, for a constant $k \geq 0$ it is possible to construct in polynomial time all binary simple level- $\leq k$ networks that represent $Cl(\mathcal{T})$ (if any exist).

Proof. For each network produced by the algorithm described in Theorem 2, we can easily check in polynomial time whether it is biconnected. \square

It is worth noting at this stage an important link with the rooted triplet literature. Recall the following proposition and lemma from [8], a paper in which the relationship between trees, clusters and triplets was discussed more broadly. Proposition 1 refers to not necessarily binary trees.

Proposition 1 (Van Iersel, Kelk [8]). For any set \mathcal{T} of trees on the same set \mathcal{X} of taxa, any phylogenetic network on \mathcal{X} representing $Cl(\mathcal{T})$ is consistent with $Tr(\mathcal{T})$.

Lemma 8 (Van Iersel, Kelk [8]). Let N be a phylogenetic network on \mathcal{X} and \mathcal{T} a set of binary trees on \mathcal{X} . Then, there exists a binary phylogenetic network N' on \mathcal{X} such that

1. N' has the same reticulation number and level as N ,
2. if N displays all trees in \mathcal{T} then so too does N' ,
3. if N is consistent with $Tr(\mathcal{T})$ then so too is N' and
4. if N represents $Cl(\mathcal{T})$ then so too does N' .

Suppose that we have an Algorithm A which, for each fixed $r \geq 0$, can construct in polynomial time every binary network consistent with $Tr(\mathcal{T})$ that has at most r reticulations, where \mathcal{T} is a set of binary trees on \mathcal{X} . Suppose A' is an algorithm that examines in turn every network output by A and rejects it if it does not represent $Cl(\mathcal{T})$ (such a filtering step can be done for each network in polynomial time for fixed r , see proof of Lemma 3). If there exists a network N on \mathcal{X} with at most r reticulations that represents $Cl(\mathcal{T})$ then by Lemma 8 there exists a binary network N' on \mathcal{X} with this property. Furthermore, in that case N' will by Proposition 1 be consistent with $Tr(\mathcal{T})$. The algorithm A' is thus guaranteed to eventually find N' . The practical consequence of this is that, simply by adding a filtering step, triplet software can in some cases easily be modified to work for clusters. Indeed, they can be used to determine whether a network with r reticulations that represents \mathcal{C} exists and if so to construct all binary networks with this property. As proof of concept, we have taken the triplet software SIMPLISTIC (based on the ideas described in [21]), removed its biconnectedness-checking subroutine so that it generates all binary networks with up to r reticulations (and not just all simple level- $\leq r$ binary networks), and added the cluster filtering step as described above. This whole process took only one day of programming, and led to the new software package CLUSTISTIC which implements




Fig. 9. Let \mathcal{T} be the set of four trees shown here. The CASS algorithm returns a network N that represents $Cl(\mathcal{T})$ where $r(N) = \ell(N) = 4$. However, Fig. 10 shows that the true value of $r(Cl(\mathcal{T})) = \ell(Cl(\mathcal{T}))$ is at most 3.

the result described in Theorem 2. This software is available for download at <http://skelk.sdf-eu.org/clustistic>.

6 WITNESSES AND A NATURAL LOWER BOUND

The results in this section will help us develop some insights into how a phylogenetic network can be efficiently dismantled (respectively, reconstructed) by iteratively removing (respectively, reattaching) reticulations. These results, and the bridges they build with other parts of the phylogenetic network literature, will be used heavily in Section 7.

Let $\mathcal{T} = \{T_1, \dots, T_m\}$ be a set of m not necessarily binary trees on \mathcal{X} . Let S be a nontrivial ST-set of $Cl(\mathcal{T})$. We know by Lemma 6 that for each $T_i \in \mathcal{T}$ there is an edge e_i or node v_i in T_i that represents S . We define a *witness* for S in T_i as follows: If T_i contains an edge $e_i = (u_i, v_i)$ that represents S , let $t_i = u_i$. Otherwise, from Lemma 6 there exists a node v_i in T_i that represents S , and let $t_i = v_i$. Then, a witness for S is any leaf descendant $w_j \in (\mathcal{X}(t_i) \setminus S)$. The only ST-sets with no witnesses are \mathcal{X} and the empty set, hence the restriction to nontrivial ST-sets. As an example consider the four trees \mathcal{T} in Fig. 9. Consider the ST-set $\{1\}$ of $Cl(\mathcal{T})$. In the top-left tree and bottom-left tree the only possible witness for this is taxon 5. In the top-right tree the only witness is taxon 3, and in the bottom-right tree the only witness is taxon 2.

Given a set of trees \mathcal{T} on \mathcal{X} and a nontrivial ST-set S of \mathcal{X} , let $W \subseteq \mathcal{X}$ be any subset of taxa such that, for each tree $T_i \in \mathcal{T}$, there exists $x \in W$ that is a witness for S in T_i . We call such a set a *witness set* of S in \mathcal{T} . Clearly, there exist W such that $|W| \leq m$. For example, for the set of trees in Fig. 9, $\{2, 3, 5\}$ is a possible witness set for $\{1\}$.

The two following simple observations are critical.

Observation 4. Let \mathcal{T} be a set of trees on \mathcal{X} , S a nontrivial ST-set of \mathcal{X} and W a witness set of S in \mathcal{T} . Then, for each $C \in \mathcal{C}(\mathcal{T})$ such that $S \subset C$, $W \cap C \neq \emptyset$.

Proof. For each cluster $C \in \mathcal{C}(\mathcal{T})$ there is at least one edge $e = (u, v)$ in some $T_i \in \mathcal{T}$ such that e represents C . Let $e_i(v_i)$ be the edge (node) in T_i that represents S . Given that $S \subset C$, all leaf descendants of v must also include all leaf descendants of the tail of $e_i(v_i)$. In particular, all possible witnesses for S in T_i . \square

To understand the meaning of Observation 4 it is helpful to again consider the example of ST-set $\{1\}$ in the context of Fig. 9. We see that any cluster that is a strict superset of $\{1\}$ must contain at least one of the taxa from $\{2, 3, 5\}$.

Observation 5. Let $\mathcal{T} = \{T_1, \dots, T_m\}$ be a set of trees on \mathcal{X} such that $m \geq 2$ and let S be a nontrivial ST-set of \mathcal{C} . Let W be a smallest cardinality witness set of S in \mathcal{T} . If $|W| = m$ then for each $C \in \mathcal{C}(\mathcal{T})$ such that $S \cap C = \emptyset$, $W \setminus C \neq \emptyset$.

Proof. Note that it is not possible for a witness for S in a tree $T_i \in \mathcal{T}$ to also be a witness for S in $T_j \neq T_i$, because then $|W| \leq m - 1$. So each witness in W comes from a different tree in \mathcal{T} . Suppose then that there is some $C \in \mathcal{C}(\mathcal{T})$ such that $S \cap C = \emptyset$ and $W \setminus C = \emptyset$. Clearly, since $W \neq \emptyset$, $W \subseteq C$, and $|C| \geq |W| \geq 2$. Furthermore, some edge e_i in some T_i represents C . Combining the fact that $S \cap C = \emptyset$ and $W \subseteq C$ leads us to the conclusion that all elements of W are possible witnesses for S in T_i . But then some element x of W is a witness for S both in T_i and also in some $T_j \neq T_i$, contradiction. \square

As mentioned in the proof of Observation 5, a witness set W (for a given ST-set S) with $m - 1$ or fewer elements exists if and only if the possible witnesses for S ranging across the different T_i are not all mutually disjoint. If a witness set with $m - 1$ or fewer elements does not exist then in the following results any witness set with m elements will turn out to be sufficient, as a consequence of Observation 5. This will become clear in due course.

6.1 A Natural Lower Bound on $r(\mathcal{T})$ that Is Tight for Clusters Obtained from Two (Not Necessarily Binary) Trees

Lemma 9. Given a set of clusters \mathcal{C} on \mathcal{X} , there exists a maximal ST-set tree sequence (S_1, S_2, \dots, S_p) such that $p \leq r(\mathcal{C})$.

Proof. The proof is equivalent to that of Lemma 7 but for the fact that no empty ST-set is inserted in the maximal ST-set tree sequence. \square

We define the *maximal ST-set lower bound* for \mathcal{C} (MST lower bound for short) as the cardinality of the smallest maximal ST-set tree sequence for \mathcal{C} . By Lemma 9 this is a genuine lower bound on $r(\mathcal{C})$. In general, it is however a rather weak lower-bound: consider the set \mathcal{C}^i of clusters on $X^i = \{r, x_1, \dots, x_i\}$ defined by $\{\{r, x_j\} \mid 1 \leq j \leq i\}$. The MST lower bound for this cluster set is always 1, while $r(\mathcal{C}^i)$ rises linearly in i . However, as we shall see the tightness of the bound is to some extent correlated with the number of trees which generate the clusters, with two trees being a special case. We first require some definitions and auxiliary lemmas.

Consider a network N on \mathcal{X} . Let $\mathcal{X}' = \{\rho\} \cup \mathcal{X}$ where $\rho \notin \mathcal{X}$ is some arbitrary symbol representing the root of N . Let T be some tree on \mathcal{X}^* where $\mathcal{X}^* \cap \mathcal{X}' = \emptyset$ and let H be a subset of \mathcal{X}' . We can obtain a new network N' on $(\mathcal{X}^* \cup \mathcal{X})$ by hanging T from H in N . Informally, N' is obtained by hanging the tree T beneath a new reticulation which has $|H|$ incoming edges, where each such incoming edge begins “just above” an element of H . Formally, the transformation is as follows: First, we add a new edge (r, r') to T , where r' is the root of T and r is a new node. For each $h \in H \setminus \{\rho\}$,

we then subdivide the unique edge entering h ; let h_p be the new parent (with indegree and outdegree 1) of h . For each $h \in H \setminus \{\rho\}$, we then add a new edge (h_p, r) . Finally, if $\rho \in H$ we also add an edge from the root of N to r ; we call this a *root edge*. It is clear that $r(N') = r(N) + |H| - 1$.

Lemma 10. Let $\mathcal{T} = \{T_1, \dots, T_m\}$ be a set of m trees on \mathcal{X} and let $\mathcal{C} = Cl(\mathcal{T})$. Let S be a maximum ST-set of \mathcal{C} . Let N be any network on $\mathcal{X} \setminus S$ that represents $\mathcal{C} \setminus S$. Then, it is possible to extend N to obtain a new network N' that represents \mathcal{C} such that $r(N') \leq r(N) + (m - 1)$.

Proof. Let T_S be the unique tree on taxa set S such that $\mathcal{C}(T_S) = \mathcal{C}|S$. Let W be a minimum-cardinality witness set for S in \mathcal{T} . Clearly, $1 \leq |W| \leq m$. If $|W| = m$, then we let N' be the network obtained by hanging T_S from W in N . If $|W| < m$ then we let N' be the network obtained by hanging T_S from $W \cup \{\rho\}$ in N . Clearly, $r(N') \leq r(N) + (m - 1)$. It remains only to show that N' represents \mathcal{C} . Consider any cluster $C \in \mathcal{C}$. There are three cases to consider 1) If $C \subseteq S$ then N' clearly represents C because T_S already represented $\mathcal{C}|S$. 2) If $S \subset C$ then consider $C' = C \setminus S$. Clearly, N represents C' . By Observation 4, there exists some $w \in W \cap C$. Since $W \cap S = \emptyset$, this implies that there exists some $w \in W \cap C'$.

To see that N' represents C consider any tree displayed by N that represents C' . We can extend this tree by “switching on” the new reticulation edge that begins above w , i.e., the edge (w_p, r) , and “switching off” the remaining reticulation edges. 3) If $S \cap C = \emptyset$, then there are two subcases. a) If $|W| < m$, then we can “switch on” the root edge that enters the reticulation above T_S , i.e., the edge (ρ, r) , and “switch off” all other reticulation edges entering T_S . b) If $|W| = m$, then by Observation 5 there exists $w \in W \setminus C$. In N' , we can thus “switch on” the new reticulation edge that begins above w , and “switch off” the rest. \square

Theorem 3. Let $\mathcal{T} = \{T_1, \dots, T_m\}$ be a set of m trees on \mathcal{X} and let $\mathcal{C} = Cl(\mathcal{T})$. Let p be the MST lower bound for \mathcal{C} . Then, $r(\mathcal{C}) \leq (m - 1)p$.

Proof. Given a tree T and a node u of T , we denote by $\mathcal{X}(T)$ the label set of T and by T_u the subtree rooted at u . From Lemma 9, we already know that $p \leq r(\mathcal{C})$. Now, let (S_1, S_2, \dots, S_p) be a maximal ST-set tree sequence for \mathcal{C} . We will complete the proof by showing how to explicitly construct a network N with reticulation number at most $(m - 1)p$ that represents \mathcal{C} . We define C_i , $1 \leq i \leq p$, as $C \setminus S_1 \setminus \dots \setminus S_i$ and C_0 as \mathcal{C} . By Lemma 6, it can be seen that for each i , $C_i = Cl(\mathcal{T}_i)$ where \mathcal{T}_i is a set of at most m trees on $\mathcal{X} \setminus S_1 \setminus \dots \setminus S_i$ and where $T_0 = \mathcal{T}$. In particular, \mathcal{T}_{i+1} can be obtained from \mathcal{T}_i as follows: Given a tree T_j in \mathcal{T}_i , let (without loss of generality) u_j be the node of T_j such that S_i is equal to the union of the clusters represented by some not necessarily strict subset of its outgoing edges. Such a u_j exists by Lemma 6. Let $Q = \{v_1, \dots, v_k\}$ be the set of children of u_j such that for each $v \in Q$, $\mathcal{X}(T_v)$ contains at least one element of S_i . The set of trees \mathcal{T}_{i+1} can be obtained from \mathcal{T}_i by computing, for each tree T_j in \mathcal{T}_i the tree $T_j \setminus T_{v_1} \setminus \dots \setminus T_{v_k}$, i.e., pruning away the subtrees corresponding to S_i and tidying up the resulting tree.

Now, consider \mathcal{C}_p . Let N_p be the unique tree such that $\mathcal{C}(N_p) = \mathcal{C}_p$; N_p will be equal to the single tree in \mathcal{T}_p . By Lemma 10, we can obtain a network N_{p-1} with $(m-1)$ reticulations that represents \mathcal{C}_{p-1} by taking $\mathcal{T} = \mathcal{T}_{p-1}$, $N = N_p$, and $S = S_p$ in the proof of that lemma. We iterate this process for $p-2, p-3, \dots, 1$. This lasts at most p iterations in total, and each iteration adds $(m-1)$ to the reticulation number, thus, yielding a network N_0 that represents \mathcal{C} with reticulation number (at most) $(m-1)p$. \square

Corollary 9. Let $\mathcal{T} = \{T_1, T_2\}$ be a set of two not necessarily binary trees on \mathcal{X} and let $\mathcal{C} = Cl(\mathcal{T})$. Let p be the MST lower bound for \mathcal{C} . Then, $r(\mathcal{C}) = p$.

In [8], it is shown that it is NP-hard and APX-hard to compute $r(\mathcal{C})$ where \mathcal{C} is the set of clusters obtained from two binary trees on \mathcal{X} . The following corollary is thus immediate.

Corollary 10. The computation of the MST lower bound is NP-hard and APX-hard.

It is interesting to note that Lemma 9 and Corollary 10 have, in some sense, already appeared in the phylogenetic network literature, albeit in the language of *recombination networks*. Specifically, in [8] we highlight that the phylogenetic network model described there (and also used here) is in a strong sense identical to the recombination network model under the assumption of an all-0 root, the infinite sites model and multiple crossover recombination. The computational lower bound described in [25, Algorithm 3] is, taking this equivalence into account, essentially identical to the MST lower bound. In [32], it is shown that computing this bound is NP-hard, by reduction from MAX-2-SAT. The same authors also give an exponential-time dynamic programming algorithm for computing the bound, improving upon an algorithm with superexponential running time given in [25].

7 THE OPTIMALITY AND NONOPTIMALITY OF CASS

The CASS algorithm for constructing simple level- k networks was presented in [27]. The algorithm was designed to produce solutions of minimum level, not of minimum reticulation number. However, when the input is a separating set \mathcal{C} of clusters on \mathcal{X} , minimizing the level or the reticulation number is equivalent. Indeed, such cluster sets have the property that any network that represents them is simple or can easily be made simple (see Lemma 1) and a simple network contains exactly one nontrivial biconnected component.

The CASS algorithm can be used as a subroutine in a divide and conquer algorithm to construct general level- k networks. We will call this more general algorithm $CASS^{DC}$. The basic idea of $CASS^{DC}$ is that it transforms each connected component of $IG(\mathcal{C})$ into a tangled set of clusters, runs CASS separately on each of these tangled sets, and combines the resulting simple networks into a single final network N (Recall that tangled cluster sets are separating). The final network N has reticulation number equal to the sum of the reticulation number of the simple

networks produced by CASS, and N has level equal to the maximum level ranging over all the simple networks. For more details on the divide and conquer strategy, see [27] and [4, Section 8.2].

In [27], the authors proved that if there is a level- ≤ 2 network that represents \mathcal{C} , then $CASS^{DC}$ will find such a solution with minimal level. Here, we clarify several other properties of the algorithm. On the negative side we show (using a special separating cluster set) that the CASS algorithm does not in general minimize level. On the positive side we show that when the input set \mathcal{C} is equal to $Cl(\{T_1, T_2\})$ for any two (not necessarily binary) trees, $CASS^{DC}$ correctly minimizes level. In fact, we show something even stronger: in this case $CASS^{DC}$ also correctly constructs networks with minimum reticulation number, which in turn is exactly equal to the hybridization number of the two input trees, i.e., the number of reticulations required to display the trees themselves. We conclude with several open questions regarding CASS.

7.1 Cass: The High-Level Idea

Let N be a network that represents a set of clusters \mathcal{C} on \mathcal{X} . Let S be a nontrivial ST-set with respect to \mathcal{C} . We say that S is *under a cut-edge* if N contains a cut-edge (u, v) such that the subnetwork rooted at v is a tree that represents $\mathcal{C}|S$.

The next two results formed (implicitly) the direct inspiration for CASS, which was designed to be a generalization of these results.

Lemma 11. Let N be a network that represents a set of clusters on \mathcal{X} . Let S be a nontrivial ST-set with respect to \mathcal{C} . Then, there exists a network N' such that $r(N') \leq r(N)$, $\ell(N') \leq \ell(N)$, S is under a cut-edge in N' and for each ST-set S' such that $S' \cap S = \emptyset$ and S' is under a cut-edge in N , S' is also under a cut-edge in N' .

Proof. Deferred to the Appendix, which can be found on the Computer Society Digital Library. \square

The following corollary follows from the fact that maximal ST-sets are disjoint:

Corollary 11. Let N be a network that represents a set of clusters \mathcal{C} . There exists a network N' such that $r(N') \leq r(N)$, $\ell(N') \leq \ell(N)$ and all maximal ST-sets (with respect to \mathcal{C}) are below cut-edges.

The pseudocode for CASS was originally given in [27]. That exposition is however rather dense and technical. See [4, Section 8.5] for a clearer detailed description. Here, we only give the core idea of the algorithm. Let us assume without loss of generality that we want to know, given a separating set of clusters, whether a simple network solution exists with reticulation number *exactly* k , for some constant k .

CASS tries to answer this by searching through the space of all maximal ST-set tree sequences of length at most k , attempting to build a network with reticulation number k from each one. It looks first at shorter maximal ST-set tree sequences, padding those of length less than k with empty ST-sets to attain a sequence of length exactly k . (As in the proof of Theorem 2, this models the situation when removing a single SBR causes the reticulation number to drop by more than 1). If there are no maximal ST-set tree

sequences of length at most k then CASS will correctly report that no solutions with reticulation number k or lower exist. Hence, CASS implicitly computes and incorporates the MST lower bound.

Assuming maximal ST-set tree sequences of length at most k do exist, CASS examines each one to determine whether it can be constructively turned into a real solution. Let $\mathcal{S} = (S_1, \dots, S_k)$ be a (possibly padded) maximal ST-set tree sequence of length- k for \mathcal{C} . As in earlier sections we define \mathcal{C}_i , $1 \leq i \leq k$, as $\mathcal{C} \setminus S_1 \setminus \dots \setminus S_i$, and we let $\mathcal{C}_0 = \mathcal{C}$. CASS does not however work with the set \mathcal{C}_i . Instead it works with \mathcal{C}'_i which is obtained from each \mathcal{C}_i by “collapsing” every maximal ST-set S with respect to \mathcal{C}_i into a single new “metataxon.” This is an extremely greedy step, and is in some sense an attempt to generalize Corollary 11. The informal motivation is this: we know from Corollary 11 that there exists some network N with a minimum number of reticulations such that all maximal ST-sets are under cut-edges. Suppose we guess an SBR T on $\mathcal{X}' \subset \mathcal{X}$ (where \mathcal{X}' is a maximal ST-set) of N ; we only have to make polynomially many guesses because there are only polynomially many maximal ST-sets. This gives a new network N' on $\mathcal{X} \setminus \mathcal{X}'$ where perhaps not all maximal ST-sets (with respect to $\mathcal{C} \setminus \mathcal{X}'$) are under cut-edges. In particular, some SBRs of N' might correspond to nonmaximal ST-sets, of which there are potentially exponentially many, so how do we efficiently guess an SBR of N' ? Fortunately, we can transform N' (in the sense of Corollary 11) to obtain a new network N'' such that $r(N'') \leq r(N')$ and where all maximal ST-sets (with respect to $\mathcal{C} \setminus \mathcal{X}'$) are under cut-edges of N'' . Hence, we know that we again only have to make polynomially many guesses to locate an SBR of N'' . Furthermore, CASS assumes that these maximal ST-sets will always remain below cut edges, so it collapses them into the aforementioned metataxa. We iterate this entire process k times, concluding the “inward” phase of CASS.

This assumption is important because it affects the “outward” phase of CASS, which begins immediately after completion of the inward phase. As in, for example, Theorem 2 the general idea is to start with a tree that represents \mathcal{C}_k and then to work backward, first trying all pairs of edges from which to “hang back” a SBR corresponding to S_k , then all pairs of edges (of the resulting network) from which to hang back an SBR corresponding to S_{k-1} , and so on, down to S_1 . However, before hanging back each S_i it first *decollapses* the maximal ST-sets that were collapsed into metataxa during the corresponding iteration of the inward phase.

In Section A.2 of the Appendix, which can be found on the Computer Society Digital Library, we will explicitly and exhaustively walk through a specific execution of the CASS algorithm and this is helpful for clarifying exactly how the algorithm works.

7.2 The CASS Algorithm Is Not Always Optimal

In this section, we present a counter example proving that the CASS algorithm does not always minimize level.

Consider the set \mathcal{T}_4 of four binary trees shown in Fig. 9. It is easy to verify that the set $\mathcal{C} = Cl(\mathcal{T})$ is separating and every network that represents \mathcal{C} is thus simple or can easily be made simple. It is also easy to verify that the simple level-3 network in Fig. 10a represents \mathcal{C} ; in the Appendix, which can be found on the Computer Society Digital




Fig. 10. (a) A simple level-3 network and (b) a simple level-4 network, both representing $Cl(\mathcal{T})$ where \mathcal{T} is defined as described in Fig. 9. The level-4 network was produced by CASS^{DC}.

Library, we prove that this is optimal by showing that any network that represents \mathcal{C} must have reticulation number 3 or higher.

However, CASS cannot find a level-3 network. To formally prove this we show in Section A.2 of the Appendix, which can be found on the Computer Society Digital Library, an exhaustive list of all possible executions of the algorithm with $k = 3$. CASS returns the simple level-4 network shown in Fig. 10b. To summarize, the problem with CASS seems to be that while the step of always collapsing at every iteration all maximal ST-sets and treating them as metataxa (in the sense of Corollary 11) is a locally optimal move, it can force us to use too many reticulation edges when hanging (trees corresponding to) maximal ST-sets back in the outward phase.

Note that this counter example fits in a tradition of highly specific and complex counter examples in the phylogenetic network literature. In particular, we note the very similar counter examples given initially in [22], and (based on this) in [7], which showed that one cannot minimize reticulation number by optimizing independently over the connected components of the incompatibility graph $IG(\mathcal{C})$.³ The relationship between these counter examples—all of which are linked in some way or other to simple level-3 networks—seems to be that networks with minimum reticulation number have a very subtle internal structure that seems impervious to locally optimal and greedy strategies, but that these properties only start emerging for level-3 and higher.

It is important to emphasize, however, that since we could not find any nonsynthetic data sets for which CASS does not find an optimal solution, we still have the feeling that CASS works quite well for real data.

7.3 CASS Is Optimal for Sets of Clusters Obtained from Two Trees

Here we show that, despite the negative news in the previous section, CASS (and more generally CASS^{DC}) correctly minimizes level in the case of clusters obtained from two not necessarily binary trees. Furthermore, the algorithms also provably minimize reticulation number. (It is interesting to note that it is not known whether CASS or CASS^{DC} correctly minimizes level in the case of clusters obtained from three trees).

3. In [22], both single crossover and multiple crossover recombination models are considered, where the latter model is very similar to the softwired cluster model, see, e.g., [8]. The network in [22] which resembles Fig. 10a is used in the context of multiple crossover recombination.

Note that the following theorem does not contradict the NP-hardness mentioned in Corollary 10 because CASS only runs in polynomial time when it bounds its search to simple level- $\leq k$ networks, for a constant k .

Theorem 4. Let $\mathcal{C} = Cl(\mathcal{T})$ be a separating set of clusters where \mathcal{T} is a set of two not necessarily binary trees on \mathcal{X} . Then CASS constructs a simple network N that represents \mathcal{C} such that $\ell(N) = r(N) = \ell(\mathcal{C}) = r(\mathcal{C})$.

Proof. Let the MST lower bound for \mathcal{C} be p . Recall that, by Corollary 9, in this case $p = r(N)$. We know that there is a maximal ST-set tree sequence (S_1, \dots, S_p) . As explained in Section 7.1, CASS will eventually find this maximal ST-set tree sequence. Now, Theorem 3 essentially works by invoking Lemma 10 p times, and in the statement of Lemma 10 there are absolutely no assumptions made about the structure of “ N ,” other than that it represents a certain set of clusters. Hence, “ N ” can just as well be one of the intermediate networks constructed by CASS. It remains only to show that CASS can simulate the hanging-back construction described in the proof of Lemma 10. This is definitely so, because the proof of Lemma 10 requires the edges entering two witnesses (or the edge entering one witness and, to simulate the attachment of a root edge, the edge connecting a dummy root to the real root) to be subdivided. CASS tries subdividing all pairs of edges, including the edge between the dummy root and the real root, and hence will eventually subdivide the correct two edges. (Note that this theorem does not hold for three or more trees because per iteration it might be necessary to subdivide three or more edges, i.e., to increase the reticulation number by more than 1). \square

The proof of Theorem 4 not only shows that CASS is optimal for sets of clusters obtained from two not necessarily binary trees, but also that in this very special case the “hanging back” (i.e., outward) phase of CASS is in some sense completely redundant. In particular: if we have already computed a maximal ST-set tree sequence, then we can easily compute a witness set for each of the maximal ST-sets in the sequence, and these witness sets directly specify a sufficient set of edges to subdivide when hanging back the maximal ST-sets. So in the case of clusters coming from two trees CASS wastes rather a lot of time trying to hang back maximal ST-sets from all possible pairs of edges, when in fact the information is already available to make this blind search unnecessary. Theorem 4 can actually be reformulated and extended to general (i.e., not necessarily separating) sets of clusters \mathcal{C} obtained from two not necessarily binary trees. Indeed, in Theorem 6, we will prove that, for such cluster sets, CASS^{DC} reconstructs a network N such that $r(N) = r(\mathcal{C})$ and $\ell(N) = \ell(\mathcal{C})$. However, we first need to prove Theorem 5 below, which is interesting in its own right, and several auxiliary results.

Observation 6. Let \mathcal{C} be a set of clusters on \mathcal{X} and let $U \subseteq \mathcal{X}$ be compatible with \mathcal{C} . Then, for any $P \subseteq \mathcal{X}$, $U \setminus P$ is compatible with $\mathcal{C}|(\mathcal{X} \setminus P)$.

Proof. U is compatible with \mathcal{C} so for each cluster $C \in \mathcal{C}$ we have either that $C \cap U = \emptyset$, $C \subseteq U$, or $U \subseteq C$. For the case $C \cap U = \emptyset$ it is clear that $(C \setminus P) \cap (U \setminus P) = \emptyset$. For the case $C \subseteq U$ we have that $(C \setminus P) \subseteq (U \setminus P)$, and for the case $U \subseteq C$ we have that $(U \setminus P) \subseteq (C \setminus P)$. \square

Observation 7. Let \mathcal{C} be a set of clusters on \mathcal{X} and let U be the union of the set of clusters in a connected component K of $IG(\mathcal{C})$. Then, U is compatible with \mathcal{C} .

Proof. Suppose U is not compatible with \mathcal{C} . Then, there must exist some $C \in \mathcal{C}$ such that $C \not\subseteq U$, $U \not\subseteq C$, and $C \cap U \neq \emptyset$. Clearly, C cannot be incompatible with any cluster in the connected component K , because then C would also be in the connected component K and thus $C \subseteq U$. Hence, every cluster in the connected component K is either disjoint from C , or a subset of it, and there is at least one of each type of cluster because $U \setminus C \neq \emptyset$ and U is the union of all the clusters in K . However, clusters in K that are disjoint from C , are always compatible with clusters in the connected component that are contained inside C , so the connected component is not connected, contradiction. \square

Observation 8. Given a set of clusters \mathcal{C} on \mathcal{X} , let S be an ST-set with respect to \mathcal{C} and let U be a set such that U is compatible with \mathcal{C} and $U \subseteq S$. Then, U is also an ST-set for \mathcal{C} .

Proof. We only need to show that all pairs of clusters in $\mathcal{C}|U$ are compatible. Clearly, for each $C \in \mathcal{C}$ we have that $C|U \subseteq U$. Now, recall that, because U is compatible with \mathcal{C} , for each $C \in \mathcal{C}$ we have either $C \subseteq U$, $U \subseteq C$ or $C \cap U = \emptyset$. Suppose by contradiction that for some $C_1 \neq C_2 \in \mathcal{C}|U$, C_1 and C_2 are incompatible. But then $\emptyset \subset C_1, C_2 \subset U$. But in that case $C_1, C_2 \in \mathcal{C}$ and $C_1, C_2 \subset S$, contradicting the fact that S was an ST-set. \square

For a set of clusters \mathcal{C} and a set U that is compatible with \mathcal{C} , $\mathcal{C}_{U \rightarrow u}$ denotes the new cluster set obtained by replacing all elements of U with a single new taxon u (i.e., “collapsing” U into a single taxon).

Theorem 5. Let $\mathcal{T} = \{T_1, T_2\}$ be two not necessarily binary trees on \mathcal{X} , and let $\mathcal{C} = Cl(\mathcal{T})$. Let $U \subseteq \mathcal{X}$ be a set compatible with \mathcal{C} . Then, $r(\mathcal{C}) = r(\mathcal{C}|U) + r(\mathcal{C}_{U \rightarrow u})$.

Proof. Let $p = r(\mathcal{C})$. We know by Corollary 9 that there exists a maximal ST-set tree sequence (S_1, \dots, S_p) . As usual we define \mathcal{C}_i , $1 \leq i \leq p$, as $\mathcal{C}|S_1 \setminus \dots \setminus S_i$, and we let $\mathcal{C}_0 = \mathcal{C}$. We let $\mathcal{X}_i = \cup_{C \in \mathcal{C}_i} C$ where $\mathcal{X}_0 = \mathcal{X}$. (Note that $\mathcal{X}_i = \mathcal{X} \setminus S_1 \setminus \dots \setminus S_i$). Since $U \setminus S_1 \setminus \dots \setminus S_i = \mathcal{X}_i \cap U$, by repeated application of Observation 6, $\mathcal{X}_i \cap U$ is compatible with \mathcal{C}_i for $0 \leq i \leq p$. Now, recall (see proof of Lemma 10) that for each S_i we can identify a set of two witnesses (where perhaps one of the witnesses is the symbol ρ representing the root).

Here we show, for each S_i , how to hang back a tree representing $\mathcal{C}_{i-1}|S_i$ from a network N representing \mathcal{C}_i to obtain a network N' representing \mathcal{C}_{i-1} where $r(N') = r(N) + 1$ and in N' the taxa $\mathcal{X}_{i-1} \cap U$ are exactly the set of taxa below some cut-edge. The witnesses of S_i guide us how to do this. If we repeat this p times we will obtain a network with p reticulations (and reticulation number p) that represents \mathcal{C} and such that the taxa in U are exactly the subset of taxa below some cut-edge. The theorem will then follow.

The first thing to do is to study the earliest point at which some elements of U are added back into the network. This is an important “base case.” Let us, thus, consider the largest value of i such that $\mathcal{X}_i \cap U \neq \emptyset$. Let i' be equal to this value. Now, suppose $i' = p$. We saw that by repeated application of Observation 6 $\mathcal{X}_p \cap U$ is compatible with \mathcal{C}_p . Furthermore, we know that the




Fig. 11. The three subcases covered by Case 1, which concerns the case when the ST-set S_i that is being hung back, is disjoint from U . The parts of the network containing elements of U are depicted in gray.

clusters C_p can be represented by a tree, so $\mathcal{X}_p \cap U$ is actually an ST-set. Hence, we can assume without loss of generality (by Lemma 11) that the tree that represents C_p , has a cut-edge such that $\mathcal{X}_p \cap U$ is exactly the set of taxa beneath it. Alternatively, suppose $i' < p$. In this case, the first elements of U that are reintroduced into the network are a (not necessarily strict) subset of $S_{i'+1}$, so we have $\mathcal{X}_{i'} \cap U = S_{i'+1} \cap U$. Given that $\mathcal{X}_{i'} \cap U$ is compatible with $C_{i'}$, and $S_{i'+1} \cap U$ is a subset of an ST-set, it follows by Observation 8 that $S_{i'+1} \cap U$ is also an ST-set. We may thus assume without loss of generality that $\mathcal{X}_{i'} \cap U$ is exactly the set of taxa below a cut-edge (again thanks to Lemma 11).

Henceforth, we may assume that the network N that we want to hang (a tree corresponding to) $C_{i-1}|S_i$ back from, contains at least one taxon of U . We will make heavy use of this fact. Let e be the cut-edge of N which the elements of $\mathcal{X}_i \cap U$ are below. Let w_1, w_2 be the two witnesses for S_i , where in some cases $w_2 = \rho$ (representing the root).

There are several cases to consider. The first case is when the tree that we are hanging back, is disjoint from U . See also Fig. 11.

Case 1: $S_i \cap U = \emptyset$.

Subcase 1.1: $\{w_1, w_2\} \cap U = \emptyset$. In this case, we can simply hang back from $\{w_1, w_2\}$ because we are not adding any new elements of U and we are not subdividing any edge reachable from e . Hence, e remains the cut-edge which all present elements of U are below.

Subcase 1.2: $\{w_1, w_2\} \subseteq U$. This case cannot actually occur because it saves a reticulation and hence implies that $r(C) < p$. Indeed, if we simply hang S_i back from $\{w_1, w_2\}$ then we obtain a network that represents C_{i-1} . However, we see from this that every cluster C in C_{i-1}

that is a strict superset of S_i , must contain at least one element of U . S_i is disjoint from U so, since U is compatible with C , C must also contain all other elements of U in the network. Hence, we do not actually need to put S_i below a reticulation at all: we can simply attach it to a single new cut-edge that subdivides e .

Subcase 1.3 (wlog): $w_1 \in U, w_2 \notin U$. Suppose we simply hang back from $\{w_1, w_2\}$, obtaining a network N' that represents C_{i-1} . Note that after doing this any cluster that passes through the reticulation edge starting just above w_1 , must (again because U is compatible with C) contain all elements of U that are in the network. So we can subdivide the cut-edge e and move the tail of that reticulation edge to the newly created vertex.

The next case is when all elements of S_i are in U , see also Fig. 12.

Case 2: $S_i \subseteq U$.

Subcase 2.1: $\{w_1, w_2\} \cap U = \emptyset$. Again, this case cannot actually happen because it saves a reticulation: see subcase 1.2. Indeed, also in this case we don't need a reticulation at all: we can just hang S_i from a single new cut-edge that subdivides e .

Subcase 2.2: $\{w_1, w_2\} \subseteq U$. This case is fine because if we hang back from w_1 and w_2 all the elements of U remain below the cut-edge e .

Subcase 2.3 (wlog): $w_1 \in U, w_2 \notin U$. Suppose we hang back from w_1 and w_2 to obtain a network that represents C_{i-1} . In this case, we could move the reticulation edge that starts just above w_2 (or at the root, in the case that $w_2 = \rho$) to subdivide the cut-edge e .

The final case is where S contains at least one element of U and at least one element not in U . See also Fig. 13.




Fig. 12. The three subcases covered by Case 2, which concerns the case when the ST-set S_i that is being hung back, is a subset of U . The parts of the network containing elements of U are depicted in gray.




Fig. 13. Case 3, which concerns the case when the ST-set S_i that is being hung back, contains elements of U and elements not in U . The parts of the network containing elements of U are depicted in gray. Note that in this case we do not care where the reticulation edges connect to the rest of the network. Here, $R = S_i \setminus U$.

Case 3: $S_i \cap U \neq \emptyset$ and $S_i \setminus U \neq \emptyset$. In this case, we will apply a transformation which brings us back into Case 2. Let $R = S_i \setminus U$. Suppose we hang S_i back from its two witnesses w_1 and w_2 , to obtain a network N' that represents \mathcal{C}_{i-1} . Consider any cluster $C \in \mathcal{C}_{i-1}$ such that $C \cap R \neq \emptyset$. Then, either $C \subseteq R$ or $\mathcal{X}_{i-1} \cap U \subseteq C$. (The second condition holds because, if $C \not\subseteq R$, then it must contain some element in $S_i \cap U$, and hence all elements of U in the network). So a cluster C that is a *strict* subset of S_i is either a subset of R or disjoint from R . Indeed, if C contains one element of R and one of the set U , C must contain all elements of U , a contradiction since C is a strict subset of S_i . It follows that R is compatible with \mathcal{C}_{i-1} so by Observation 8, R is an ST-set. Hence, we may assume (without loss of generality) that R is the taxa set of some subtree T' below a cut-edge in the tree representing $\mathcal{C}_{i-1}|S_i$ that we hung back. Now, we can prune T' and regraft it back onto the network at a new vertex obtained by subdividing e . It can be verified that after this prune/regraft move the resulting network still represents \mathcal{C}_{i-1} . It remains only to apply Case 2, taking w_1, w_2 as the witnesses and $S_i \setminus R$ as the ST-set that we want to hang back. \square

Theorem 6. Let $\mathcal{T} = \{T_1, T_2\}$ be two not necessarily binary trees on \mathcal{X} , and let $\mathcal{C} = Cl(\mathcal{T})$. When given \mathcal{C} as input, $CASS^{DC}$ computes a level- $\ell(\mathcal{C})$ network N with reticulation number $r(\mathcal{C})$ that represents \mathcal{C} .

Proof. Observation 7 and Theorem 5 ensure that we can analyze each connected component of the incompatibility graph $IG(\mathcal{C})$ separately, which (as mentioned) is exactly what $CASS^{DC}$ does. (To see this it is helpful to note that any subset of \mathcal{C} can also, with the possible exception of some superfluous singleton clusters, be expressed as the set of clusters in two trees).

Let K be a connected component and denote by \mathcal{C}_K the set of clusters in K . Let \mathcal{X}_K be the set of taxa equal to the union of all clusters in \mathcal{C}_K . Note that \mathcal{C}_K is not necessarily a tangled set. Indeed, while $IG(\mathcal{C}_K)$ is connected, the second property of a tangled set (every size-2 subset of \mathcal{X}_K is incompatible with \mathcal{C}_K , Section 2) does not always hold. To ensure that the latter condition holds $CASS^{DC}$ simply computes all the maximal ST-sets $\{S_1, \dots, S_k\}$ for \mathcal{C}_K and for each of them replaces all elements of S_i with a single new taxon s_i in \mathcal{C}_K , to obtain a new cluster set \mathcal{C}'_K that is tangled.

To see that the constructed network has minimum level, observe that (by Theorem 4) $CASS$ correctly computes minimum level solutions for the \mathcal{C}'_K cluster sets mentioned above. In [27], it is proven that combining minimum-level solutions for the various \mathcal{C}'_K yields a minimum-level solution for \mathcal{C} .

We now need to prove that the constructed network has reticulation number $r(\mathcal{C})$. Since maximal ST-sets of \mathcal{C}_K are compatible with \mathcal{C}_K and all mutually disjoint (see Corollary 4), it follows from Theorem 5 that

$$r(\mathcal{C}_K) = \sum_{\substack{S \text{ is a maximal} \\ \text{ST-set of } \mathcal{C}_K}} r(\mathcal{C}_K|S) + r(\mathcal{C}'_K).$$

Since $r(\mathcal{C}_K|S)$ always equals zero when S is an ST-set, $r(\mathcal{C}_K) = r(\mathcal{C}'_K)$. Moreover, since the sets \mathcal{X}_K are compatible with \mathcal{C} (from Observation 7) and all mutually disjoint, then from Theorem 5 we have that

$$r(\mathcal{C}) = \sum_{\substack{K \text{ is a connected} \\ \text{component of } IG(\mathcal{C})}} r(\mathcal{C}|X_K) + r(\mathcal{C}'),$$

where \mathcal{C}' is obtained from \mathcal{C} by replacing all elements of \mathcal{X}_K with a single new taxon x_K . Obviously, $r(\mathcal{C}') = 0$. Since $r(\mathcal{C}|X_K) = r(\mathcal{C}_K) = r(\mathcal{C}'_K)$, and $r(\mathcal{C}'_K) = \ell(\mathcal{C}'_K)$ because \mathcal{C}'_K is separating, this concludes the proof that the constructed network has reticulation number $r(\mathcal{C})$. \square

7.4 CASS Can be Used to Compute the Hybridization Number of Two Not Necessarily Binary Trees

It is important to understand the relationship between the results presented in the previous section and the extensive literature on computing the hybridization distance of two trees [14], [15], [16], [17], [29], i.e., the problem of displaying the trees themselves, and not just their clusters.

For two binary trees $\mathcal{T} = \{T_1, T_2\}$ on \mathcal{X} the hybridization distance $h(\mathcal{T})$ is defined as the minimum reticulation number $r_t(\mathcal{T})$ of any network that displays both the trees in \mathcal{T} . In [8], we showed that $r_t(\mathcal{T}) = r(Cl(\mathcal{T}))$. As discussed in [8] that means that both positive and negative results for the computation of $r_t(\mathcal{T})$ transfer automatically to computation of $r(Cl(\mathcal{T}))$ (for two binary trees). Negative results are NP-hardness and APX-hardness; positive results include fixed parameter tractability and running time improvements based on an increasingly deep understanding of maximum acyclic agreement forests.

Before proceeding there are some technical issues regarding the definition of hybridization number of a set \mathcal{T} of two *not necessarily binary* trees on \mathcal{X} . This can be attributed to the fact that several different definitions have appeared in the literature:

1. $h^+(\mathcal{T})$: the minimum reticulation number of N ranging over all networks N that display in a strict topological sense all the trees in \mathcal{T} . This is exactly the definition of *display* given in Section 2. In [8, Fig. 8] this strict definition was used.
2. $h^0(\mathcal{T})$: the minimum reticulation number of N ranging over all networks N that display some not necessarily binary refinement of each tree in \mathcal{T} . Recall that a (binary) *refinement* of a tree T on \mathcal{X} is any (binary) tree T' on \mathcal{X} such that $Cl(T) \subseteq Cl(T')$. (Note that a tree is generically considered to be a refinement of itself).

3. $h^-(\mathcal{T})$: the minimum reticulation number of N ranging over all networks N that display a binary refinement of each tree in \mathcal{T} . This was the definition used in [29].

Two of the definitions, h^0 and h^- , turn out to be equivalent. We clarify this, extend an equivalence result from [8] and thus show that CASS^{DC} correctly computes the hybridization number of two not necessarily binary trees in the sense of h^0 , equivalently h^- . In [8, Fig. 8] a set \mathcal{T} of two trees is given, one of which is nonbinary, such that $r(Cl(\mathcal{T})) < h^+(\mathcal{T})$. A result showing that CASS computes h^+ was thus already excluded.

Observation 9. $h^0(\mathcal{T}) = h^-(\mathcal{T})$ for all sets \mathcal{T} of not necessarily binary trees on the same taxa set \mathcal{X} .

Proof. The fact that $h^0(\mathcal{T}) \leq h^-(\mathcal{T})$ follows immediately from the definitions. Suppose by way of contradiction that there exists a set of not necessarily binary trees \mathcal{T} such that $h^-(\mathcal{T}) > h^0(\mathcal{T})$. Let N be any network with reticulation number $h^0(\mathcal{T})$ that displays some refinement of each tree in \mathcal{T} . Now, it is not too difficult to see (using, for example, the transformation described in [8, Lemma 2]) that we can create a binary network N' such that $r(N') = r(N)$ and such that N' displays a binary refinement of each of the trees in \mathcal{T} , yielding a contradiction. \square

Given a set of trees \mathcal{T} , recall that we define $r_{tr}(\mathcal{T})$ to be the minimum reticulation number of any network that displays $Tr(\mathcal{T})$ (i.e., all the rooted triplets in the input trees). We have the following result:

Theorem 7. Let $\mathcal{T} = \{T_1, T_2\}$ be two not necessarily binary trees on \mathcal{X} . Then, $h^0(\mathcal{T}) = r_{tr}(\mathcal{T})$.

Proof. Obviously, for any set \mathcal{T} of not necessarily binary trees on \mathcal{X} , $r_{tr}(\mathcal{T}) \leq h^0(\mathcal{T})$. It remains to show that this inequality is always tight. Suppose that it is not always tight. Let then $\mathcal{T} = \{T_1, T_2\}$ be two smallest (in terms of the size of $|\mathcal{X}| = n$) trees such that $h^0(\mathcal{T}) > r_{tr}(\mathcal{T})$. Clearly $n > 2$. Now, let N_{tr} be any network that is consistent with $Tr(\mathcal{T}) = R$ such that $r(N_{tr}) = r_{tr}(R)$. If $r_{tr}(R) = 0$ we have a contradiction because this only occurs if T_1 and T_2 have a common refinement, in which case $h^0(\mathcal{T}) = 0$. Hence, $r_{tr}(R) > 0$. This means that N_{tr} has at least one SBR with taxa set S . Now, we claim that, for each $i \in \{1, 2\}$, S is the set of taxa reachable from an edge e_i (in T_i) or as the set of taxa reachable from some subset of the children of some node u_i (in T_i). If this is not so then there exists some triplet $xy|z$ in R such that $x \notin S$ and $y, z \in S$. However, N_{tr} cannot be consistent with $xy|z$ since S is the taxa set of a SBR, which by definition sits below a cut-edge, yielding a contradiction. (If e_i exists assume without loss of generality that u_i is its head). Let $Q' = \{v_1, \dots, v_k\}$ be the set of children of u_i such that for each $v \in Q'$, $\mathcal{X}(T_v)$ contains at least one element of S . Now let T'_i be the tree $T_i \setminus T_{v_1} \dots \setminus T_{v_k}$ and let $\mathcal{T}' = \{T'_1, T'_2\}$. It is clear that $r_{tr}(\mathcal{T}') \leq r_{tr}(R) - 1$. By the assumption of minimality on $|\mathcal{X}|$ we have that $h^0(\mathcal{T}') = r_{tr}(\mathcal{T}')$. Now, let N^* be a network with a minimum reticulation number that displays some refinement of each tree in \mathcal{T}' . Observe that there exists a binary tree T_S on taxa set S such that T_S is a binary

refinement of both $T_1|S$ and $T_2|S$, otherwise $R|S$ would not be an SBR (i.e., it would contain at least one reticulation).

But we can obtain a network with $r(N^*) + 1$ reticulations that displays some refinement of T_1 and T_2 by hanging T_S back below a single new reticulation in N^* such that the tails of the two reticulation edges extend the embeddings of the refinements of T'_1 and T'_2 in N^* . Hence, $h^0(\mathcal{T}) \leq r_t(\mathcal{T}') + 1 \leq r_t(\mathcal{T}) < h^0(\mathcal{T})$, contradiction. \square

The following lemma extends a result from [8]:

Lemma 12. If \mathcal{T} consists of two not necessarily binary phylogenetic trees on the same set of taxa, $r_{tr}(\mathcal{T}) = h^0(\mathcal{T}) = r(Cl(\mathcal{T}))$.

Proof. In [8], it is proven that for a not necessarily binary set of trees \mathcal{T} on the same set of taxa, $r_{tr}(\mathcal{T}) \leq r(Cl(\mathcal{T}))$. Now, if a network displays some refinement of a not necessarily binary tree T , then it represents all the clusters in $Cl(T)$ (and possibly more). Hence, we also have that $r(Cl(\mathcal{T})) \leq h^0(\mathcal{T})$. Combining this with Theorem 7 gives the result. \square

Combining all these results we finally obtain the following theorem.

Theorem 8. Let $\mathcal{T} = \{T_1, T_2\}$ be two not necessarily binary trees on \mathcal{X} , and let $\mathcal{C} = Cl(\mathcal{T})$. When given \mathcal{C} as input, CASS^{DC} computes a network N such that $r(N) = h^0(\mathcal{T}) = h^-(\mathcal{T})$.

8 CONCLUSION AND OPEN PROBLEMS

The largest open problem emerging from this paper is whether there exists a “reasonable” polynomial-time algorithm for constructing phylogenetic networks of bounded reticulation number (or level) that represent a given set of clusters. The result in Section 3, which shows a theoretical polynomial-time algorithm for constructing networks of bounded level, does not lend itself to a real-world implementation. On the other hand, we have seen that for clusters obtained from *binary* trees a relatively simple and efficient algorithm can be used. At the moment, however, there is no reasonable polynomial-time algorithm for general cluster sets, i.e., those obtained from sets of potentially nonbinary trees. We have shown that CASS, which has a reasonable running time, is in general not optimal (although we had to explicitly engineer a highly synthetic counter example to determine this). CASS is, however, an *extremely* greedy algorithm, in the sense that at every iteration it assumes that *all* maximal ST-sets are below cut edges. Can we relax this assumption in some way to yield a slightly less greedy version of CASS that is optimal? A less urgent problem, but nevertheless very interesting, is the question whether CASS is optimal for clusters obtained from exactly three binary trees on \mathcal{X} .

ACKNOWLEDGMENTS

S. Kelk was partly funded by a Computational Life Sciences grant of The Netherlands Organisation for Scientific Research (NWO). L. van Iersel was funded by the Allan Wilson Centre for Molecular Ecology and Evolution.

REFERENCES

- [1] C. Semple and M. Steel, *Phylogenetics*. Oxford Univ. Press, 2003.
- [2] *Mathematics of Evolution and Phylogeny*, O. Gascuel, ed. Oxford Univ. Press, Inc., <http://portal.acm.org/citation.cfm?id=1557209>, 2005.
- [3] *Reconstructing Evolution: New Mathematical and Computational Advances*, O. Gascuel and M. Steel, eds. Oxford Univ. Press, <http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0199208220>, 2007.
- [4] D.H. Huson, R. Rupp, and C. Scornavacca, *Phylogenetic Networks: Concepts, Algorithms and Applications*. Cambridge Univ. Press, 2011.
- [5] L. Nakhleh, "Evolutionary Phylogenetic Networks: Models and Issues," *The Problem Solving Handbook for Computational Biology and Bioinformatics*, Springer, 2009.
- [6] C. Semple, "Hybridization Networks," *Reconstructing Evolution - New Mathematical and Computational Advances*, Oxford Univ. Press, 2007.
- [7] D.H. Huson, R. Rupp, V. Berry, P. Gambette, and C. Paul, "Computing Galled Networks from Real Data," *Bioinformatics*, vol. 25, no. 12, pp. i85-i93, 2009.
- [8] L.J.J. van Iersel and S.M. Kelk, "When Two Trees Go to War," *J. Theoretical Biology*, vol. 269, no. 1, pp. 245-255, 2011.
- [9] J. Jansson and W.-K. Sung, "Inferring a Level-1 Phylogenetic Network from a Dense Set of Rooted Triplets," *Theoretical Computer Science*, vol. 363, no. 1, pp. 60-68, 2006.
- [10] J. Jansson, N.B. Nguyen, and W.-K. Sung, "Algorithms for Combining Rooted Triplets into a Galled Phylogenetic Network," *SIAM J. Computing*, vol. 35, no. 5, pp. 1098-1121, 2006.
- [11] L.J.J. van Iersel, J.C.M. Keijsper, S.M. Kelk, L. Stougie, F. Hagen, and T. Boekhout, "Constructing Level-2 Phylogenetic Networks from Triplets," *IEEE/ACM Trans. Computational Biology and Bioinformatics*, vol. 6, no. 4, pp. 667-681, Oct.-Dec. 2009.
- [12] L.J.J. van Iersel, S.M. Kelk, and M. Mnich, "Uniqueness, Intractability and Exact Algorithms: Reflections on Level- k Phylogenetic Networks," *J. Bioinformatics and Computational Biology*, vol. 7, no. 2, pp. 597-623, 2009.
- [13] T.-H. To and M. Habib, "Level- k Phylogenetic Networks Are Constructable from a Dense Triplet Set in Polynomial Time," *Proc. Conf. Combinatorial Pattern Matching (CPM '09)*, pp. 275-288, 2009.
- [14] M. Bordewich and C. Semple, "Computing the Minimum Number of Hybridization Events for a Consistent Evolutionary History," *Discrete Applied Math.*, vol. 155, no. 8, pp. 914-928, 2007.
- [15] M. Bordewich, S. Linz, K.S. John, and C. Semple, "A Reduction Algorithm for Computing the Hybridization Number of Two Trees," *Evolutionary Bioinformatics*, vol. 3, pp. 86-98, 2007.
- [16] M. Bordewich and C. Semple, "Computing the Hybridization Number of Two Phylogenetic Trees Is Fixed-Parameter Tractable," *IEEE/ACM Trans. Computational Biology and Bioinformatics*, vol. 4, no. 3, pp. 458-466, July-Sept. 2007.
- [17] J. Collins, S. Linz, and C. Semple, "Quantifying Hybridization in Realistic Time," *J. Computational Biology*, vol. 18, no. 10, pp. 1305-1318, 2011.
- [18] Y. Wu and W. Jiayin, "Fast Computation of the Exact Hybridization Number of Two Phylogenetic Trees," *Proc. Sixth Int'l Symp. Bioinformatics Research and Applications (ISBRA)*, vol. 6053, pp. 203-214, 2010.
- [19] Y. Wu, "Close Lower and Upper Bounds for the Minimum Reticulate Network of Multiple Phylogenetic Trees," *Bioinformatics*, vol. 26, pp. i140-i148, 2010.
- [20] T. Huynh, J. Jansson, N. Nguyen, and W.-K. Sung, "Constructing a Smallest Refining Galled Phylogenetic Network," *Proc. Int'l Conf. Research in Computational Molecular Biology (RECOMB)*, pp. 265-280, 2005.
- [21] L. van Iersel and S. Kelk, "Constructing the Simplest Possible Phylogenetic Network from Triplets," *Algorithmica*, pp. 1-29, Springer, <http://dx.doi.org/10.1007/s00453-009-9333-0>, 2009.
- [22] D. Gusfield, V. Bansal, V. Bafna, and Y. Song, "A Decomposition Theory for Phylogenetic Networks and Incompatible Characters," *J. Computational Biology*, vol. 14, no. 10, pp. 1247-1272, 2007.
- [23] D. Gusfield, D. Hickerson, and S. Eddhu, "An Efficiently Computed Lower Bound on the Number of Recombinations in Phylogenetic Networks: Theory and Empirical Study," *Discrete Applied Math.*, vol. 155, nos. 6/7, pp. 806-830, 2007.
- [24] Y. Wu and D. Gusfield, "A New Recombination Lower Bound and the Minimum Perfect Phylogenetic Forest Problem," *J. Combinatorial Optimization*, vol. 16, no. 3, pp. 229-247, 2008.
- [25] S.R. Myers and R.C. Griffiths, "Bounds on the Minimum Number of Recombination Events in a Sample History," *Genetics*, vol. 163, pp. 375-394, 2003.
- [26] D.H. Huson and T.H. Klöpper, "Beyond Galled Trees—Decomposition and Computation of Galled Networks," *Proc. Int'l Conf. Research in Computational Molecular Biology (RECOMB)*, pp. 211-225, 2007.
- [27] L.J.J. van Iersel, S.M. Kelk, R. Rupp, and D.H. Huson, "Phylogenetic Networks Do Not Need to be Complex: Using Fewer Reticulations to Represent Conflicting Clusters," *Bioinformatics*, vol. 26, pp. i124-i131, 2010.
- [28] D. Huson and C. Scornavacca, "Dendroscope 3—A Program for Computing and Drawing Rooted Phylogenetic Trees and Networks," in Preparation. Software Available from: www.dendroscope.org, 2011.
- [29] S. Linz and C. Semple, "Hybridization in Non-Binary Trees," *IEEE/ACM Trans. Computational Biology and Bioinformatics*, vol. 6, no. 1, pp. 30-45, Jan.-Mar. 2009.
- [30] P. Gambette, V. Berry, and C. Paul, "The Structure of Level- k Phylogenetic Networks," *Proc. 20th Ann. Symp. Combinatorial Pattern Matching*, pp. 289-300, http://dx.doi.org/10.1007/978-3-642-02441-2_26, 2009.
- [31] P. Gambette, "Generators - Building Level- k Generators," <http://www.lirmm.fr/gambette/ProgGenerators.php>, 2011.
- [32] V. Bafna and V. Bansal, "Inference about Recombination from Haplotype Data: Lower Bounds and Recombination Hotspots," *J. Computational Biology*, vol. 13, pp. 501-521, 2006.



Steven M. Kelk is currently an assistant professor in the Department of Knowledge Engineering (DKE) at Maastricht University in the Netherlands, where he works on applications of combinatorial optimization and discrete mathematics to computational biology. His main research interests are approximation algorithms, phylogenetic networks, metabolic networks and exploring the interface between discrete, and continuous mathematical biology.



Celine Scornavacca received the PhD degree in computer science from Montpellier University in 2009. After a postdoctoral position at the University of Tübingen, she joined the Institut des Sciences de l'Évolution de Montpellier (ISEM) as research associate in October 2011. Her research interests include parameterized complexity, supertree, and network methods for phylogenetics, combinatorics, and bioinformatics.



Leo van Iersel received the MSc degree in 2004 at Twente University and the PhD degree in 2009 at Eindhoven University, both in The Netherlands. Subsequently, he worked as a postdoctoral researcher at the University of Canterbury in Christchurch, New Zealand, until mid 2010. After one year teaching in East Africa, he will start another postdoctoral research position at CWI in Amsterdam, The Netherlands. His research is concerned with the construction of phylogenetic trees and networks from biological data.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.