



On the Complexity of Constructing Evolutionary Trees

LESZEK GAŚNIENIEC

Department of Computer Science, University of Liverpool, Peach Street, L69 7ZF, Liverpool, UK

leszek@csc.liv.ac.uk

JESPER JANSSON

Department of Computer Science, Lund Institute of Technology, Box 118, 221 00 Lund, Sweden

Jesper.Jansson@dna.lth.se

ANDRZEJ LINGAS

ANNA ÖSTLIN

Department of Computer Science, Lund University, Box 118, 221 00 Lund, Sweden

Andrzej.Lingas@dna.lth.se

Anna.Ostlin@dna.lth.se

Received November 25, 1995; Revised March 30, 1996

Abstract. In this paper we study a few important tree optimization problems with applications to computational biology. These problems ask for trees that are consistent with an as large part of the given data as possible. We show that the maximum homeomorphic agreement subtree problem cannot be approximated within a factor of N^ϵ , where N is the input size, for any $0 \leq \epsilon < \frac{1}{9}$ in polynomial time unless $P = NP$, even if all the given trees are of height 2. On the other hand, we present an $O(N \log N)$ -time heuristic for the restriction of this problem to instances with $O(1)$ trees of height $O(1)$ yielding solutions within a constant factor of the optimum. We prove that the maximum inferred consensus tree problem is NP-complete, and provide a simple, fast heuristic for it yielding solutions within one third of the optimum. We also present a more specialized polynomial-time heuristic for the maximum inferred local consensus tree problem.

Keywords: algorithm, time complexity, evolutionary tree, homeomorphism, consensus tree

1. Introduction

An evolutionary tree models how different species in a given set have evolved. The leaves in an evolutionary tree correspond to species and internal nodes represent the species' ancestors.

The problem of constructing a reliable evolutionary tree has been studied extensively (Farach et al., 1995; Farach and Thorup, 1994a, 1994b; Hein et al., 1996; Henzinger et al., 1996; Kannan et al., 1995; Kao et al., 1997; Keselman and Amir, 1994; Lam et al., 1996; Phillips and Warnow, 1996; Steel and Warnow, 1993). There are many different approaches, depending among other things on what kind of data that is available. Therefore, various versions of this problem arise in, for example, computational biology when one wants to find out how different species are related, and comparative linguistics, where it is central to find out how different languages have evolved. In this paper, we look at some of these problems.

Given a set of alternative evolutionary trees describing possible evolutions for a fixed set of species, one might want to identify a subtree contained within every given tree such

that the number of leaves labeled by species is maximized. This problem is known as *the maximum homeomorphic agreement subtree problem* (MHT) (Keselman and Amir, 1994). More formally, it is defined as follows. Given k rooted trees T_1, T_2, \dots, T_k , each with n leaves labeled distinctly with elements chosen from a set A of cardinality n , find a maximum cardinality subset B of A such that the minimal homeomorphic subtrees of T_1, T_2, \dots, T_k (i.e., with all degree 2 nodes except for the root contracted) containing exactly the leaves labeled by B are isomorphic. To measure the input size of an instance of MHT, we let N denote the total number of nodes contained in the given trees. MHT restricted to instances with two trees is frequently called MAST; algorithms for MAST have been developed since 1985 (Finden and Gordon, 1985). It has been shown to be solvable in polynomial time, both for rooted trees (Farach and Thorup, 1994a, 1994b) and for UMAST, a variant of MAST with unrooted trees (Farach and Thorup, 1994a; Kao et al., 1997; Lam et al., 1996; Steel and Warnow, 1993). In practice, however, the number of trees is often much larger than two (Keselman and Amir, 1994). For the special case of MHT in which at least one of the given trees has bounded degree, there exist polynomial-time algorithms (Farach et al., 1995; Keselman and Amir, 1994). In contrast, MHT is known to be NP-complete even for instances with three trees of unbounded degree (Keselman and Amir, 1994). The first non-approximability result for MHT was published in (Hein et al., 1996). It states that for three trees with unbounded degree, MHT cannot be approximated within ratio $2^{\log^\delta n}$ in polynomial time for any $\delta < 1$ unless $\text{NP} \subseteq \text{DTIME}[2^{\text{poly} \log n}]$. Here we prove that, unless $\text{P} = \text{NP}$, MHT cannot be approximated within a factor of N^ϵ , for any $0 \leq \epsilon < \frac{1}{9}$ in polynomial time even for instances containing only trees of height 2; see Section 2. On the other hand, in Section 3, we show that MHT for instances with $O(1)$ number of trees of height $O(1)$ can be approximated within a constant factor in time $O(N \log N)$. Similar results also hold for the unrooted version of MHT which is at least as hard as MHT (this can be seen by an argument analogous to that in (Farach and Thorup, 1994a) for MAST and UMAST).

Usually MHT instances do not admit a solution containing all the members of the species set. Therefore, in some applications, other methods may be preferred. One alternative approach is to attempt to construct an evolutionary tree from a set of constraints that relate the species to each other. Already during the early eighties, Aho et al. (1981) studied the problem of inferring a tree from constraints on its lowest common ancestors in the context of relational data bases. They defined it as follows: Given a set of constraints of the form $\{i, j\} < \{k, l\}$, where $\{i, j, k, l\} \subset \{1, 2, \dots, n\}$, if possible construct a tree on the set of leaves $\{1, 2, \dots, n\}$ such that for each constraint of the aforementioned form, the lowest common ancestor of i and j is a proper descendant of the lowest common ancestor of k and l . Aho et al. showed how to decide whether an instance of this problem admits a solution, and if so, how to construct it, both in time $O(mn \log n)$, where m denotes the number of constraints.

Recently, many authors have studied the related problem of constructing the so-called consensus tree or local consensus tree (Henzinger et al., 1996; Kannan et al., 1995; Phillips and Warnow, 1996). For a set of binary rooted trees $\{T_1, T_2, \dots, T_k\}$, each one leaf-labeled by a subset $L(T_i)$ of $\{1, 2, \dots, n\}$, the consensus tree problem asks whether or not there is a tree T such that for $i = 1, 2, \dots, k$, T_i is homeomorphic to the subtree of T induced

by the nodes in $L(T_i)$ and their ancestors. If the input trees are of constant size it is termed the local consensus tree problem. A constraint of the form $\{i, j\} < \{i, k\}$ (denoted $(\{i, j\}, k)$ for short) is easily seen to be equivalent to the constraint imposed by a full binary tree on the leaves i, j, k in the local consensus tree problem. For this reason, we shall call the tree inferring problem posed in (Aho et al., 1981) *the inferred consensus tree problem*.

Unfortunately, it is often impossible to construct an exact consensus tree. This creates a need for an optimization version of the inferred consensus tree problem whose objective is to find a consensus tree for an as large as possible subset of the input set of constraints of the form $\{i, j\} < \{k, l\}$. For brevity, we term this optimization problem *the maximum inferred consensus tree problem* (MICT for short). We also distinguish the restricted case of MICT where all constraints are of the form $(\{i, j\}, k)$ and call it *the maximum inferred local consensus tree problem* (MILCT).

In Section 4 we provide an NP-completeness proof for MICT. Section 5 contains a simple $O((n + m) \log n)$ -time heuristic for MICT yielding solutions within one third of the optimum and a more involved polynomial-time heuristic for MILCT. Both heuristics work equally well for the weighted versions of MICT and/or MILCT where the objective is to find a consensus tree for a subset of the input constraints of maximum total weight.

2. MHT is hard to approximate

Our main result in this section is the following theorem.

Theorem 1. *For any $0 \leq \epsilon < \frac{1}{9}$, MHT, even if restricted to trees of height 2, cannot be approximated within a factor of N^ϵ in polynomial time, unless $P = NP$.*

Proof: First, we describe a reduction from the maximum independent set problem to MHT. Next, we show that if MHT can be approximated within a factor of N^ϵ in polynomial time then the problem of finding a maximum independent set in a graph with l nodes can be approximated within a factor of $l^{3\epsilon+o(1)}$. Finally, we apply known results about the inapproximability of the maximum independent set problem to get our result. In part, our reduction can be seen as a generalization of the reduction of three-dimensional perfect matching to MHT restricted to instances with three trees used in (Keselman and Amir, 1994).

Let $G = (V, E)$ be a graph where $V = \{v_1, \dots, v_l\}$ and $E = \{e_1, \dots, e_k\}$ with $k > 1$. Construct k rooted trees T_1, \dots, T_k on $l + q$ labeled leaves that contain all the adjacency information about the nodes of G as follows. For each edge $e_i = (v_a, v_b) \in E$, build a rooted tree T_i on the set of leaves labeled by $w_1, \dots, w_l, w_{l+1}, \dots, w_{l+q}$. Let the root r_i of T_i be the parent of $(l - 1) + q$ children, where the first child (“the non-leaf child”) is a node with two children leaves labeled w_a and w_b , and the remaining children of r_i are leaves labeled by the elements in $\{w_j \mid 1 \leq j \leq l + q \text{ and } j \notin \{a, b\}\}$. Thus, r_i has exactly one pair of grandchildren, and we write $GC(T_i) = \{w_a, w_b\}$.

Now, let T be a maximum homeomorphic agreement subtree of T_1, \dots, T_k . We choose q large enough to guarantee that each of the roots r_1, \dots, r_k will correspond to the root r of T . Actually, $q = 2$ is sufficient. (To see this, assume that the non-leaf child of r_i turned out to be the root for some i . By the construction above, all non-leaf children have two leaf children, so the number of leaves in this agreement subtree can be no larger than two. But we can always find an agreement subtree with three leaves by selecting r_i as root and including w_{l+1} and w_{l+2} in addition to the path from the root to a fixed leaf w_j , where $1 \leq j \leq l$.)

T has no non-leaf children because if it did, then there would exist some x and y such that for each i , where $1 \leq i \leq k$, $GC(T_i)$ would be equal to $\{w_x, w_y\}$. Consequently, G would have only one edge which contradicts $k > 1$. The children of T are $m + q (= m + 2)$ leaves labeled $w_{\mu_1}, w_{\mu_2}, \dots, w_{\mu_m}, w_{l+1}, w_{l+2}$. If v_a is adjacent to v_b in G then at most one of w_a and w_b can be a child of T . Otherwise, $GC(T_i)$ wouldn't be equal to $\{w_a, w_b\}$ for any T_i . Consequently, $e_i \neq (v_a, v_b)$ would hold for all i , contradicting the adjacency of v_a and v_b in G . Thus, the nodes $v_{\mu_1}, v_{\mu_2}, \dots, v_{\mu_m}$ form an independent set in G . Conversely, given an independent set I of nodes in G , we can easily construct an agreement subtree T_i in the form of a rooted tree with $|I| + 2$ leaves uniquely labeled with w_j , where $v_j \in I$, and w_{l+1}, w_{l+2} . By the maximality of T , m equals the cardinality of the maximum independent set of G . Thus, an algorithm for MHT would immediately imply an algorithm for the maximum independent set problem. See figures 1–3 for an example of the reduction.

The total size N of the trees T_1, \dots, T_k is $k \cdot O(l) = O(l^3) = l^{3+o(1)}$. Clearly, they can be constructed from G in polynomial time. Also, note that they are of height 2. Below we will only consider approximations that can be carried out in polynomial time. If MHT could be approximated within a factor of N^ϵ , then $\frac{OPT+2}{s+2} \leq N^\epsilon$, where $OPT + 2$ refers to the number of leaves in an optimal solution for a given instance of MHT and $s + 2$ is the number of leaves in its corresponding, approximative solution. For $s \geq 1$, it follows that $\frac{OPT}{s} \leq 3 \cdot \frac{OPT+2}{s+2} \leq 3N^\epsilon = l^{3\epsilon+o(1)}$, which would imply that the problem of finding a maximum independent set in a graph could be approximated within a factor of $l^{3\epsilon+o(1)}$. However, Håstad (1996) proved that this problem isn't approximable within $l^{1/3-\delta}$ for any $\delta > 0$, unless $P = NP$. Hence, if $P \neq NP$, MHT cannot be approximated within a factor of N^ϵ for any $0 \leq \epsilon \leq \frac{1}{9} - o(1)$. Finally, since $\frac{1}{9} - o(1)$ can be made arbitrarily close to $\frac{1}{9}$ by choosing N large enough, there exist instances of MHT which cannot be approximated within a factor of N^ϵ for any constant $0 \leq \epsilon < \frac{1}{9}$ in polynomial time (unless $P = NP$). \square




Figure 1. An instance of the maximum independent set problem with $l = 7$ and $k = 8$.




Figure 2. The trees T_i corresponding to the graph in figure 1.




Figure 3. The maximum homeomorphic agreement subtree of T_1, \dots, T_8 tells us that $\{v_2, v_3, v_4, v_7\}$ is a maximum independent set of the graph in figure 1.

3. Approximations of MHT with $O(1)$ trees of height $O(1)$

We know that MHT is hard to approximate, both for instances with three trees (Hein et al., 1996) and for instances with an arbitrary number of trees of height 2 or more by Theorem 1. The natural question arises whether or not MHT for instances with a bounded number of trees, each one of bounded height, can be tightly approximated in polynomial time. The following result, together with Theorem 1, yields a characterization of the approximability of MHT restricted to instances with trees of $O(1)$ height.

Theorem 2. *MHT restricted to instances with k trees of height not exceeding h can be approximated within a factor of k^h in time $O(n \log n)$.*

To begin the proof of Theorem 2, we need to introduce the following notation. For a tree T , $V(T)$ stands for the set of nodes of T . Let v be a node of a rooted tree T . The minimal subtree of T rooted at v , including v and all its descendants is denoted by T_v . $L(T_v)$ stands for the set of labels of the leaves in T_v . The set of children of v in T is denoted by $C(v)$. Furthermore, by a k -partite hypergraph H we shall mean a pair $(V_1 \cup \dots \cup V_k, E)$ where V_1 through V_k are pairwise disjoint sets and E is a subset of $V_1 \times \dots \times V_k$. The elements of $V_1 \cup \dots \cup V_k$ are called the *nodes* of H whereas the elements of E are called the *edges* of H . A *matching* of H is a subset of E in which no pair of edges includes a common node.

Let T_1, \dots, T_k be the input trees. For (v_1, \dots, v_k) , where $v_i \in V(T_i)$ for $i = 1, \dots, k$, let $Mht(v_1, \dots, v_k)$ denote the maximum size of an agreement subtree of the trees T_1, \dots, T_k restricted to $B = L((T_1)_{v_1}) \cap \dots \cap L((T_k)_{v_k})$. We can view $Mht(v_1, \dots, v_k)$ as the solution of MHT for $(T_1)_{v_1}, \dots, (T_k)_{v_k}$. Next, let $H(v_1, \dots, v_k)$ denote the k -partite hypergraph $(C(v_1) \cup \dots \cup C(v_k), C(v_1) \times \dots \times C(v_k))$ whose edges (w_1, \dots, w_k) have weight $Mht(w_1, \dots, w_k)$. Finally, let $Match(v_1, \dots, v_k)$ be the maximum weight of a matching in hypergraph $H(v_1, \dots, v_k)$ and $Diag(v_1, \dots, v_k) = \max\{Mht(w_1, \dots, w_k) \mid (w_1, \dots, w_k) \in (\{v_1\} \cup C(v_1)) \times \dots \times (\{v_k\} \cup C(v_k)) - \{(v_1, \dots, v_k)\}\}$.

Intuitively, in the final agreement subtree of $(T_1)_{v_1}, \dots, (T_k)_{v_k}$ either the roots of the trees, i.e., v_1 through v_k , are matched together which forces their children to be optimally matched together (*Match*), or only some of the roots are matched together with some children of the remaining roots (*Diag*). This yields the following lemma which is a straightforward generalization of the basic lemma in the dynamic programming approach to MAST in (Steel and Warnow, 1993) (see also Farach and Thorup, 1994a).

Lemma 1. *For any (v_1, \dots, v_k) , where $v_i \in V(T_i)$ for $i = 1, \dots, k$, if at least one of the v_i 's is a leaf then $Mht(v_1, \dots, v_k) = |L((T_1)_{v_1}) \cap \dots \cap L((T_k)_{v_k})|$ else $Mht(v_1, \dots, v_k) = \max\{Match(v_1, \dots, v_k), Diag(v_1, \dots, v_k)\}$.*

It is easy to see that the recursive computation of $Mht(v_1, \dots, v_k)$ for $(v_1, \dots, v_k) \in V(T_1) \times \dots \times V(T_k)$ used in Lemma 1 can be bottom-up ordered by $Hs(v_1, \dots, v_k) = \sum_{i=1}^k height((T_i)_{v_i})$. Hence, we have the following algorithm for MHT.

Algorithm 1.

1. *input* T_1, \dots, T_k
2. *for each* $(v_1, \dots, v_k) \in V(T_1) \times \dots \times V(T_k)$, in increasing order of $Hs(v_1, \dots, v_k)$
do compute $Mht(v_1, \dots, v_k)$ by using the expression in Lemma 1.
3. *output* $Mht(r_1, \dots, r_k)$ where r_i is the root of T_i , for $i = 1, \dots, k$.

It is hard to compute the exact value of $Match(v_1, \dots, v_k)$ in the expression of Lemma 1 since the problem of computing maximum matching in a 3-partite hypergraph is NP-complete (Papadimitriou, 1994). For this reason, we shall rely on a greedy method for approximating $Match(v_1, \dots, v_k)$ yielding an approximation of $Mht(v_1, \dots, v_k)$. The greedy method consists of repeatedly picking the heaviest edge e and removing all edges overlapping e . It can be implemented easily using a priority queue. e can overlap with at most k edges in an optimum solution, and since their total weight $\leq k \cdot weight(e)$, we obtain:

Lemma 2. *Let $H = (V, E)$ be a k -partite hypergraph on m edges with positive integer weights. A matching in H of total weight within a factor k of the maximum can be constructed in a greedy fashion in time $O(k|E| + |V| + m \log m)$.*

Interestingly, in the unweighted case there are known (much slower, but still) polynomial-time heuristics yielding solutions within almost $\frac{k}{2}$ of the optimum (Hurkens and Schrijver, 1989).

By combining the scheme of Algorithm 1 with the greedy method for approximating $Match(v_1, \dots, v_k)$, we obtain the following lemma, yielding Theorem 2.

Lemma 3. *For all $(v_1, \dots, v_k) \in V(T_1) \times \dots \times V(T_k)$, we can approximate $Mht(v_1, \dots, v_k)$ within a factor of k^h , where $h = \max\{height((T_i)_{v_i}) \mid 1 \leq i \leq k\}$ in time $O(n \log n)$.*

Proof: For $(v_1, \dots, v_k) \in V(T_1) \times \dots \times V(T_k)$, let $s(v_1, \dots, v_k)$ denote the size of the intersections $L((T_1)_{v_1}) \cap \dots \cap L((T_k)_{v_k})$. Clearly, we have $Mht(v_1, \dots, v_k) \leq s(v_1, \dots, v_k)$, and in particular if one of the v_i 's is a leaf then $Mht(v_1, \dots, v_k) = s(v_1, \dots, v_k)$. For a leaf label j , we determine all k -tuples (v_1, \dots, v_k) for which $j \in L((T_1)_{v_1}) \cap \dots \cap L((T_k)_{v_k})$ by finding, in each T_i , $i = 1, \dots, k$, the nodes on the path of length $\leq h$ from the leaf labeled j to the root. It follows that the number of these tuples is $(h+1)^k$. Consequently, the set L of all k -tuples for which $s(v_1, \dots, v_k) > 0$ has size not exceeding $n(h+1)^k$. To list L efficiently, we sort the pointers to leaves in T_1 through T_k by the leaf labels. Such a sorted list of pointers can be produced in time $O(|V(T_1)| + \dots + |V(T_k)|)$. Using it, we can generate L by finding appropriate tree paths in time $O(|V(T_1)| + \dots + |V(T_k)| + (h+1)^k)$.

For the k -tuples (v_1, \dots, v_k) in set L that include at least one leaf, we clearly have $s(v_1, \dots, v_k) = 1$ and $Mht(v_1, \dots, v_k) = 1$. To compute approximations of $Mht(v_1, \dots, v_k)$ for the remaining k -tuples in L , we build a balanced search tree S_L for L , with respect to the lexicographic order of k -tuples in $V(T_1) \times \dots \times V(T_k)$, in time $O(|L| \log |L|)$. Next, we follow the scheme of Algorithm 1 using the greedy method to approximate

$Match(v_1, \dots, v_k)$ in the hypergraph $H_L(v_1, \dots, v_k)$ which is the hypergraph $H(v_1, \dots, v_k)$ defined in Lemma 1 restricted to edges in L .

Each k -tuple $(w_1, \dots, w_k) \in L$ occurs at most once as an edge in the hypergraphs $H_L(v_1, \dots, v_k)$ for $(v_1, \dots, v_k) \in L$ (only when $w_i \in C(v_i)$, for $i = 1, \dots, k$). Hence, the hypergraphs $H_L(v_1, \dots, v_k)$ for $(w_1, \dots, w_k) \in L$, have no more than $|L|$ edges totally and can be constructed (without weights) by scanning L and using S_L in total time $O(|L| \log |L|)$. Clearly, each $H_L(v_1, \dots, v_k)$ has at most $s(v_1, \dots, v_k)$ edges with positive weights. For each of its edges (w_1, \dots, w_k) , we have $Hs(w_1, \dots, w_k) = Hs(v_1, \dots, v_k) - k$ and $\max\{\text{height}(w_i) \mid 1 \leq i \leq k\} = h - 1$. Hence, we may inductively assume that we have already k^{h-1} approximations of $Mht(w_1, \dots, w_k)$, i.e., of the weights of (w_1, \dots, w_k) in the hypergraph. Consequently, we obtain an approximation of $Match(v_1, \dots, v_k)$ within a factor of k^h by applying the greedy method. Due to Lemma 2, the total time complexity of the greedy method is bounded by $O(k|L| + (\sum_{(v_1, \dots, v_k) \in L} s(v_1, \dots, v_k)) \log n)$. By induction on $Hs(v_1, \dots, v_k)$, we also obtain an approximation of $Diag(v_1, \dots, v_k)$ within a factor of k^h by considering solely k -tuples (w_1, \dots, w_k) in $L \cap ((\{v_1\} \cup C(v_1)) \times \dots \times (\{v_k\} \cup C(v_k)) - \{(v_1, \dots, v_k)\})$. Each $(w_1, \dots, w_k) \in L$ can contribute to the value of $Diag(v_1, \dots, v_k)$ for at most $2^k - 1$ k -tuples $(v_1, \dots, v_k) \in L$. Hence, the total size of the subsets of L contributing to $Diag(v_1, \dots, v_k)$ over all $(v_1, \dots, v_k) \in L$, and consequently the total cost of finding maxima of Mht -approximations over these subsets, is $O(2^k |L|)$. We can build these subsets, again by scanning L and using S_L , in total time $O(2^k |L| \log |L|)$.

Each of the trees T_1 through T_k has size not exceeding $2n$ by its binarity. Hence, by $|L| \leq (h+1)^k n$, $\sum_{(v_1, \dots, v_k) \in L} s(v_1, \dots, v_k) \leq (h+1)^k n$, $h = O(1)$, $k = O(1)$, and straightforward calculations, we obtain the $O(n \log n)$ bound. \square

4. MICT is NP-complete

The problem of deciding whether or not a 3-partite hypergraph (V, E) has a perfect matching (3PM), i.e., if V is covered by a subset of pairwise disjoint edges in E , is known to be NP-complete (Papadimitriou, 1994). To show the NP-completeness of MICT, we provide a reduction of 3PM to MICT.

Let $H = (V, E)$ be a 3-partite hypergraph and k a parameter that will be specified later on. We let each vertex in V label one leaf. Also, for each edge $e \in E$, we introduce $k+2$ leaves labeled e_i , where $i = 0, \dots, k+1$. Let C be the minimal set of constraints satisfying:

1. for each $e, f \in E$ with $e \neq f$: the constraints $(\{e_i, e_l\}, f_j) \in C$, where $i = 0, 1, l = 2, \dots, k+1$, and $j = 2, \dots, k+1$.
2. for each $e = (a, b, c) \in E$: the three constraints $\{a, b\} < \{e_0, e_1\}$, $\{a, c\} < \{e_0, e_1\}$, and $\{b, c\} < \{e_0, e_1\} \in C$.

Thus, C consists of $2k^2(|E|^2 - |E|)$ constraints of the first type and $3|E|$ constraints of the second type.

To characterize consensus trees for large subsets of C , we need the following definitions.

Definition 1. In a rooted tree T , the lowest common ancestor of a sequence of nodes v_1, \dots, v_m will be denoted by $\text{lca}(v_1, \dots, v_m)$. Furthermore, the path from a node v to the root of T will be denoted by $R(v)$. The subtree of T induced by a sequence of nodes v_1, \dots, v_m is the smallest subtree of T including the paths $R(v_i)$, $i = 1, \dots, m$.

Definition 2. The full binary tree on four leaves a, b, c, d , where $\text{lca}(a, b)$ and $\text{lca}(c, d)$ form the intermediate level, will be denoted by $B_4(a, b, c, d)$.

Lemma 4. *If T is a consensus tree for at least $|C| - k^2 + 1$ constraints in C , then for each $e, f \in E$ with $e \neq f$, the subtree induced of T by $\{e_0, e_1, f_0, f_1\}$ is homeomorphic to $B_4(e_0, e_1, f_0, f_1)$.*

Proof: By the assumption on the number of constraints satisfied by T , for each $e, f \in E$ with $e \neq f$, there are indices $l, j \in \{2, \dots, k + 1\}$ such that for $i = 0, 1$, the constraints $(\{e_i, e_l\}, f_j)$, $(\{f_i, f_j\}, e_l)$ are satisfied by T .

By $(\{e_0, e_l\}, f_j)$ and $(\{e_1, e_l\}, f_j)$, the path $R(\text{lca}(e_0, e_1, e_l))$ cannot be included in the path $R(f_j)$. Thus, $R(\text{lca}(e_0, e_1, e_l)) \not\subseteq R(\text{lca}(f_0, f_1, f_j))$. Similarly, by $(\{f_0, f_j\}, e_l)$ and $(\{f_1, f_j\}, e_l)$, we have $R(\text{lca}(f_0, f_1, f_j)) \not\subseteq R(\text{lca}(e_0, e_1, e_l))$. This means that the paths from $\text{lca}(e_0, e_1, e_l)$ to $\text{lca}(e_0, e_1, e_l, f_0, f_1, f_j)$ as well as from $\text{lca}(f_0, f_1, f_j)$ to $\text{lca}(e_0, e_1, e_l, f_0, f_1, f_j)$ must be edge-disjoint. \square

Corollary 1. *Let T be a consensus tree for at least $|C| - k^2 + 1$ constraints in C . For each node $a \in V$ and two different edges $e, f \in E$, if T satisfies a constraint of the form $\{a, \cdot\} < \{e_0, e_1\}$ then T cannot satisfy any constraints of the form $\{a, \cdot\} < \{f_0, f_1\}$.*

Lemma 5. *Let $k > \sqrt{3|E| - |V|}$. The hypergraph H has a perfect matching iff there is a consensus tree for a subset of $2k^2(|E|^2 - |E|) + |V|$ constraints in C .*

Proof: Suppose first that H has a perfect matching M . We can construct a consensus tree T satisfying at least $2k^2(|E|^2 - |E|) + |V|$ of the constraints in C as follows. The root of T has $|E|$ children which are in one-to-one correspondence with the edges in E . For every $e \in E$, a subtree rooted in the corresponding child has as children the leaves e_0, e_1, \dots, e_{k+1} . Furthermore, if $e = \{a, b, c\}$ is in M , then the subtree has another child which in turn is the parent of the leaves labeled a, b, c .

Suppose in turn that there is a consensus tree T satisfying $2k^2(|E|^2 - |E|) + |V|$ constraints in C . The total number of constraints in C is $2k^2(|E|^2 - |E|) + 3|E|$. It follows by $k > \sqrt{3|E| - |V|}$ that T satisfies at least $|C| - k^2 + 1$ constraints. Thus, by Corollary 1, for each node $a \in V$, there is at most one edge $e \in E$ such that some constraint of the form $\{a, \cdot\} < \{e_0, e_1\}$ is satisfied by T . On the other hand, for a given node a and a given edge e , at most two constraints of the form $\{a, \cdot\} < \{e_0, e_1\}$ can be satisfied by T by the construction of C . Consequently, V can be partitioned into three disjoint subsets V_r , $r = 0, 1, 2$, respectively consisting of nodes $a \in V$ for which T satisfies r constraints of the form $\{a, \cdot\} < \{\cdot, \cdot\}$. At most $|V_2| + \frac{|V_1|}{2}$ constraints of the form $\{\cdot, \cdot\} < \{\cdot, \cdot\}$ are satisfied by T , so since there are only $2k^2(|E|^2 - |E|)$ constraints of the form $(\{\cdot, \cdot\}, \cdot)$ in C , we conclude that V_2 has to be as large as possible, i.e., $V_2 = V$. It follows that for each edge $e \in E$, if a

constraint of the form $\{\cdot, \cdot\} < \{e_0, e_1\}$ is satisfied by T , then all the three constraints of this form are satisfied by T . Hence, H has a perfect matching. \square

The construction of C for k equal, say, to $|E|$ can easily be done in polynomial time. Hence, MICT is NP-hard by the NP-completeness of 3PM and Lemma 5. The membership of MICT in NP is obvious.

Theorem 3. *MICT is NP-complete.*

5. Approximation heuristics for MICT

Our heuristics in fact work for the generalization of MICT where with each input constraint c a positive weight $w(c)$ is associated, and the objective is to construct a consensus tree for a subset of constraints of maximum total weight.

5.1. Heuristic 1

For a constraint $\{i, j\} < \{k, l\}$, where all the leaves are different, k and l are said to have an *upper occurrence* in the constraint, and i and j are said to have a *lower occurrence* in the constraint. For a constraint $\{i, j\} < \{i, k\}$, where i, j, k are different, i and j are said to have a lower occurrence in the constraint and k is said to have an upper occurrence in the constraint. The *total weight* of upper (or lower) occurrences for a leaf l is equal to the sum of the weights of all constraints in which l has upper (or lower) occurrences.

Lemma 6. *For any instance of MICT, the sum of all leaves' total weights of upper occurrences is at least one third (one half if all constraints contain four leaves) of the sum of all leaves' total weight of upper and lower occurrences.*

Heuristic 1.

input: a set C of m weighted constraints on leaves 1 through n ;

output: a consensus tree T for a subset of C whose weight is at least one third (one half if all constraints contain four leaves) of the total weight of the constraints in C ;

1. $LEFT \leftarrow C$;
 $LEAVES \leftarrow \{1, \dots, n\}$;
 $T \leftarrow \{v\}$;
2. if $LEFT = \emptyset$ then extend T by adding $|LEAVES|$ children to v , label them uniquely with elements in $LEAVES$, and return T ;
3. pick a leaf y in $LEAVES$ which achieves the maximum ratio between the total weight of y 's upper occurrences and the total weight of y 's lower occurrences in the constraints in $LEFT$;
4. set Y to the set of constraints in $LEFT$ which contains y ;
5. $LEFT \leftarrow LEFT \setminus Y$;

6. $LEAVES \leftarrow LEAVES \setminus \{y\}$;
7. extend T by adding two children to v ; label the first child by y ; set v to the second child;
8. go to 2

Theorem 4. *Heuristic 1 constructs a consensus tree for a subset of the input set of constraints C , whose total weight is at least one third (one half if all constraints contain four leaves) of the total weight of C , in time $O((m + n) \log n)$.*

Proof: By Lemma 6 and the choice of y , the ratio between the total weight of upper occurrences and lower occurrences of y in the constraints in $LEFT$ is at least one third. All the constraints in Y in which y has an upper occurrence are satisfied by T by the construction of T .

To implement Steps 3 and 6 efficiently, we arrange $LEAVES$ in a priority queue partially ordered by the ratio between the total weight of their upper and lower occurrences in constraints in $LEFT$. All the priority queue operations, i.e., creating the priority queue, picking the y 's, updating the priority queue after Step 5, take a total of $O((n + m) \log n)$ time.

To implement Steps 4 and 5, we lexicographically sort C four times according to four cyclic permutations of the four leaves in each constraint. For $i = 1, \dots, 4$, the i th permutation puts the i th leaf as the first, the $i + 1$ st (in the cyclic order) as the second etc. Next, four search trees are built on the basis of the sorted lists. Using the search trees, we can find Y in $LEFT$ and remove it from $LEFT$ in time $O(|Y| \log n)$. We conclude that Steps 4 and 5 totally take time $O((m + n) \log n)$ (inclusive the preprocessing). \square

The absolute factors of one third and half respectively provided by Heuristic 1 are worst-case optimal. For example, any consensus tree can satisfy at most one constraint from each consecutive triple of constraints in a sequence $(\{a_i, b_i\}, c_i), (\{b_i, c_i\}, a_i), (\{c_i, a_i\}, b_i), i = 1, \dots, k$. In case all constraints contain four leaves, the sequence $\{a_i, b_i\} < \{c_i, d_i\}, \{c_i, d_i\} < \{a_i, b_i\}, i = 1, \dots, k$, causes the lower bound $\frac{1}{2}$.

The consensus tree produced by Heuristic 1 has the form of a linear chain with singular leaves pending, where only the last chain node can have larger degree. It is easy to slightly modify Heuristic 1 to output a subset of the input constraints (a priori) satisfied by the tree. A minimum height consensus tree for at least one third of the input constraints is then obtained in time $O(mn \log n)$ by running the algorithm of Aho et al. (1981) for the inferred consensus tree problem on this set.

In case the minimum number of constraints necessary to delete in order to build a consensus tree for the remaining part is very small, and the number m of constraints relative to the number of leaves is high (it is always $O(n^4)$), an approach different from that of Heuristic 1 might be more useful.

5.2. Heuristic 2

Heuristic 2 for MILCT simply mimics the algorithm of Aho et al. (1981) for the inferred consensus tree problem restricted to constraints of the form $(\{i, j\}, k)$. Their basic idea

is simple. The input set of leaves $1, 2, \dots, n$ is partitioned into a minimal set of blocks satisfying the following requirement:

(*) If $(\{i, j\}, k)$ is a constraint then i and j are in the same block.

Now, if the number of blocks in the minimal set is at least two, the algorithm of Aho et al. creates the consensus tree by connecting the roots of the consensus trees recursively computed for the respective blocks with a common parent root node. Otherwise, the number is one, and it returns a null consensus tree.

For a subset S of leaves, let $G(S)$ denote the auxiliary graph on S where the edges are induced by the requirement (*), and their weights are equal to the total weight of the constraints inducing them.

Whenever the algorithm of Aho et al. is stuck at a non-divisible subset S of the set of leaves and has to return a null tree, Heuristic 2 simply finds a minimum weight edge cut of the auxiliary graph $G(S)$ (with respect to the current set of constraints). Next, the edges of the min-cut are deleted from $G(S)$ and the connected components of $G(S)$ are computed. Consequently, the constraints corresponding to the edges of the min-cut are also deleted. Finally, the approximation consensus trees for the connected components are recursively computed and connected by a common parent node.

Using recent dynamic data structures for graph connectivity, Henzinger et al. gave efficient implementations of the algorithm of Aho et al. restricted to constraints of the form $(\{i, j\}, k)$ (Henzinger et al., 1996). Their randomized implementation takes $O(m \log^3 n)$ expected time. They use the undirected graph U and the directed graph D defined as follows.

- $U = (V, E)$ with V equal to the set $\{1, 2, \dots, n\}$ of leaves and where for each constraint $(\{a, b\}, c)$ in C edges $\{a, b\}$ and $\{b, c\}$ are in E .
- $D = (V', A)$ where for each constraint $(\{a, b\}, c)$ in C nodes $\{a, b\}$ and $\{b, c\}$ are in V' and $\{a, b\} \rightarrow \{b, c\}$ is in A .

At the beginning the graph U is colored yellow. The graph U is used for finding yellow components. The consensus tree returned is found by combining the trees constructed for the yellow components. The graph D is used for finding edges in U that can be colored red, these edges correspond to the so-called maximal nodes in D .

A *maximal node* in D is a node with no outgoing edges, and a red edge whose endpoints are in different yellow components is called a *separable red edge*.

By slightly modifying the algorithm of Henzinger et al. and combining it with an algorithm for minimum weight edge cut (Karger, 1996), we can implement Heuristic 2 as follows.

Heuristic 2.

1. Construct U and D . Add weights to the edges in U . The weight of an edge $\{a, b\}$ in U is equal to the sum of the weights of constraints of the form $(\{a, b\}, \cdot)$. Color all nodes in D and edges in U yellow.

2. Identify maximal nodes in D . Recolor these nodes and the corresponding edges of U red.
3. If U has no edges, then return the consensus tree T with a root and all nodes in U children of the root. Otherwise, compute yellow components of U . If there is only one yellow component then find a minimum weight edge cut, delete the edges in the cut from U and the corresponding nodes from D , and recompute the yellow components. Let C_1, \dots, C_k be the current yellow components. Form a tree T by creating the root of T and connecting the root of the consensus trees recursively computed for the components C_1, \dots, C_k to it. For each current yellow component, identify the set E_{sep} of separable red edges incident to that new component. Delete these edges from U and the corresponding nodes from D . Go to Step 2.

Lemma 7. *Heuristic 2 can be implemented to run in expected time $O(n^3 \log n + m \log^3 n)$.*

Proof: A minimum weight edge cut can be computed with high probability in time $O(n^2 \log n)$ (Karger, 1996). In the worst case it has to be done n times; hence, the calls to minimum weight edge cut take a total of $O(n^3 \log n)$ expected time. All other operations can be performed in expected time $O(m \log^3 n)$ like in the algorithm of Henzinger et al. Thus, the total expected time is $O(n^3 \log n + m \log^3 n)$. \square

Lemma 8. *Let I be an instance of MICT, and let T be the tree produced by Heuristic 2 for I . The total weight of constraints in I not satisfied by T is at most $\text{height}(T)$ times the minimum.*

Proof: Let J be a subset of I of minimum total weight such that $I \setminus J$ has a consensus tree. Next, let D be the set of connected components in the auxiliary graph where edges corresponding to the constraints in J are deleted. Suppose that Heuristic 2 at some stage finds a min-cut in a currently connected fragment C . Clearly, C cannot be a subset of a simple component in D since then there wouldn't exist a consensus tree for $I \setminus J$. Hence, there is a subset J_C of J such that the set of edges corresponding to the constraints in J_C disconnects $G(C)$ into disjoint components. Clearly, the total weight of J_C is not smaller than the weight of a min-cut of $G(C)$. Now, it is sufficient to observe that the subsets J_C for distinct C 's on the same recursion level of Heuristic 2 are pairwise disjoint. \square

Theorem 5. *Let n , w , t be respectively the number of leaves, the total weight of constraints, and the minimum total weight of the constraints to remove in an instance I of MILCT. Heuristic 2 constructs a consensus tree for a subset of the constraints in I whose total weight is not smaller than $w - nt$.*

Note that the number of constraints in I might be cubic in n and that Heuristic 2 yields a better approximation factor than Heuristic 1 for MILCT whenever $t < \frac{2w}{3n}$.

6. Open problems

We do not know whether or not it is possible to find a polynomial-time approximation scheme for instances of MHT with $O(1)$ trees of height $O(1)$.

It follows from Theorem 3 and the definition of MICT that MICT is strongly NP-complete. Hence, it cannot admit a fully polynomial-time approximation scheme (Papadimitriou, 1994). However, it is an open question whether it admits a polynomial-time approximation scheme or at least a polynomial-time heuristic with a smaller approximation factor.

The complexity status of MILCT is also an interesting open question. If MILCT is NP-complete, does it admit a polynomial-time approximation scheme?

On a high level, the definitions of MICT and MILCT resemble those of MAX SAT and MAX k-SAT (Hochbaum, 1995; Johnson, 1974). In the design of Heuristic 1 we have utilized this similarity taking inspiration from the early heuristic for MAX k-SAT due to Johnson (1974). Recently, substantial progress in approximating MAX SAT and MAX k-SAT has been made by using linear programming, semidefinite programming, and randomized rounding (Goemans and Williamson, 1994; Hochbaum, 1995). One of the main obstacles in applying these techniques to MICT is the complexity of “arithmetization” of the proper-descendant lowest-common-ancestor relation (the case of MILCT seems more promising).

Acknowledgments

The authors were supported in part by TFR (Swedish Research Council for Engineering Sciences) and in part by NUF-NAL (The Nuffield Foundation Awards to Newly Appointed Lecturers).

References

- A.V. Aho, Y. Sagiv, T.G. Szymanski, and J.D. Ullman, “Inferring a tree from lowest common ancestors with an application to the optimization of relational expressions,” *SIAM Journal of Computing*, vol. 10, no. 3, pp. 405–421, 1981.
- M. Farach, T. Przytycka, and M. Thorup, “Computing the agreement of trees with bounded degrees,” in *Proc. of the 3rd ESA*, 1995, pp. 381–393.
- M. Farach and M. Thorup, “Fast comparison of evolutionary trees,” in *Proc. of the 5th ACM-SIAM SODA*, 1994a, pp. 481–488.
- M. Farach and M. Thorup, “Optimal evolutionary tree comparison by sparse dynamic programming,” in *Proc. of the 35th FOCS*, 1994b, pp. 770–779.
- C.R. Finden and A.D. Gordon, “Obtaining common pruned trees,” *Journal of Classification*, vol. 2, pp. 255–276, 1985.
- M.X. Goemans and D.P. Williamson, “New $\frac{3}{4}$ -approximation algorithms for MAX SAT,” *SIAM Journal of Discrete Mathematics*, vol. 7, pp. 656–666, 1994.
- J. Håstad, “Testing of the long code and hardness for clique,” in *Proc. of the 28th ACM STOC*, 1996, pp. 11–19.
- J. Hein, T. Jiang, L. Wang, and K. Zhang, “On the complexity of comparing evolutionary trees,” *Discrete Applied Mathematics*, vol. 71, pp. 153–169, 1996.
- M.R. Henzinger, V. King, and T. Warnow, “Constructing a tree from homeomorphic subtrees, with applications to computational biology,” in *Proc. of the 7th ACM-SIAM SODA*, 1996, pp. 333–340.

- D.S. Hochbaum, Ed., *Approximation Algorithms for NP-hard Problems*, PWS Publishing Company: Boston, 1995.
- C.A.J. Hurkens and A. Schrijver, "On the size of systems of sets every t of which have an SDR, with an application to the worst-case ratio of heuristics for packing problems," *SIAM Journal of Discrete Mathematics*, vol. 2, no. 1, pp. 68–72, 1989.
- D.S. Johnson, "Approximation algorithms for combinatorial problems," *Journal of Computer and System Sciences*, vol. 9, pp. 256–278, 1974.
- S. Kannan, T. Warnow, and S. Yoosseph, "Computing the local consensus of trees," in *Proc. of the 6th ACM-SIAM SODA*, 1995, pp. 68–77.
- M.Y. Kao, T.W. Lam, T. Przytycka, W.K. Sung, and H.F. Ting, "General techniques for comparing unrooted evolutionary trees," in *Proc. of the 29th ACM STOC*, 1997, pp. 54–65.
- D.R. Karger, "Minimum cuts in near-linear time," in *Proc. of the 28th ACM STOC*, 1996, pp. 56–63.
- D. Keselman and A. Amir, "Maximum agreement subtree in a set of evolutionary trees—Metrics and efficient algorithms," in *Proc. of the 35th FOCS*, 1994, pp. 758–769.
- T.W. Lam, W.K. Sung, and H.F. Ting, "Computing the unrooted maximum agreement subtree in sub-quadratic time," in *Proc. of the 5th SWAT*, 1996, pp. 124–135.
- C.H. Papadimitriou, "Computational Complexity," Addison-Wesley: Reading, 1994.
- C. Phillips and T.J. Warnow, "The asymmetric median tree—a new model for building consensus Trees," in *Proc. of the 7th CPM, LNCS 1075*, 1996, pp. 234–252.
- M. Steel and T. Warnow, "Kaikoura tree theorems: Computing the maximum agreement subtree," *Information Processing Letters*, vol. 48, pp. 77–82, 1993.