

UC Davis

UC Davis Previously Published Works

Title

Analyzing and reconstructing reticulation networks under timing constraints

Permalink

<https://escholarship.org/uc/item/1q65240w>

Journal

Journal of Mathematical Biology, 61(5)

ISSN

1432-1416

Authors

Linz, Simone
Semple, Charles
Stadler, Tanja

Publication Date

2010-11-01

DOI

10.1007/s00285-009-0319-y

Peer reviewed

Analyzing and reconstructing reticulation networks under timing constraints

Simone Linz · Charles Semple · Tanja Stadler

Received: 13 May 2009 / Revised: 19 November 2009 / Published online: 27 December 2009
© The Author(s) 2009. This article is published with open access at Springerlink.com

Abstract Reticulation networks are now frequently used to model the history of life for various groups of species whose evolutionary past is likely to include reticulation events such as horizontal gene transfer or hybridization. However, the reconstructed networks are rarely guaranteed to be temporal. If a reticulation network is temporal, then it satisfies the two biologically motivated timing constraints of instantaneously occurring reticulation events and successively occurring speciation events. On the other hand, if a reticulation network is not temporal, it is always possible to make it temporal by adding a number of additional unsampled or extinct taxa. In the first half of the paper, we show that deciding whether a given number of additional taxa is sufficient to transform a non-temporal reticulation network into a temporal one is an NP-complete problem. As one is often given a set of gene trees instead of a network in the context of hybridization, this motivates the second half of the paper which provides an algorithm, called TEMPORALHYBRID, for reconstructing a temporal hybridization network that simultaneously explains the ancestral history of two trees or indicates that no such network exists. We further derive two methods to decide whether or not a temporal hybridization network exists for two given trees and illustrate one of the methods on a grass data set.

All authors contributed equally to this work.

S. Linz

Department of Computer Science, University of California, Davis, CA, USA
e-mail: linzs@cs.ucdavis.edu

C. Semple

Biomathematics Research Centre, Department of Mathematics and Statistics,
University of Canterbury, Christchurch, New Zealand
e-mail: c.semple@math.canterbury.ac.nz

T. Stadler (✉)

Institute of Integrative Biology, ETH Zürich, Zürich, Switzerland
e-mail: tanja.stadler@env.ethz.ch

Mathematics Subject Classification (2000) 05C05 · 05C20 · 92D15


1 Introduction

Evolution is often regarded as a tree-like process in which an ancestral species evolves to a set of present-day species via a sequence of speciation events. This approach is well-suited to tackle various questions arising from evolutionary studies. However, reticulation has now been shown to be an important process in the evolution of various ancestral and present-day species (for example, see [Mallet 2005](#); [Ochman et al. 2000](#) and references therein). Two major reticulation scenarios that are discussed in this paper are horizontal gene transfer (HGT) and hybridization. In the case of the latter, two organisms of different ancestral species combine their DNA to create a new species. This process is common in certain groups of plants and fish ([Mallet 2005](#)). On the other hand, in the case of HGT, which is widely observed among bacteria ([Ochman et al. 2000](#)), a piece of DNA (e.g. a gene) is transferred from one organism to another which is not its offspring. Consequently, if the genome of a species is chimeric as a result of one or more HGT or hybridization events, its evolutionary history can often be better represented by using a reticulation network rather than a phylogenetic tree.

Recently, a lot of effort has been put into the development of algorithms to reconstruct reticulation networks for a set of present-day species (for example, see [Makarenkov et al. 2006](#) and references therein). However, as pointed out in [Maddison \(1997\)](#), although a reticulation network might explain several conflicting signals in a data set, there may be no process of instantaneously occurring reticulation events that realizes this network. Consequently, the two (resp. three) species that are involved in an HGT (resp. a hybridization) event are not guaranteed to exist contemporaneously. Roughly speaking, we say that a reticulation network is *temporal* if each reticulation event can be realized between coexisting ancestral species while speciation events occur successively. Beside the reconstruction of possibly non-temporal reticulation networks, there exists a number of algorithms that do calculations on networks and implicitly assume that the input consists of a temporal reticulation network. For example, [Jin et al. \(2007\)](#) have developed an algorithm for computing the parsimony score of a temporal reticulation network.

A reticulation network that is not temporal does not necessarily imply that the network is incorrect. By allowing for additional taxa that for instance correspond to unsampled or extinct taxa one can always transform a non-temporal reticulation network into a temporal one without introducing any new reticulation events ([Baroni et al. 2006](#); [Moret et al. 2004](#)). For example, consider the reticulation network shown in [Fig. 1](#), where the non-arrowed arcs are directed down the page. If one ignores the dashed arc and its end vertices, then the network is not temporal. However, by allowing for this dashed arc and the taxon x , the resulting network is temporal. Given this, a natural task in the study of biologically meaningful reticulation networks is to calculate the minimum number of additional taxa that must be allowed for, so that the resulting network is temporal. We call the analogous decision problem ADDTAXA and show in the first half of this paper that this problem is NP-complete.

Fig. 1 A temporal hybridization network with one additional taxa x



In the second half of this paper, we focus our attention on hybridization and consider a task that is one step closer to the initial biological data. Instead of being given a reticulation network, one frequently starts with a set of gene trees. For example, due to hybridization, these gene trees—reconstructed for different genetic loci—often reveal inconsistencies. Here, a fundamental problem is to calculate the smallest number of hybridization events needed to simultaneously explain the set of gene trees. While this problem is NP-hard, [Bordewich et al. \(2007\)](#) and, more recently, [Collins et al. \(2009\)](#) have implemented a fixed-parameter algorithm for solving it when the initial set consists of two gene trees, \mathcal{T} and \mathcal{T}' say. This algorithm is dependent on finding an associated optimal agreement forest. For the purposes of the introduction, simply think of the forest as a smallest collection of (disjoint) subtrees common to \mathcal{T} and \mathcal{T}' . From this forest, the algorithm HYBRIDPHYLOGENY ([Baroni et al. 2006](#)) can be applied to reconstruct a hybridization network that explains \mathcal{T} and \mathcal{T}' , and in which the number of hybridization events equates to the size of the forest. However, despite its practical application, HYBRIDPHYLOGENY does not guarantee that the resulting network is temporal. In the second half of the paper, we provide an algorithm, called TEMPORALHYBRID, that constructs temporal hybridization networks from agreement forests. It is worth noting here that there is no guarantee that such networks exist. Furthermore, two applications of TEMPORALHYBRID are given with the second application being used to analyze a grass data set.

The paper is organized as follows. The next section contains some notation and terminology that is used throughout the paper and formally defines the decision problem ADDTAXA. In Sect. 3, we show that ADDTAXA is NP-complete. As fixed-parameter algorithms have been shown to often be a valuable tool for attacking instances of an NP-complete problem that are otherwise not solvable in reasonable time, we end Sect. 3 by showing how a fixed-parameter algorithm for another NP-complete problem can be used to solve ADDTAXA. In the first part of Sect. 4, we present the algorithm TEMPORALHYBRID. Two methods to decide whether or not a temporal hybridization network for two trees exists are presented in the end of Sects. 4.2 and 4.3. The second method is illustrated on a grass data set. Section 5 summarizes the paper.

2 Modeling reticulate evolution

In this section, we give some preliminary definitions for acyclic digraphs that are commonly used to model reticulate evolution, and formally state the decision problem ADDTAXA.

Let X be an arbitrary finite set. A *reticulation network* \mathcal{N} on X is a rooted acyclic digraph with the following properties:

- (i) the *root* has indegree 0 and outdegree 2;
- (ii) X is the set of *leaves* of the network, i.e. vertices with outdegree 0 and indegree 1;
- (iii) all remaining vertices are *interior vertices*, and each such vertex either has indegree 1 and outdegree 2 or has indegree 2 and outdegree 1;
- (iv) either one arc or both arcs ending in a vertex with indegree 2 are *reticulation arcs*, all other arcs in the network are *tree arcs*; and
- (v) every interior vertex has at least one outgoing tree arc.

The set X represents in our context a collection of present-day taxa. Furthermore, vertices of \mathcal{N} with indegree 1 are referred to as *tree vertices* while vertices of \mathcal{N} with indegree 2 are referred to as *reticulation vertices*. In this paper, a reticulation vertex represents either a hybrid species or a species which evolved due to HGT. Property (v) in the definition of a reticulation network guarantees that every species that arises from either a speciation or a reticulation event exists for a certain time before possibly going extinct. For example, an ancestral species does not yield two new hybrid species and simultaneously becomes extinct. This is biologically well-motivated since, although hybridization can result in the extinction of one or both hybridizing species, this process takes at least a few generations, and hybridization and extinction are often only locally observed (Martinsen et al. 2001; Wolf et al. 2001). Ignoring the dashed arc and its end vertices for the moment, Fig. 1 shows a reticulation network, where $X = \{a, b, c, d\}$. Here, as in all figures in the paper, the direction of any non-arrowed arc is down the page.

Let \mathcal{N} be a reticulation network with vertex set V and arc set A , and let $u, v \in V$. If there is a directed path from u to v and $u \neq v$, we write $u < v$ and refer to u as an *ancestor* of v and to v as a *descendant* of u . If (u, v) is an arc of \mathcal{N} , we call (u, v) a *parent arc* of v , and say that u is a *parent* of v and that v is a *child* of u .

A reticulation vertex due to hybridization is called a *hybridization vertex* and the parent arcs of such a vertex are called *hybridization arcs*. Similarly, a reticulation vertex due to HGT is called an *HGT vertex* and exactly one of its two parent arcs is an *HGT arc* with this arc indicating the direction of the DNA transfer. Hybridization and HGT arcs are collectively referred to as *reticulation arcs* (see definition of a reticulation network). A *hybridization network* is a reticulation network where all reticulation events are due to hybridization while an *HGT network* is a reticulation network where all reticulation events are due to HGT.

As only one parent arc of an HGT vertex is an HGT arc, a reticulation network can lead to different HGT networks depending on which parent arc is the HGT arc. In contrast, as both parent arcs of a hybridization vertex are hybridization arcs, a reticulation network leads to a unique hybridization network. The discrimination between HGT and hybridization arcs becomes crucial when considering *temporal* reticulation networks whose formal definition is given in the next paragraph. Throughout the paper, whenever we refer to a reticulation network we mean that the network is either a hybridization network or an HGT network. Furthermore, in the context of figures, reticulation arcs are always drawn with an arrow while tree arcs are drawn without an arrow. Again ignoring the dashed arc and its end vertices, Fig. 1 and the left figure





Fig. 2 A non-temporal HGT network (*left*) and its associated critical graph (*right*)

Fig. 3 A temporal hybridization network for the 4 taxa $a, b, c,$ and d with a temporal labeling



in Fig. 2 show a hybridization network and an HGT network, respectively. Lastly, we remark that our definition of a hybridization network \mathcal{N} coincides with that of a so-called tree-child phylogenetic network, which has been introduced by Cardona et al. (2009). But note that an HGT network is not necessarily a tree-child network as the two children of a tree vertex can both be HGT vertices, see Fig. 2.

We next formalize the notion of assigning dates to the vertices of a reticulation network. Let \mathcal{N} be a reticulation network with vertex set V and arc set A , and let V' be a subset of V . Let $f : V' \rightarrow \mathbb{R}^+$ be a map such that, for all $s, t \in V'$, we have $f(s) = f(t)$ whenever (s, t) is a reticulation arc, and $f(s) < f(t)$ whenever there is a directed path from s to t that contains a tree arc. Then f is a *partial temporal labeling* of \mathcal{N} . If $V' = V$, then f is a *temporal labeling* of \mathcal{N} , and we say that \mathcal{N} is *temporal* or has a *temporal representation*. As an example, consider the hybridization network shown in Fig. 3 which illustrates a temporal labeling f of this network.

Not all reticulation networks have a temporal representation. However, as noted in Baroni et al. (2006), one can always make such a network temporal by allowing for additional taxa that do not introduce new reticulation vertices. Such taxa may correspond to unsampled or extinct taxa. This can be done as follows. Let \mathcal{N} be a reticulation network, and let $e = (u, v)$ be an arc of \mathcal{N} . Consider the reticulation network obtained from \mathcal{N} by replacing e with a 2-arc directed path consisting of (u, z) and (z, v) , and then adjoining a new taxa x via the new arc (z, x) . We say that the resulting reticulation network has been obtained from \mathcal{N} by *adding a new taxa x across (u, v)* . As an example, consider Fig. 1. The hybridization network shown (including the dashed arc) has been obtained from the underlying solid-arc hybridization network by adding x across (u, v) . We will soon see that by adding new taxa in this way to \mathcal{N} it is always

possible to produce a reticulation network that has a temporal representation. This motivates the following decision problem which is the main focus of the first part of this paper.

Problem: $\text{ADDTAXA}(\mathcal{N}, k)$

Instance: A reticulation network \mathcal{N} and a positive integer k .

Question: Is there a temporal reticulation network that can be obtained from \mathcal{N} by adding at most k new taxa?

If an instance of $\text{ADDTAXA}(\mathcal{N}, k)$ is restricted to HGT or hybridization networks, we refer to the resulting decision problems as $\text{ADDTAXA}_{\text{HGT}}(\mathcal{N}, k)$ and $\text{ADDTAXA}_{\text{Hybrid}}(\mathcal{N}, k)$, respectively.

3 ADDTAXA is NP-complete

In this section, we show that both $\text{ADDTAXA}_{\text{HGT}}(\mathcal{N}, k)$ and $\text{ADDTAXA}_{\text{Hybrid}}(\mathcal{N}, k)$ are NP-complete. We begin with some further definitions and establish several preliminary results. Let \mathcal{N} be a reticulation network, and let (v, v_1) be a reticulation arc. If there exists a reticulation vertex w with $w \neq v_1$ such that $v < w$, but $v_1 \not< w$, then v is said to be *critical*, in which case, v_1 is a *critical reticulation vertex*. Furthermore, we call v_1 a *critical HGT vertex* or a *critical hybridization vertex* if \mathcal{N} is an HGT or a hybridization network, respectively. To illustrate, consider Fig. 2. In the left-hand figure, u is a critical vertex as the reticulation vertex v_1 is a descendant of u , but is not a descendant of u_1 .

For a reticulation network \mathcal{N} , let $\mathcal{C}_{\mathcal{N}}$ denote the graph whose vertex set is the set of critical vertices of \mathcal{N} and whose arc set is

$$\{(u, v) : u < v_1, \text{ where } (v, v_1) \text{ is a reticulation arc in } \mathcal{N}\}.$$

The graph $\mathcal{C}_{\mathcal{N}}$ is called the *critical graph* of \mathcal{N} . As an example, the critical graph of the HGT network shown in Fig. 2 is shown in the right of that figure.

Lemma 3.1 *Let \mathcal{N} be a reticulation network. Then \mathcal{N} has a temporal representation if and only if there exists a partial temporal labeling for the set consisting of the critical vertices and the critical reticulation vertices of \mathcal{N} .*

Proof Evidently, if \mathcal{N} has a temporal representation, then, by restricting the labeling of such a representation to the critical vertices and the critical reticulation vertices, we have a partial temporal labeling of these vertices. Next we prove the converse for hybridization networks and then use this fact to prove the converse for HGT networks.

Let \mathcal{N} be a hybridization network. Suppose that we have a partial temporal labeling f_c of the critical vertices and critical hybridization vertices of \mathcal{N} . By assigning values to the other vertices of \mathcal{N} , we extend the labeling of these vertices under f_c to a temporal labeling of \mathcal{N} . We do this in two steps; we first extend to all hybridization vertices and their parents, and then, second, extend to all remaining tree vertices.

For a hybridization arc (v, v_1) in which v is critical, it is possible that the other parent of v_1 , say v' , is not critical. In this case, extend f_c by assigning v' the same value as v and v_1 . Since v' is not critical, every hybridization vertex that is a descendant of v' is also a descendant of v_1 . It is now easily seen that this extension is a partial temporal labeling of \mathcal{N} . Continuing in this way, we eventually assign appropriate values to all parents of critical hybridization vertices.

Together with the assignment given in the last paragraph, we now extend the partial temporal labeling f_c of a hybridization network to all remaining hybridization vertices and their respective parents so that the resulting extension is a partial temporal labeling of \mathcal{N} . We do this using induction on the number k of hybridization vertices not currently assigned a temporal label. If $k = 0$, then this extension is vacuous. Now suppose that we can extend f_c to include an additional $k - 1$ hybridization vertices and their respective parents, where $k \geq 1$. Let v_1 be a hybridization vertex without a label, and suppose that its parents are v and v' . Let t_a denote the maximum value of a temporal label assigned to an ancestor of v_1 . If no ancestor of v_1 has been assigned such a label, set $t_a = 0$. Let t_d denote the minimum value of a temporal label assigned to a descendant of v_1 . If no descendant of v_1 has been assigned such a label, set $t_d = \infty$. Note that, since v_1 is unlabeled, neither v nor v' is critical, and so every hybridization vertex that is a descendant of v or v' is a descendant of v_1 . Therefore, t_d is also the minimum value of a temporal label assigned to a descendant of either v or v' . Now $t_a < t_d$; otherwise we contradict the induction assumption that we can extend f_c to $k - 1$ additional hybridization vertices and their respective parents. It now follows that any real number in the interval (t_a, t_d) can be assigned to v_1 , v , and v' to obtain an appropriate extension of f_c to k hybridization vertices and their parents. Thus, by induction, f_c can be extended to all such hybridization vertices and their parents.

Now that all hybridization vertices and their respective parents are labeled appropriately, we next label the remaining tree vertices of \mathcal{N} . We first partition the set of tree vertices (including the root) of \mathcal{N} as follows. Let C_1, C_2, \dots, C_k denote the maximal connected subgraphs of \mathcal{N} whose vertex sets consist entirely of tree vertices. For all i , the subgraph C_i is a rooted tree as each vertex is a tree vertex. Without loss of generality, we may assume that C_1, C_2, \dots, C_k is ordered so that $i < j$ if and only if either the root of C_i is an ancestor of the root of C_j or neither the root of C_i nor the root of C_j is an ancestor of the other. Thus C_1 contains the root of \mathcal{N} . Beginning with the vertices in C_1 and using this ordering, we can systematically label the remaining unlabeled tree vertices of \mathcal{N} as follows. For each i , label the root with a real number bigger than that of any of its ancestors, but smaller than that of any descendant that is labeled and, label each of its leaves with a real number bigger than that of any currently labeled ancestor, but smaller than that of any descendant that is labeled. Now label the rest of the vertices of C_i appropriately. Note that C_i may contain a vertex v that has already been assigned a label; in which case v is a parent of a hybridization vertex. However, this is not problematic as it simply means that each ancestral vertex of v must be labeled with a real number smaller than the label assigned to v , and each descendant vertex of v must be labeled with a real number bigger than the label assigned to v . As C_i is a rooted tree, this is always possible. The resulting labeling is a temporal labeling of \mathcal{N} .

Now, let \mathcal{N} be an HGT network and suppose that we have a partial temporal labeling f_c of the critical vertices and the critical HGT vertices of \mathcal{N} . Let \mathcal{N}' be the hybridization network that is obtained from \mathcal{N} by adding a new taxa across the parent tree arc of each HGT vertex and viewing each HGT vertex as a hybridization vertex. By construction, a vertex is critical in \mathcal{N}' if and only if it is critical in \mathcal{N} . Thus f_c is also a partial temporal labeling of the critical vertices and critical hybridization vertices of \mathcal{N}' . By assigning values to the other vertices of \mathcal{N}' as described in the proof for hybridization networks, we can extend f_c to a temporal labeling f' of \mathcal{N}' . Restricting f' to the vertices of \mathcal{N} , it is easily checked that this restriction gives a temporal labeling of \mathcal{N} . This completes the proof of the lemma.

Remark We remarked prior to the formal description of ADDTAXA that a reticulation network \mathcal{N} can always be made temporal by adding new taxa in a certain way. Indeed, because of Lemma 3.1, we can do this as follows. Suppose that (v, v_1) is a reticulation arc in \mathcal{N} such that v is critical, and consider the reticulation network \mathcal{N}' obtained from \mathcal{N} by adjoining a new taxa x across (v, v_1) . Let (z, x) denote the adjoining arc. Since (v, z) is not a reticulation arc, v is not critical in \mathcal{N}' . Furthermore, z is not critical in \mathcal{N}' . So \mathcal{N}' has strictly less critical vertices than \mathcal{N} . Continuing in this way, we eventually obtain a reticulation network with no critical vertices and thus, by Lemma 3.1, a network that is temporal.

We next restrict our attention to HGT networks. The reason for this is that we first prove that $\text{ADDTAXA}_{\text{HGT}}(\mathcal{N}, k)$ is NP-complete and then use this result to show that $\text{ADDTAXA}_{\text{Hybrid}}(\mathcal{N}, k)$ is NP-complete.

Lemma 3.2 *Let \mathcal{N} be an HGT network, and let (v, v_1) be an HGT arc of \mathcal{N} such that v is critical. Let \mathcal{N}' be the HGT network obtained from \mathcal{N} by adding a new taxa across (v, v_1) . Then the graphs $\mathcal{C}_{\mathcal{N}'}$ and $\mathcal{C}_{\mathcal{N}' \setminus v}$ are equal.*

Proof First observe that, as v is not incident with an HGT arc in \mathcal{N}' , it is not critical. Moreover, the parent vertex, z say, of the new taxa in \mathcal{N}' is also not critical since, except for the new taxa, z and v_1 have the same descendants. It now follows that the vertex sets of $\mathcal{C}_{\mathcal{N}'}$ and $\mathcal{C}_{\mathcal{N}' \setminus v}$ are equal. Furthermore, for all vertices u and w in $\mathcal{C}_{\mathcal{N}'}$, we have that (u, w) is an arc in $\mathcal{C}_{\mathcal{N}'}$ if and only if it is an arc in $\mathcal{C}_{\mathcal{N}' \setminus v}$. Thus the graphs $\mathcal{C}_{\mathcal{N}'}$ and $\mathcal{C}_{\mathcal{N}' \setminus v}$ are equal.

Proposition 3.3 *Let \mathcal{N} be an HGT network. Then \mathcal{N} has a temporal representation if and only if $\mathcal{C}_{\mathcal{N}}$ is acyclic.*

Proof First suppose that $\mathcal{C}_{\mathcal{N}}$ is acyclic. To show that \mathcal{N} has a temporal representation, it suffices to show by Lemma 3.1 that there is a partial temporal labeling f_c of the critical vertices and critical HGT vertices of \mathcal{N} . We define such a labeling as follows. It is well-known and easily proved that, as $\mathcal{C}_{\mathcal{N}}$ is acyclic, it has a vertex v with indegree zero. Let (v, v_1) be the HGT arc incident with v in \mathcal{N} . Since v has indegree zero, no vertex in $\mathcal{C}_{\mathcal{N}}$ is an ancestor of v or v_1 in \mathcal{N} . Set $f_c(v) = f_c(v_1) = 1$. Now delete v in $\mathcal{C}_{\mathcal{N}}$ and consider the resulting graph. Since this graph is acyclic, it has a vertex of indegree zero. Repeat the above process for this vertex, but assign it and its child HGT vertex value 2 under f_c . By deleting this vertex and continuing in this way, we

eventually assign all critical vertices and critical HGT vertices of \mathcal{N} a value under f_c . Moreover, f_c is a partial temporal labeling of the critical vertices and critical HGT vertices of \mathcal{N} and so, by Lemma 3.1, \mathcal{N} has a temporal representation.

It remains to show that \mathcal{N} having a temporal representation implies $\mathcal{C}_{\mathcal{N}}$ being acyclic. We establish this statement by showing that if $\mathcal{C}_{\mathcal{N}}$ has a directed cycle, then \mathcal{N} does not have a temporal representation. Suppose that $\mathcal{C}_{\mathcal{N}}$ has a directed cycle $v_0, v_1, \dots, v_{m-1}, v_0$. For all i , let v_{i1} denote the child HGT vertex of v_i in \mathcal{N} . Then, for all i modulo m , there is a directed path from v_{i-1} to v_{i1} containing at least one tree arc. It follows that \mathcal{N} has no temporal representation. This completes the proof of the proposition.

Corollary 3.4 *Let \mathcal{N} be an HGT network, and let V_c be a subset of the vertex set of $\mathcal{C}_{\mathcal{N}}$. Then the HGT network obtained from \mathcal{N} by adding, for each $v \in V_c$, a new taxa across the HGT arc incident with v has a temporal representation if and only if $\mathcal{C}_{\mathcal{N}} \setminus V_c$ is acyclic.*

Proof Combining Lemma 3.2 and Proposition 3.3 gives the desired result.

We now show that $\text{ADDTAXA}_{\text{HGT}}(\mathcal{N}, k)$ is NP-complete. The NP-complete problem that we use for the polynomial-time reduction is FEEDBACKVERTEXSET (Karp 1972):

Problem: $\text{FEEDBACKVERTEXSET}(G, m)$

Instance: Directed graph $G = (V, E)$ and a positive integer $m \leq |V|$.

Question: Is there a subset $V' \subseteq V$ with $|V'| \leq m$ such that $G \setminus V'$ is acyclic?

Remarks 1. Observe that if \mathcal{N} is an HGT network, then, by Corollary 3.4, the answer to $\text{ADDTAXA}_{\text{HGT}}(\mathcal{N}, k)$ is *yes* if and only if the answer to $\text{FEEDBACKVERTEXSET}(\mathcal{C}_{\mathcal{N}}, k)$ is *yes*.

2. Baroni et al. (2006) have provided the algorithm TEMPREP which determines whether a given hybridization network \mathcal{N} has a temporal representation or not. To describe the idea of TEMPREP , let V be the vertex set of \mathcal{N} . Ignoring the direction of the arcs of \mathcal{N} , an equivalence relation on V is now defined by setting

$$[v] = \{v\} \cup \{u \in V : \text{there is a path of reticulation arcs from } u \text{ to } v \text{ in } \mathcal{N}\}.$$

Observe that if v is not incident with a reticulation arc, then $[v] = \{v\}$. Set

$$[V] = \{[v] : v \in V\}.$$

Essentially, calling TEMPREP for \mathcal{N} results in the following two steps. First, a digraph $D_{\mathcal{N}}$ is constructed whose vertex set is $[V]$, and for which $[u]$ and $[v]$ are joined by an arc $([u], [v])$ precisely if there exist $a \in [u]$ and $b \in [v]$ such that (a, b) is a tree arc in \mathcal{N} . Second, it is checked whether $D_{\mathcal{N}}$ is acyclic or not. In Baroni et al. (2006, Theorem 3), it has been shown that \mathcal{N} has a temporal representation if and only if $D_{\mathcal{N}}$ is acyclic. Inspecting the associated proof reveals




Fig. 4 The base case for the proof of Theorem 3.5

that the same applies for when \mathcal{N} is an HGT network. Note that both steps of the algorithm can be computed in polynomial time. Thus, **TEMPREP** is a polynomial-time algorithm, and $\text{ADDTAXA}_{\text{HGT}}(\mathcal{N}, k)$ and $\text{ADDTAXA}_{\text{Hybrid}}(\mathcal{N}, k)$ are consequently in NP.

Using the two previous remarks, we next prove the main result of this section.

Theorem 3.5 *The decision problem $\text{ADDTAXA}_{\text{HGT}}(\mathcal{N}, k)$ is NP-complete.*

Proof By Remark 2, first note that $\text{ADDTAXA}_{\text{HGT}}(\mathcal{N}, k)$ is in NP as we can apply the polynomial-time algorithm **TEMPREP** to determine whether a given HGT network has a temporal representation or not. Now, making use of Remark 1, to complete the proof it is sufficient to show that, given an instance G of **FEEDBACKVERTEXSET**, we can construct in polynomial time an HGT network \mathcal{N} whose size is polynomial in the size of G and for which $\mathcal{C}_{\mathcal{N}}$ and G are the same. The proof that such an HGT network always exists is by induction on the number k of arcs of G .

If $k = 0$, then G consists of only isolated vertices, u_1, u_2, \dots, u_n (where $n \in \mathbb{N}_0$ arbitrary), and the HGT network shown in Fig. 4 has the property that its critical graph consists of isolated vertices u_1, u_2, \dots, u_n . Since the size of this network is polynomial in the size of G , this establishes the base case of the induction. Now assume that, if an instance of **FEEDBACKVERTEXSET** has $k - 1$ arcs, where $k \geq 1$, then there is an HGT network with the desired properties.

Let G' denote the directed graph obtained from G by deleting an arbitrary arc (u, v) . By the induction assumption, there is an HGT network \mathcal{N}' that can be constructed in polynomial time and is polynomial in the size of G' , and has the property that $\mathcal{C}_{\mathcal{N}'}$ and G' are the same. As u and v are vertices in G' , they are parents of HGT vertices in \mathcal{N}' . Call the corresponding HGT vertices u_1 and v_1 . Note that u is not an ancestor of v_1 (and therefore also not of v) since (u, v) is not an arc in G' . Furthermore, note that u and v are tree vertices. A generic picture of \mathcal{N}' is shown in Fig. 5. We complete the proof by modifying \mathcal{N}' to obtain an HGT network \mathcal{N} such that $\mathcal{C}_{\mathcal{N}}$ and G are the same. $\mathcal{C}_{\mathcal{N}}$ and G being equal implies in particular that u becomes an ancestor of v_1 in \mathcal{N} .

Fig. 5 The *bold triangle* represents the HGT network \mathcal{N}' in the proof of Theorem 3.5. Only the vertices $u, u_1, v,$ and v_1 are drawn explicitly as we need them to obtain \mathcal{N} from \mathcal{N}' (see Fig. 6). Note that u is not an ancestor of v_1 , however, v is an ancestor of u_1 if (v, u) is an arc in G'

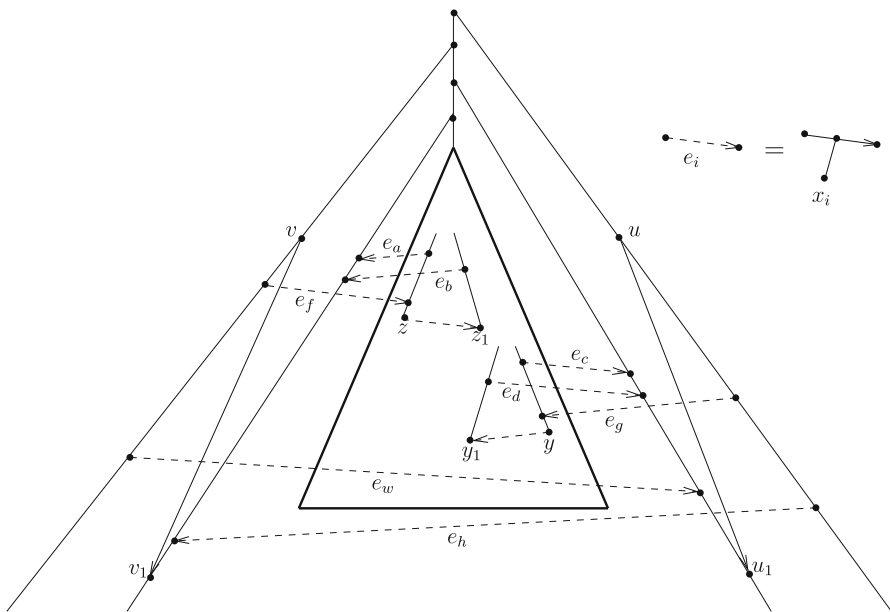



Fig. 6 The HGT network \mathcal{N} constructed from \mathcal{N}' in the proof of Theorem 3.5. The part of the network enclosed by the *solid triangle* is the same as that shown in Fig. 5 except that vertices $u, u_1, v,$ and v_1 have been renamed as $y, y_1, z,$ and z_1 , respectively. Each *dashed line* e_i with $i \in \{a, b, c, d, f, g, h, w\}$ represents a 2-arc directed path (with the middle vertex omitted) and the attachment of a new taxa x_i (also omitted) adjoined to the middle vertex on the path, see *top right* of the figure. This means that the endpoint of each *dashed line* corresponds to a non-critical HGT vertex

We modify \mathcal{N}' in the following way: Rename the vertices $u, u_1, v,$ and v_1 to $y, y_1, z,$ and z_1 , respectively. Further, add a new taxa across (y, y_1) and (z, z_1) , so that y and z are not critical vertices. We add an arc ancestral to the root in \mathcal{N}' , which has four new arcs descending (see Fig. 6). The vertices $u, u_1, v,$ and v_1 are placed on these four arcs (one on each) and HGT arcs (u, u_1) and (v, v_1) are added. In order to preserve all relationships between critical vertices occurring in \mathcal{N}' , we add the seven HGT arcs

$e_a, e_b, e_c, e_d, e_f, e_g$ (and possibly the arc e_w if (v, u) is an arc in G'), each with a new taxa across the arc so that the added HGT vertices are not critical (see Fig. 6):

- The arcs e_a and e_b are included so that each ancestor of the critical HGT vertex v_1 in \mathcal{N}' is also an ancestor of the critical HGT vertex v_1 in the modified HGT network.
- The arcs e_c and e_d are included so that each ancestor of the critical HGT vertex u_1 in \mathcal{N}' is also an ancestor of the critical HGT vertex u_1 in the modified HGT network.
- The arc e_f is included so that each descendant of the critical vertex v in \mathcal{N}' is also a descendant of the critical vertex v in the modified HGT network.
- The arc e_g is included so that each descendant of the critical vertex u in \mathcal{N}' is also a descendant of the critical vertex u in the modified HGT network.
- If (v, u) is an arc in G , then the arc e_w is also included into the construction so that v is an ancestor of u_1 in the modified HGT network.

Lastly, as (u, v) is an arc in G , the arc e_h is added with a new taxa across it so that u is an ancestor of v_1 in the modified HGT network.

We call this resulting HGT network \mathcal{N} . Clearly, the construction of \mathcal{N} from \mathcal{N}' takes polynomial time and so, by the induction assumption the construction of \mathcal{N} from G also takes polynomial time. Moreover, the size of \mathcal{N} is polynomial in the size of \mathcal{N}' and so, again by the induction assumption, the size of \mathcal{N} is polynomial in the size of G .

It remains to verify that the critical graph of \mathcal{N} is G . Because of the way in which \mathcal{N} is constructed from \mathcal{N}' and noting that neither y nor z is critical in \mathcal{N} , the set of critical vertices of \mathcal{N} is exactly the same as the set of critical vertices of \mathcal{N}' . Thus the vertex sets of $\mathcal{C}_{\mathcal{N}}$ and G are the same. Furthermore, by construction, if w and x are critical vertices of \mathcal{N} and $\{w, x\} \cap \{u, v\} = \emptyset$, then (w, x) is an arc in $\mathcal{C}_{\mathcal{N}}$ precisely if (w, x) is an arc in G . Thus to show that $\mathcal{C}_{\mathcal{N}}$ and G are the same, it remains to check, for each $j \in \{u, v\}$, that (w, j) is an arc in $\mathcal{C}_{\mathcal{N}}$ precisely if (w, j) is an arc in G and that (j, w) is an arc in $\mathcal{C}_{\mathcal{N}}$ precisely if (j, w) is an arc in G . First observe that because of the addition of e_h , we have (u, v) in $\mathcal{C}_{\mathcal{N}}$. Furthermore, if (v, u) is in G , then e_w is added and so (v, u) is in $\mathcal{C}_{\mathcal{N}}$. We next complete the check for when $j = v$. The check for when $j = u$ is similar and omitted.

Since any ancestor of v in \mathcal{N}' remains an ancestor of v_1 in \mathcal{N} and any ancestor of v_1 in \mathcal{N}' remains an ancestor of v_1 in \mathcal{N} because of the addition of e_a and e_b , respectively, it follows by the induction assumption that (w, v) is an arc in $\mathcal{C}_{\mathcal{N}}$ precisely if (w, v) is an arc in G . Furthermore, since any descendant of v in \mathcal{N}' remains a descendant of v in \mathcal{N} because of the addition of e_f , it follows by the induction assumption that (v, w) is an arc in $\mathcal{C}_{\mathcal{N}}$ precisely if (v, w) is an arc in G . It now follows that the two graphs $\mathcal{C}_{\mathcal{N}}$ and G are equal. This completes the proof of the theorem.

Given an instance of $\text{ADDTAXA}_{\text{HGT}}(\mathcal{N}, k)$, consider the instance of $\text{ADDTAXA}_{\text{Hybrid}}(\mathcal{N}', k)$, where \mathcal{N}' is obtained from \mathcal{N} by adding a new taxa across the parent tree arc of each HGT vertex and then viewing each such vertex as a hybridization vertex. Observe that \mathcal{N} is a temporal HGT network if and only if \mathcal{N}' is a temporal hybridization network. Moreover, if \mathcal{N} is not temporal, it is easily seen that the minimum number of new taxa one needs to add to \mathcal{N} so that the resulting HGT network is temporal is equal to the minimum number of new taxa one needs to add to \mathcal{N}' so that the resulting hybridization network is temporal. Thus the answer to

$\text{ADDTAXA}_{\text{HGT}}(\mathcal{N}, k)$ is *yes* if and only if the answer to $\text{ADDTAXA}_{\text{Hybrid}}(\mathcal{N}', k)$ is *yes*. Since the construction of \mathcal{N}' from \mathcal{N} takes polynomial time in the size of \mathcal{N} and since, by Remark 2 prior to the statement of Theorem 3.5, we can apply the polynomial-time algorithm TEMPREP to determine whether a given hybridization network has a temporal representation, the following corollary now follows from Theorem 3.5.

Corollary 3.6 *The decision problem $\text{ADDTAXA}_{\text{Hybrid}}(\mathcal{N}, k)$ is NP-complete.*

We end this section with a brief discussion on solving ADDTAXA in reasonable time. For an HGT network \mathcal{N} , the proof of Theorem 3.5 provides an algorithm for solving $\text{ADDTAXA}(\mathcal{N}, k)$. In particular, construct the critical graph of \mathcal{N} and use an exact algorithm for $\text{FEEDBACKVERTEXSET}(\mathcal{C}_{\mathcal{N}}, k)$. For hybridization networks, we can use the same approach as the analogous results for HGT networks also hold in this setting.

It is straightforward to check that Lemma 3.2 and Proposition 3.3 hold for hybridization networks. For the analogous proof of Proposition 3.3, note that, in the context of hybridization networks, if v_1 is a hybridization vertex with parents v and v' , and both v and v' are in the critical graph, then (u, v) is an arc in $\mathcal{C}_{\mathcal{N}}$ if and only if (u, v') is an arc in $\mathcal{C}_{\mathcal{N}}$. Thus v and v' can be labeled appropriately, similar to the procedure described in the first part of the proof of this proposition for labeling an HGT network. Since Lemma 3.2 and Proposition 3.3 hold for hybridization networks, Corollary 3.4 holds for hybridization networks. Hence, one can also solve $\text{ADDTAXA}(\mathcal{N}, k)$ for when \mathcal{N} is a hybridization network by constructing $\mathcal{C}_{\mathcal{N}}$ and using an exact algorithm for $\text{FEEDBACKVERTEXSET}(\mathcal{C}_{\mathcal{N}}, k)$.

It is shown in Chen et al. (2008) that FEEDBACKVERTEXSET for directed graphs is fixed-parameter tractable. The authors provide an algorithm which solves $\text{FEEDBACKVERTEXSET}(G, k)$ in $O(4^k k! n^{O(1)})$ where n is the number of vertices in G . Thus if k is relatively small, this algorithm will work reasonably quickly in practice regardless of the size of G .

For a reticulation network \mathcal{N} , the size of $\mathcal{C}_{\mathcal{N}}$ is determined by the number of critical vertices of \mathcal{N} . This number is less than the number of reticulation vertices and we would expect this latter number to be much less than the size of \mathcal{N} . Consequently, k may typically be relatively small and so the algorithm in Chen et al. (2008) should work well in helping to find the solution for many instances of $\text{ADDTAXA}(\mathcal{N}, k)$.

4 Constructing temporal hybridization networks for two trees

In the previous section, we determined the smallest number of taxa that must be added to a reticulation network so that the resulting network is temporal. However, for many evolutionary studies, we are initially given a set of gene trees rather than a reticulation network. In such a case, in particular, when there are no unsampled taxa, it is of importance to decide whether a temporal reticulation network exists that simultaneously explains the evolutionary histories of the gene trees under consideration. An example showing that the existence of such a network is not guaranteed is given after some preliminaries.

In this section, we analyze the construction of temporal hybridization networks from so-called agreement forests. Such forests are frequently used to model reticulate




Fig. 7 A rooted binary phylogenetic X -tree \mathcal{T} , and the two subtrees $\mathcal{T}(Y)$ and $\mathcal{T}|Y$ with $Y = \{a, d, e\}$

evolution for when two rooted binary phylogenetic trees \mathcal{T} and \mathcal{T}' are given (see for example Baroni et al. 2005; Bordewich and Semple 2007; Hein et al. 1996). Our analysis centers around a new algorithm, called TEMPORALHYBRID, which reconstructs a temporal hybridization network for \mathcal{T} and \mathcal{T}' if one exists. Two applications of the algorithm are given at the end of this section, with the second application being illustrated on a grass data set. Throughout this section, we restrict our attention to hybridization networks.

4.1 The hybridization number

Before detailing TEMPORALHYBRID, we need some additional definitions.

4.1.1 Hybrid phylogenies and the hybridization number

A rooted binary phylogenetic X -tree \mathcal{T} is a rooted tree whose root has degree two and all other internal vertices have degree three, and whose leaf set is X . The set X is called the *label set* of \mathcal{T} and is denoted by $\mathcal{L}(\mathcal{T})$. For a subset A of X , the minimal rooted subtree of \mathcal{T} that connects all the elements in A is denoted by $\mathcal{T}(A)$. Furthermore, the *restriction of \mathcal{T} to A* , denoted by $\mathcal{T}|A$, is the rooted phylogenetic tree obtained from $\mathcal{T}(A)$ by contracting every vertex of degree-2 apart from the root. An example of both types of subtrees is shown in Fig. 7.

Let \mathcal{T} be a rooted binary phylogenetic X' -tree, and let \mathcal{N} be a hybridization network with label set X , where $X' \subseteq X$. We say that \mathcal{N} *displays* \mathcal{T} if all of the ancestral relationships described in \mathcal{T} are preserved by \mathcal{N} . Formally, \mathcal{N} displays \mathcal{T} if \mathcal{T} can be obtained from \mathcal{N} by deleting a subset of arcs and vertices of \mathcal{N} , and contracting any resulting degree-2 vertices apart from the root.

A fundamental problem for biologists studying a set of present-day species whose evolutionary history includes hybridization is to determine the extent to which hybridization has influenced their past. For two rooted binary phylogenetic X -trees \mathcal{T} and \mathcal{T}' , a common way to quantify this extent is by determining the value

$$h(\mathcal{T}, \mathcal{T}') = \min\{h(\mathcal{N}) : \mathcal{N} \text{ is a hybridization network on } X \text{ that displays } \mathcal{T} \text{ and } \mathcal{T}'\},$$

where $h(\mathcal{N})$ is the number of hybridization vertices of \mathcal{N} . However, Bordewich and Semple (2007) showed that determining this number is an NP-hard problem.

We next give an example of two trees which cannot be displayed in a temporal hybridization network. Consider the two rooted binary phylogenetic X -trees \mathcal{T}




Fig. 8 Two rooted binary phylogenetic X -trees \mathcal{T} and \mathcal{T}' , and all nine hybridization networks on X displaying both trees with two hybridization vertices. None of the hybridization networks is temporal

and \mathcal{T}' shown at the top of Fig. 8. It is easily seen that every hybridization network that displays \mathcal{T} and \mathcal{T}' has at least two reticulation vertices. There are exactly nine hybridization networks on X , each with two hybridization vertices, displaying \mathcal{T} and \mathcal{T}' . These networks are shown in the bottom part of Fig. 8. None of these networks is temporal. Moreover, a straightforward check shows that any hybridization network that displays \mathcal{T} and \mathcal{T}' with more than two hybridization vertices is not temporal either. Thus, there is no temporal hybridization network on X that displays \mathcal{T} and \mathcal{T}' .

4.1.2 Agreement forests

The approach taken by TEMPORALHYBRID (see Sect. 4.2) is based on the concept of a so-called acyclic-agreement forest for two rooted binary phylogenetic X -trees \mathcal{T} and \mathcal{T}' . Loosely speaking, the smallest size of such a forest for \mathcal{T} and \mathcal{T}' equates with $h(\mathcal{T}, \mathcal{T}')$. We next make this precise.

Let \mathcal{T} and \mathcal{T}' be two rooted binary phylogenetic X -trees. For the upcoming definitions, we regard the roots of both \mathcal{T} and \mathcal{T}' as an extra vertex ρ adjoined to the original root by an additional edge and $\mathcal{L}(\mathcal{T}) = \mathcal{L}(\mathcal{T}') = X \cup \{\rho\}$. An *agreement forest* for \mathcal{T} and \mathcal{T}' is a collection $\mathcal{F} = \{\mathcal{S}_\rho, \mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_k\}$ of rooted leaf-labeled trees, where \mathcal{S}_ρ is a rooted tree whose label set \mathcal{L}_ρ contains ρ and $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_k$

are rooted binary phylogenetic trees with label sets $\mathcal{L}_1, \mathcal{L}_2, \dots, \mathcal{L}_k$, respectively, such that the following properties hold:

- (i) The label sets $\mathcal{L}_\rho, \mathcal{L}_1, \mathcal{L}_2, \dots, \mathcal{L}_k$ partition $X \cup \{\rho\}$.
- (ii) For all $i \in \{1, 2, \dots, k\}$, $\mathcal{S}_i \cong \mathcal{T}|_{\mathcal{L}_i}$ and $\mathcal{S}_i \cong \mathcal{T}'|_{\mathcal{L}_i}$.
- (iii) The trees in $\{\mathcal{T}(\mathcal{L}_i) : i \in \{1, 2, \dots, k\}\}$ and $\{\mathcal{T}'(\mathcal{L}_i) : i \in \{1, 2, \dots, k\}\}$ are vertex-disjoint rooted subtrees of \mathcal{T} and \mathcal{T}' , respectively.

We sometimes refer to the component \mathcal{S}_ρ of \mathcal{F} as the *root component* of \mathcal{F} . To illustrate, consider the two trees \mathcal{T} and \mathcal{T}' in Fig. 8. Viewing the root of each of \mathcal{T} and \mathcal{T}' as an extra vertex ρ adjoined to the original root as described above, the restrictions of \mathcal{T} (and \mathcal{T}') to the label sets $\{\rho, a, d\}$, $\{b\}$, and $\{c\}$, respectively, are the components of an agreement forest of \mathcal{T} and \mathcal{T}' .

To make the connection between hybridization networks and agreement forests, we need to extend the definition of an agreement forest. This extension allows for the biological constraint that species cannot inherit genetic material from their own descendants. Let $G_{\mathcal{F}}$ be the directed graph whose vertex set is \mathcal{F} and for which $(\mathcal{S}_i, \mathcal{S}_j)$ is an arc precisely if $i \neq j$ and

- (I) the root of $\mathcal{T}(\mathcal{L}_i)$ is an ancestor of the root of $\mathcal{T}(\mathcal{L}_j)$ or
- (II) the root of $\mathcal{T}'(\mathcal{L}_i)$ is an ancestor of the root of $\mathcal{T}'(\mathcal{L}_j)$.

We say that \mathcal{F} is *acyclic* precisely if $G_{\mathcal{F}}$ is acyclic. If \mathcal{F} is acyclic and it has the smallest number of components amongst all such forests of \mathcal{T} and \mathcal{T}' , then \mathcal{F} is a *maximum-acyclic-agreement forest* of \mathcal{T} and \mathcal{T}' , in which case we denote $|\mathcal{F}| - 1$ by $m_a(\mathcal{T}, \mathcal{T}')$.

The minimum number $h(\mathcal{T}, \mathcal{T}')$ of hybridization events and the size of a maximum-acyclic-agreement forest of \mathcal{T} and \mathcal{T}' are related through the following theorem (Baroni et al. 2005).

Theorem 4.1 *Let \mathcal{T} and \mathcal{T}' be two rooted binary phylogenetic X -trees. Then*

$$h(\mathcal{T}, \mathcal{T}') = m_a(\mathcal{T}, \mathcal{T}').$$

We now define an *ordering* of an acyclic-agreement forest which will be used as an input to the algorithm TEMPORALHYBRID. This algorithm constructs a temporal hybridization network if one exists by iterating through such an ordering. For an acyclic-agreement forest \mathcal{F} , a tuple $\mathcal{O} = (\mathcal{S}_\rho, \mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_k)$ is an *ordering* of the components of \mathcal{F} if, for each i , the vertex \mathcal{S}_i has indegree 0 in the graph obtained from $G_{\mathcal{F}}$ by deleting the vertices $\mathcal{S}_\rho, \mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_{i-1}$ and their incident arcs. Since \mathcal{F} is acyclic, such an ordering always exists.

Although not explicitly stated in Baroni et al. (2005), one direction of Theorem 4.1 is essentially established by proving that, if \mathcal{N} is a hybridization network on X that displays \mathcal{T} and \mathcal{T}' , then there is an acyclic-agreement forest \mathcal{F} for \mathcal{T} and \mathcal{T}' such that $|\mathcal{F}| \leq h(\mathcal{N}) + 1$. Intuitively, one takes \mathcal{N} and iteratively cuts off rooted subtrees by deleting hybridization vertices and their three incident arcs. By viewing the root of \mathcal{N} as a vertex at the end of a pendant arc adjoined to the original root, we obtain an acyclic-agreement forest \mathcal{F} of \mathcal{T} and \mathcal{T}' , and so $|\mathcal{F}| - 1 \leq h(\mathcal{N})$. Now, if we extend their argument and suppose that \mathcal{N} is temporal, then we can obtain an

acyclic-agreement forest for \mathcal{T} and \mathcal{T}' as follows. First, select a vertex v that either (i) is a hybridization vertex and, except for v itself, its parents have no hybridization vertex as a descendant or (ii) is a tree vertex whose parent is a tree vertex that has no hybridization vertex as a descendant. Now delete v and its three incident arcs if v is a hybridization vertex or delete the parent arc of v if v is a tree vertex, and contract any non-root degree-2 vertex. Repeating this process so that every hybridization vertex is selected, and then reversing the order of the selections, we eventually obtain an ordering of an acyclic-agreement forest for \mathcal{T} and \mathcal{T}' . Given an ordering \mathcal{O} of an arbitrary acyclic-agreement forest \mathcal{F} of two rooted binary phylogenetic X -trees, we say that a temporal hybridization network \mathcal{N} *preserves* \mathcal{O} if \mathcal{O} can be obtained from \mathcal{N} in the above way.

4.2 The algorithm TEMPORALHYBRID

In this section, we present the algorithm TEMPORALHYBRID. This algorithm takes two rooted binary phylogenetic X -trees \mathcal{T} and \mathcal{T}' , and an ordering \mathcal{O} of an acyclic-agreement forest \mathcal{F} of \mathcal{T} and \mathcal{T}' as input, and outputs either

- (i) a temporal hybridization network on X that displays \mathcal{T} and \mathcal{T}' , and preserves \mathcal{O} , or
- (ii) a statement indicating that there is no such network.

Baroni et al. (2006), have previously presented a similar algorithm. Called HYBRIDPHYLOGENY, this algorithm has the same input as TEMPORALHYBRID but without an ordering \mathcal{O} of \mathcal{F} . The task for HYBRIDPHYLOGENY is simply to construct one of potentially many hybridization networks that display \mathcal{T} and \mathcal{T}' . However, in doing so, there is no guarantee that the resulting network is temporal. The correctness of TEMPORALHYBRID is established after some remarks following the description of the algorithm.

Algorithm: TEMPORALHYBRID($\mathcal{T}, \mathcal{T}', \mathcal{O}$)

Input: Two rooted binary phylogenetic X -trees \mathcal{T} and \mathcal{T}' , and an ordering $\mathcal{O} = (\mathcal{S}_\rho, \mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_k)$ of an acyclic-agreement forest \mathcal{F} for \mathcal{T} and \mathcal{T}' .

Output: A temporal hybridization network on X with at most k hybridization vertices that displays \mathcal{T} and \mathcal{T}' and preserves \mathcal{O} , or a statement indicating that there is no such network.

1. For each $i \in \{1, \dots, k\}$, regard ρ_i as the label of an extra vertex adjoined to the original root of \mathcal{S}_i by an additional edge.
2. Set $\mathcal{N}_0 = \mathcal{S}_\rho$, and set $i = 1$.
3. Attach \mathcal{S}_i to \mathcal{N}_{i-1} via at most two new arcs. Each new arc joins the vertex labeled ρ_i to a new vertex which subdivides an arc of \mathcal{N}_{i-1} . The subdivided arcs $e = (u, v)$ and $e' = (u', v')$ are chosen so that
 - (i) the resulting network displays $\mathcal{T}|(\mathcal{L}(\mathcal{N}_{i-1}) \cup \mathcal{L}(\mathcal{S}_i))$ and $\mathcal{T}'|(\mathcal{L}(\mathcal{N}_{i-1}) \cup \mathcal{L}(\mathcal{S}_i))$, where $\mathcal{L}(\mathcal{N}_{i-1}) = \mathcal{L}(\mathcal{S}_\rho) \cup \mathcal{L}(\mathcal{S}_1) \cup \dots \cup \mathcal{L}(\mathcal{S}_{i-1})$, and

- (ii) no element of $\{\rho_1, \rho_2, \dots, \rho_{i-1}\}$ is a descendant of either v or v' and, if $e \neq e'$, then e and e' are not on the same path from the root of \mathcal{N}_{i-1} to one of its leaves.

Set \mathcal{N}_i to be the resulting network. If there is no such attachment for \mathcal{S}_i , then stop and return *there is no temporal hybridization network on X that displays \mathcal{T} and \mathcal{T}' , and preserves \mathcal{O}* .

4. If $i = k$, remove the arc incident with the vertex labeled ρ and remove all labels in $\{\rho, \rho_1, \rho_2, \dots, \rho_k\}$ from \mathcal{N}_k , contract any resulting degree-0 and degree-2 vertices apart from the root, and return the obtained network. If $i < k$, increment i by 1 and return to Step 3.

Remarks 1. If \mathcal{N}_0 consists of an isolated vertex labeled ρ , a slight complication arises in Step 3 of the algorithm since no arc can be subdivided by a new vertex. In this case, \mathcal{S}_1 is adjoined to \mathcal{N}_0 by adding precisely one new arc that joins the vertex labeled ρ_1 with the vertex labeled ρ .

2. For all $\mathcal{S}_i \in \mathcal{F}$, the algorithm TEMPORALHYBRID potentially checks each arc of \mathcal{N}_{i-1} to decide whether \mathcal{S}_i can be appropriately attached to \mathcal{N}_{i-1} . Thus, the running time of the algorithm is at most $O(kn)$, where $n = |X|$.
3. If \mathcal{F} is a maximum-acyclic-agreement forest, then the attachment of a new component \mathcal{S}_i to \mathcal{N}_{i-1} always requires two new arcs in Step 3 of the algorithm; otherwise, \mathcal{F} is not optimal.
4. If TEMPORALHYBRID returns a temporal hybridization network, then this is the unique such network on X that displays \mathcal{T} and \mathcal{T}' , and preserves \mathcal{O} (see Theorem 4.2).

We next prove the correctness of TEMPORALHYBRID. In doing this, we additionally show that if the algorithm returns a hybridization network, then this network is unique relative to the ordering of \mathcal{O} .

Theorem 4.2 *Let \mathcal{T} and \mathcal{T}' be two rooted binary phylogenetic X -trees, and let \mathcal{O} be an ordering of an acyclic-agreement forest \mathcal{F} for \mathcal{T} and \mathcal{T}' with $|\mathcal{F}| - 1 = k$. Then*

- (i) TEMPORALHYBRID returns a temporal hybridization network \mathcal{N} that displays \mathcal{T} and \mathcal{T}' with $h(\mathcal{N}) \leq k$ and that preserves \mathcal{O} if there exists such a network, or
- (ii) TEMPORALHYBRID returns a statement indicating that there is no such network.

Moreover, if TEMPORALHYBRID returns a temporal hybridization network, then, up to isomorphism, this is the unique network with the properties in (i).

Proof Without loss of generality, let $\mathcal{O} = (\mathcal{S}_\rho, \mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_k)$. The proof is by induction on k . If $k = 0$, then $\mathcal{S}_\rho \cong \mathcal{T} \cong \mathcal{T}'$ and the theorem trivially holds. Now assume that $k > 0$ and that the theorem holds for all orderings of all acyclic-agreement forests with at most k components of two rooted binary phylogenetic trees. Let \mathcal{T}_1 be the rooted binary phylogenetic tree $\mathcal{T}|(X - \mathcal{L}(\mathcal{S}_k))$, and let \mathcal{T}'_1 be the rooted binary phylogenetic tree $\mathcal{T}'|(X - \mathcal{L}(\mathcal{S}_k))$. Since \mathcal{S}_k is the last coordinate in \mathcal{O} , the trees \mathcal{T}_1 and \mathcal{T}'_1 can also be obtained from \mathcal{T} and \mathcal{T}' , respectively, by deleting a single edge and contracting the resulting degree-2 vertex. Let $\mathcal{F}_1 = \mathcal{F} - \{\mathcal{S}_k\}$. As \mathcal{O} is an ordering of \mathcal{F} , it follows that $\mathcal{O}_1 = (\mathcal{S}_\rho, \mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_{k-1})$ is an ordering of \mathcal{F}_1 . Now observe that

the workings of TEMPORALHYBRID applied to $(\mathcal{T}_1, \mathcal{T}'_1, \mathcal{O}_1)$ and applied to $(\mathcal{T}, \mathcal{T}', \mathcal{O})$ up to considering \mathcal{S}_k are identical. This observation is used in the rest of the proof.

Since $|\mathcal{O}_1| < |\mathcal{O}|$, it follows by the induction assumption that TEMPORALHYBRID $(\mathcal{T}_1, \mathcal{T}'_1, \mathcal{O}_1)$ either returns, up to isomorphism, a unique temporal hybridization network, \mathcal{N}_{k-1} say, that displays \mathcal{T}_1 and \mathcal{T}'_1 and preserves \mathcal{O}_1 , or returns a statement indicating that there is no such network. First suppose that TEMPORALHYBRID $(\mathcal{T}_1, \mathcal{T}'_1, \mathcal{O}_1)$ returns the latter. If there is a temporal hybridization network that displays \mathcal{T} and \mathcal{T}' , and preserves \mathcal{O} , let v be the vertex of \mathcal{N} that corresponds to the root of \mathcal{S}_k . Then by deleting v and its incident arcs or the parent arc of v , depending on whether v is a hybridization vertex or a tree vertex, respectively, and contracting any resulting degree-2 vertices apart from the root, we have a temporal hybridization network that displays \mathcal{T}_1 and \mathcal{T}'_1 , and preserves \mathcal{O}_1 ; a contradiction. Thus, if TEMPORALHYBRID $(\mathcal{T}, \mathcal{T}', \mathcal{O})$ returns a statement that there is no such network prior to considering \mathcal{S}_k , then it returns correctly. Therefore, we may suppose that TEMPORALHYBRID $(\mathcal{T}_1, \mathcal{T}'_1, \mathcal{O}_1)$ returns \mathcal{N}_{k-1} . By the observation at the end of the last paragraph, this means that \mathcal{N}_{k-1} is constructed immediately prior to considering \mathcal{S}_k in TEMPORALHYBRID $(\mathcal{T}, \mathcal{T}', \mathcal{O})$.

If TEMPORALHYBRID $(\mathcal{T}, \mathcal{T}', \mathcal{O})$ returns a statement indicating that there is no appropriate network, then, because of the uniqueness of \mathcal{N}_{k-1} and the fact that (i) and (ii) in Step 3 of the algorithm are necessary conditions for the placement of \mathcal{S}_k , the algorithm performs correctly. On the other hand, if TEMPORALHYBRID $(\mathcal{T}, \mathcal{T}', \mathcal{O})$ returns a network \mathcal{N} , it follows by (i) and (ii) in Step 3 that \mathcal{N} displays \mathcal{T} and \mathcal{T}' , and \mathcal{N} is temporal and preserves \mathcal{O} . Furthermore, let $e = (u, v)$ and $e' = (u', v')$ denote the arcs of \mathcal{N}_{k-1} that are subdivided in the attachment of \mathcal{S}_k . By (ii) in Step 3, no hybridization vertex is a descendant of v or v' , so the choice of e and e' is unique. By the uniqueness of \mathcal{N}_{k-1} , it follows that \mathcal{N} is the unique temporal hybridization network that displays \mathcal{T} and \mathcal{T}' , and preserves \mathcal{O} . Note that this argument also holds if $e = e'$ in which case \mathcal{S}_k is attached by a single arc. This completes the proof of the theorem.

We saw earlier that, for two rooted binary phylogenetic X -trees, there may not exist a temporal hybridization network on X that displays both trees. In general, how does one decide such an outcome for two rooted binary phylogenetic X -trees \mathcal{T} and \mathcal{T}' ? Of course, by considering all orderings of all acyclic-agreement forests for \mathcal{T} and \mathcal{T}' , we can repeatedly use TEMPORALHYBRID to decide whether or not such a hybridization network exists. While still exponential in time, we can do much better than this by only considering certain orderings and certain forests, as shown in the next proposition.

For two rooted binary phylogenetic X -trees \mathcal{T} and \mathcal{T}' , an acyclic-agreement forest $\mathcal{F} = \{\mathcal{S}_\rho, \mathcal{S}_1, \dots, \mathcal{S}_k\}$ of \mathcal{T} and \mathcal{T}' is *trivial* if $|\mathcal{L}(\mathcal{S}_\rho)| = 3$ and each of $\mathcal{S}_1, \dots, \mathcal{S}_k$ consists of an isolated vertex. Provided \mathcal{S}_ρ is the first coordinate, any $(k + 1)$ -tuple of \mathcal{F} is an ordering of \mathcal{F} . The next proposition shows that it is sufficient to consider each such ordering of \mathcal{F} to decide whether or not there exists a temporal hybridization network that displays \mathcal{T} and \mathcal{T}' .

Proposition 4.3 *Let \mathcal{T} and \mathcal{T}' be two rooted binary phylogenetic X -trees, and let $|X| = n$. Deciding whether there is a temporal hybridization network on X that displays \mathcal{T} and \mathcal{T}' takes at most time $O(n! \cdot p(a))$, where $p(a)$ is the polynomial running time of the algorithm TEMPORALHYBRID.*

Proof Suppose there is a temporal hybridization network \mathcal{N} on X that displays \mathcal{T} and \mathcal{T}' . Let $X = \{a_1, a_2, \dots, a_n\}$. We next construct an ordering of a trivial forest of \mathcal{T} and \mathcal{T}' , that is preserved by \mathcal{N} . Select an element of X so that its parent is a tree vertex and does not have a hybridization vertex as a descendant. If this is not possible, then, as \mathcal{N} is acyclic, there exists a hybridization vertex whose child is an element of X and whose parents have no hybridization vertex as a descendant. Without loss of generality, we may assume that the selected element is a_n . If the parent of a_n is a tree vertex, then delete the parent arc of a_n and contract the resulting degree-2 vertex. Since the parent of a_n has no hybridization vertex as a descendant, the resulting network is again a hybridization network. If the parent of a_n is a hybridization vertex, then delete the parent vertex of a_n and its three incident arcs, and contract the resulting degree-2 vertices. By property (iv) in the definition of a reticulation network, it is easily checked that the arcs incident with the contracted degree-2 vertices are tree arcs, and thus, the resulting network is again a hybridization network. Furthermore, in both cases, the resulting hybridization network is temporal as \mathcal{N} is temporal. Selecting another vertex and continuing in this way, we eventually select (in order) the vertices a_n, a_{n-1}, \dots, a_3 say. This gives a trivial forest \mathcal{F} of \mathcal{T} and \mathcal{T}' , where the label set of the root component is $\{\rho, a_1, a_2\}$ and the remaining components consist of isolated vertices labeled a_3, a_4, \dots, a_n . Furthermore, the tuple beginning with the root component and followed (in order) by the components whose label sets are $\{a_3\}, \{a_4\}, \dots, \{a_n\}$ is an ordering \mathcal{O} of \mathcal{F} and, by construction, \mathcal{N} preserves \mathcal{O} . Now, by Theorem 4.2, calling `TEMPORALHYBRID`($\mathcal{T}, \mathcal{T}', \mathcal{O}$) returns \mathcal{N} . Thus to decide whether or not there is a temporal hybridization network on X that displays \mathcal{T} and \mathcal{T}' , it suffices to consider all possible acyclic-agreement forests of \mathcal{T} and \mathcal{T}' that are trivial. Since there are $\binom{n}{2} \cdot (n-2)! = \frac{1}{2}n!$ such forests, the proposition now follows.

While the above approach is not fast because of the number of orderings to consider, we can do much better in practice if we restrict our attention to maximum-acyclic-agreement forests. We do this in the next section.

4.3 Minimal temporal hybridization networks

To provide an indication of the significance of hybridization, biologists are often interested in reconstructing (temporal) hybridization networks that explain the ancestral history of the species under consideration and simultaneously minimize the number of hybridization events. This minimum number provides a lower bound on the number of such events, thus it gives an indication of the role that hybridization has had on the evolution of the present-day species. In this section, we consider an approach to this task with the view of constructing hybridization networks that are temporal.

Let \mathcal{T} and \mathcal{T}' be two rooted binary phylogenetic X -trees. A (temporal) hybridization network \mathcal{N} that displays \mathcal{T} and \mathcal{T}' is *minimal* if $h(\mathcal{N}) = m_a(\mathcal{T}, \mathcal{T}')$. Since we are using a combinatorial framework to calculate the number of hybridization events, it is likely that there exist several maximum-acyclic-agreement forests for \mathcal{T} and \mathcal{T}' . For example, for the grass (*Poaceae*) data set that has been analyzed in [Bordewich et al. \(2007\)](#), the associated gene loci for 12 gene tree pairs, the minimum number of hybridization events, and the number of maximum-acyclic-agreement forests are

Table 1 Results for the *Poaceae* data set

Pairwise combination		Hybridization number $h(\mathcal{T}, \mathcal{T}')$	# MAAFs	# MAAFs with a proper root component
<i>ndhF</i>	<i>phyB</i>	14	2,268	0
<i>ndhF</i>	<i>rbcL</i>	13	48	0
<i>ndhF</i>	<i>rpoC2</i>	12	27	0
<i>ndhF</i>	<i>waxy</i>	9	396	18
<i>phyB</i>	<i>rbcL</i>	4	4	4
<i>phyB</i>	<i>rpoC2</i>	7	1	0
<i>phyB</i>	<i>waxy</i>	3	6	6
<i>phyB</i>	<i>ITS</i>	8	9	9
<i>rbcL</i>	<i>rpoC2</i>	13	9	0
<i>rbcL</i>	<i>waxy</i>	7	35	0
<i>rpoC2</i>	<i>waxy</i>	1	1	1
<i>waxy</i>	<i>ITS</i>	8	18	0

MAAFs maximum-acyclic-agreement forests

given in Table 1. This data set was originally provided by the [Grass Phylogeny Working Group \(2001\)](#) and contains DNA sequences for the six genetic loci *ndhF*, *phyB*, *rbcL*, *rpoC2*, *waxy*, and *ITS*. For each locus, up to 65 taxa were sequenced and a maximum likelihood gene tree was reconstructed (for details, see [Bordewich et al. 2007](#) and references therein).

A natural way to decide whether there exists a minimal temporal hybridization network for two rooted binary phylogenetic X -trees \mathcal{T} and \mathcal{T}' is to apply TEMPORAL-HYBRID to all orderings of each maximum-acyclic-agreement forest for \mathcal{T} and \mathcal{T}' . However, since the number of maximum-acyclic-agreement forests can still be quite large (see Table 1), we next establish a quick test that significantly reduces the number of such forests that one needs to consider.

To describe the test, we need one further definition. Let \mathcal{F} be an acyclic-agreement forest for two rooted binary phylogenetic X -trees \mathcal{T} and \mathcal{T}' , and let \mathcal{S}_ρ be the root component of \mathcal{F} . If the roots of $\mathcal{T}(\mathcal{L}_\rho - \{\rho\})$ and $\mathcal{T}'(\mathcal{L}_\rho - \{\rho\})$ coincide with the original roots of \mathcal{T} and \mathcal{T}' , respectively, \mathcal{S}_ρ is said to be *proper*.

Proposition 4.4 *Let \mathcal{T} and \mathcal{T}' be two rooted binary phylogenetic X -trees, and let \mathcal{F} be a maximum-acyclic-agreement forest of \mathcal{T} and \mathcal{T}' with root component \mathcal{S}_ρ . Let \mathcal{O} be an ordering of \mathcal{F} . If \mathcal{S}_ρ is not proper, then TEMPORALHYBRID applied to $(\mathcal{T}, \mathcal{T}', \mathcal{O})$ returns a statement indicating that there is no temporal hybridization network on X that displays \mathcal{T} and \mathcal{T}' , and preserves \mathcal{O} . In particular, there is no minimal temporal hybridization network on X that displays \mathcal{T} and \mathcal{T}' , and preserves \mathcal{O} .*

Proof Suppose that \mathcal{S}_ρ is not proper, and consider TEMPORALHYBRID applied to $(\mathcal{T}, \mathcal{T}', \mathcal{O})$. If this application returns a hybridization network \mathcal{N} , then, as \mathcal{F} is maximum, \mathcal{N} has exactly $|\mathcal{F}| - 1$ hybridization vertices and so, at each iteration i , two (distinct) new arcs are used to adjoin \mathcal{S}_i to \mathcal{N}_{i-1} . Now, at some iteration i , the component \mathcal{S}_i gets adjoined to \mathcal{N}_{i-1} via an arc that is incident with a new vertex that

subdivides the arc leaving the vertex labeled ρ . But this means that wherever the second new arc is placed to adjoin \mathcal{S}_i to \mathcal{N}_{i-1} we contradict (ii) in Step 3. Thus, by Theorem 4.2, there is no temporal hybridization network on X that displays \mathcal{T} and \mathcal{T}' , and preserves \mathcal{O} . The proposition now follows.

By checking which maximum-acyclic-agreement forests for a given pair of gene trees have a proper root component, the number of such forests that can yield a minimal temporal hybridization network can be reduced significantly. For example, in reference to Table 1, the highest number of maximum-acyclic-agreement forests for a pair of trees is 2,268. However, as shown in the last column, none of these forests has a proper root component. Hence, for the first pair of gene trees (*ndhF* and *phyB*) in this table, there is no minimal temporal hybridization network that displays the two trees. Moreover, for the 12 analyzed gene tree pairs of the grass data set, at most 18 maximum-acyclic-agreement forests need to be checked in order to determine whether any associated ordering leads to a minimal temporal hybridization network for the gene trees under consideration. In general, we can check in time $O(rk! \cdot p(a))$ if a minimal temporal hybridization network exists, where r is the number of maximum-acyclic-agreement forests (with a proper root component) for a pair of trees, each such forest consists of $k + 1$ components, and $p(a)$ is the polynomial running time of the algorithm TEMPORALHYBRID. Note that $k!$ is an upper bound on the number of orderings that are associated with a maximum-acyclic-agreement forest containing $k + 1$ components. Furthermore, we are assuming here that we have the list of maximum-acyclic-agreement forests with a proper root component. Such a list can be found by using the recently implemented extended version (unpublished) of the fixed-parameter algorithm HYBRIDINTERLEAVE (Collins 2009; Collins et al. 2009). Despite the theoretical exponential time of calculating this list, the practical running times presented in Collins et al. (2009) essentially show that computing maximum-acyclic-agreement forests is remarkably quick for many biological instances.

5 Summary

In the first part of this paper, we showed that ADDTAXA—the decision problem associated with determining the minimum number of taxa to add to a reticulation network so that the resulting network has a temporal representation—is an NP-complete problem. However, in establishing the result, this determination comes down to finding the minimum number of vertices to delete so that the associated critical graph is acyclic. In practice, we expect this graph to be relatively small for many biological instances and thus even brute-force algorithms might be feasible.

In the second part of this paper, we presented the polynomial-time algorithm TEMPORALHYBRID. This algorithm takes as input two rooted binary phylogenetic X -trees \mathcal{T} and \mathcal{T}' and an ordering \mathcal{O} of an associated acyclic-agreement forest, and outputs a temporal hybridization network that displays \mathcal{T} and \mathcal{T}' and preserves \mathcal{O} , or the statement that no such network exists. As many biological studies consider the reconstruction of minimal hybridization networks to provide an indication of the significance of hybridization in evolution, we focused our attention to the potentially exponential-time task of finding a temporal hybridization network whose number of

hybridization vertices is minimized in the last section. By using a simple and quick check prior to the application of TEMPORALHYBRID, we showed that—applied to a grass data set—the number of maximum-acyclic-agreement forests that need to be considered for finding a minimal temporal hybridization network is significantly reduced and that most inferred minimal hybridization networks are not temporal.

Acknowledgments We thank two anonymous reviewers for their helpful comments. S.L. was supported by NSF grants SEI-BIO 0513910 and IIS-0803564, and the New Zealand Marsden Fund. C.S. thanks the New Zealand Marsden Fund for supporting this work. T.S. was funded by the Deutsche Forschungsgemeinschaft through the graduate program “Angewandte Algorithmische Mathematik” at the Munich University of Technology. All authors thank the Allan Wilson Centre for Molecular Ecology and Evolution for its support.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Baroni M, Grünwald S, Moulton V, Semple C (2005) Bounding the number of hybridization events for a consistent evolutionary history. *J Math Biol* 51:171–182
- Baroni M, Semple C, Steel M (2006) Hybrids in real time. *Syst Biol* 44:46–56
- Bordewich M, Semple C (2007) Computing the minimum number of hybridization events for a consistent evolutionary history. *Discret Appl Math* 155:914–928
- Bordewich M, Linz S, John K, Semple C (2007) A reduction algorithm for computing the hybridization number of two trees. *Evol Bioinform* 3:86–98
- Cardona G, Rossello F, Valiente G (2009) Comparison of tree-child phylogenetic networks. *IEEE/ACM Trans Comput Biol Bioinform* 6:552–569
- Chen J, Liu Y, Lu S, O’Sullivan B, Razgon I (2008) A fixed-parameter algorithm for the directed feedback vertex set problem. In: *Proceedings of the fortieth annual ACM symposium on theory of computing*, pp 177–186
- Collins J (2009) Rekernelisation algorithms in hybrid phylogenies. MSc Thesis, University of Canterbury, Christchurch, New Zealand
- Collins J, Linz S, Semple C (2009) Quantifying hybridization in realistic time (submitted)
- Grass Phylogeny Working Group (2001) Phylogeny and subfamilial classification of the grasses *Poaceae*. *Ann Mo Botanical Gard* 88:373–457
- Hein J, Jiang T, Wang L, Zhang K (1996) On the complexity of comparing evolutionary trees. *Discret Appl Math* 71:153–169
- Jin G, Nakhleh L, Snir S, Tuller T (2007) Efficient parsimony-based methods for phylogenetic network reconstruction. *Bioinformatics* 23:e123–e128
- Karp RM (1972) Reducibility among combinatorial problems. In: *Complexity of computer computations*. Plenum Press, New York, pp 85–103
- Maddison W (1997) Gene trees in species trees. *Syst Biol* 46:523–536
- Makarenkov V, Kevorkov D, Legendre P (2006) Phylogenetic network construction approaches. In: *Applied mycology and biotechnology*. International Elsevier Series 6, Bioinformatics. Elsevier, Amsterdam, pp 61–97
- Mallet J (2005) Hybridization as an invasion of the genome. *Trends Ecol Evol* 20:229–237
- Martinsen G, Whitham T, Turek R, Keim P (2001) Hybrid populations selectively filter gene introgression between species. *Evolution* 55:1325–1335
- Moret BME, Nakhleh L, Warnow T, Linder CR, Tholse A, Padohina A, Sun J, Timme R (2004) Phylogenetic networks: modeling, reconstructibility, and accuracy. *Trans Comput Biol Bioinform* 1:13–23
- Ochman H, Lawrence J, Groisman E (2000) Lateral gene transfer and the nature of bacterial innovation. *Nature* 405:299–304
- Wolf D, Takebayashi N, Rieseberg L (2001) Predicting the risk of extinction through hybridization. *Conserv Biol* 15:1039–1053