The HKU Scholars Hub   The University of Hong Kong   香港大學學術庫

Title	**Cavity matchings, label compressions, and unrooted evolutionary trees**
Author(s)	**Kao, MY; Lam, TW; Sung, WK; Ting, HF**
Citation	**Siam Journal On Computing, 2000, v. 30 n. 2, p. 602-624**
Issued Date	**2000**
URL	**http://hdl.handle.net/10722/43657**
Rights	**Creative Commons: Attribution 3.0 Hong Kong License**

CAVITY MATCHINGS, LABEL COMPRESSIONS, AND UNROOTED EVOLUTIONARY TREES*

MING-YANG KAO†, TAK-WAH LAM‡, WING-KIN SUNG‡, AND HING-FUNG TING‡

Abstract. We present an algorithm for computing a maximum agreement subtree of two unrooted evolutionary trees. It takes $O(n^{1.5} \log n)$ time for trees with unbounded degrees, matching the best known time complexity for the rooted case. Our algorithm allows the input trees to be mixed trees, i.e., trees that may contain directed and undirected edges at the same time. Our algorithm adopts a recursive strategy exploiting a technique called label compression. The backbone of this technique is an algorithm that computes the maximum weight matchings over many subgraphs of a bipartite graph as fast as it takes to compute a single matching.

Key words. computational biology, evolutionary trees, all-cavity maximum weight matchings, label compressions, unrooted trees, mixed trees

AMS subject classifications. 05C05, 05C85, 05C90, 68Q25, 92B05

PII. S0097539797332275

1. Introduction. An *evolutionary* tree is one whose leaves are labeled with distinct symbols representing species. Evolutionary trees are useful for modeling the evolutionary relationship of species [1, 4, 6, 16, 17, 25]. An *agreement subtree* of two evolutionary trees is an evolutionary tree that is also a topological subtree of the two given trees. A *maximum* agreement subtree is one with the largest possible number of leaves. Different models about the evolutionary relationship of the same species may result in different evolutionary trees. A fundamental problem in computational biology is to determine how much two models of evolution have in common. To a certain extent, this problem can be solved by computing a maximum agreement subtree of two given evolutionary trees [12].

Algorithms for computing a maximum agreement subtree of two unrooted evolutionary trees as well as two rooted trees have been studied intensively in the past few years. The unrooted case is more difficult than the rooted case. There is indeed a linear-time reduction from the rooted case to the unrooted one, but the reverse is not known. Steel and Warnow [24] gave the first polynomial-time algorithm for unrooted trees, which runs in $O(n^{4.5} \log n)$ time. Farach and Thorup reduced the time to $O(n^{2+o(1)})$ for unrooted trees [10] and $O(n^{1.5} \log n)$ for rooted trees [11]. For the unrooted case, the time was improved by Lam, Sung, and Ting [22] to $O(n^{1.75+o(1)})$. Algorithms that work well for rooted trees with degrees bounded by a constant have also been revealed recently. The algorithm of Farach, Przytycka, and Thorup [9] takes $O(n \log^3 n)$ time, and that of Kao [20] takes $O(n \log^2 n)$ time. Cole and Hariharan [7]

gave an $O(n \log n)$-time algorithm for the case where the input is further restricted to binary rooted trees.

This paper presents an algorithm for computing a maximum agreement subtree of two unrooted trees. It takes $O(n^{1.5} \log n)$ time for trees with unbounded degrees, matching the best known time complexity for the rooted case [11]. If the degrees are bounded by a constant, the running time is only $O(n \log^4 n)$. We omit the details of this reduction since Przytycka [23] recently devised an $O(n \log n)$-time algorithm for the same case.

Our algorithm allows the input trees to be *mixed* trees, i.e., trees that may contain directed and undirected edges at the same time [15, 18]. Such trees can handle a broader range of information than rooted and unrooted trees. To simplify the discussion, this paper focuses on unrooted trees. Our subtree algorithm adopts a conceptually simple recursive strategy exploiting a novel technique called *label compression*. This technique enables our algorithm to process overlapping subtrees iteratively while keeping the total tree size very close to the original input size. Label compression builds on an unexpectedly fast algorithm for the *all-cavity maximum weight matching* problem [21], which asks for the weight of a maximum weight matching in $G - \{u\}$ for each node u of a bipartite graph G with integer edge weights. If G has n nodes, m edges and maximum edge weight N, the algorithm takes $O(\sqrt{n}m \log(nN))$ time, which matches the best known time bound for computing a single maximum weight matching of G, due to Gabow and Tarjan [13].

In section 2, we solve the all-cavity matching problem. In section 3, we formally define maximum agreement subtrees and outline our recursive strategy for computing them. We describe label compression in section 4, detail our subtree algorithm in section 5, and discuss how to compute auxiliary information for label compression in sections 6 and 7. We conclude by extending the subtree algorithm to mixed trees in section 8.

2. All-cavity maximum weight matching. Let $G = (X, Y, E)$ be a bipartite graph with n nodes and m edges where each edge (u, v) has a positive integer weight $w(u, v) \leq N$. Let $\text{mwm}(G)$ denote the weight of a maximum weight matching in G. The all-cavity matching problem asks for $\text{mwm}(G - \{u\})$ for all $u \in X \cup Y$. A naive approach to solve this problem is to compute $\text{mwm}(G - \{u\})$ separately for each u using the fastest algorithm for computing a single maximum weight matching [13], thus taking $O(n^{1.5}m \log(nN))$ total time. A main finding of this paper is that the matchings in different subgraphs $G - \{u\}$ are closely related and can be represented succinctly. From this representation, we can solve the problem in $O(\sqrt{n}m \log(nN))$ time. By symmetry, we detail only how to compute $\text{mwm}(G - \{u\})$ for all $u \in X$. Below we assume $m \geq n/2$; otherwise, we remove the degree-zero nodes and work on the smaller resultant graph.

A node v of G is *matched* by a matching of G if v is an endpoint of an edge in the matching. In the remainder of this section, let M be a fixed maximum weight matching of G; also let $w(H)$ be the total weight of a set H of edges. An *alternating path* is a simple path P in G such that (1) P starts with an edge in M; (2) the edges of P alternate between M and $E - M$; and (3) if P ends at an edge $(u, v) \notin M$, then v is not matched by M. An *alternating cycle* is a simple cycle C in G whose edges alternate between M and $E - M$. P (respectively, C) can *transform* M to another matching $M' = P \cup M - P \cap M$ (respectively, $C \cup M - C \cap M$). The *net change* induced by P, denoted by $\Delta(P)$, is $w(M') - w(M)$, i.e., the total weight of the edges of P in $E - M$ minus that of the edges of P in M. The *net change* induced by C is

FIG. 1. (a) *a bipartite graph G; (b) the corresponding directed graph D.*

defined similarly.

The next lemma divides the computation of $\mathrm{mwm}(G - \{u\})$ into two cases.

LEMMA 2.1. *Let $u \in X$.*

1. *If u is not matched by M, then M is also a maximum weight matching in $G - \{u\}$ and $\mathrm{mwm}(G - \{u\}) = \mathrm{mwm}(G)$.*
2. *If u is matched by M, then G contains an alternating path P starting from u, which can transform M to a maximum weight matching in $G - \{u\}$.*

Proof. Statement 1 is straightforward. To prove statement 2, let M' be a maximum weight matching in $G - \{u\}$. Consider the edges in $M \cup M' - M \cap M'$. They form a set S of alternating paths and cycles. Since u is matched by M but not by M', u is of degree 1 in $M \cup M' - M \cap M'$. Let P be the alternating path in S with u as an endpoint. Let M'' be the matching obtained by transforming M only with P. Since u is not matched by M'', M'' is a matching in $G - \{u\}$. M' can be obtained by further transforming M'' with the remaining alternating paths and cycles in S. The net change induced by each of these alternating paths and cycles is nonpositive; otherwise, such a path or cycle can improve M and we obtain a contradiction. Therefore, $w(M'') \geq w(M')$, i.e., both M' and M'' are maximum weight matchings in $G - \{u\}$. □

By Lemma 2.1(2), we can compute $\mathrm{mwm}(G - \{u\})$ for any $u \in X$ matched by M by finding the alternating path starting from u with the largest net change. Below we construct a directed graph D, which enables us to identify such an alternating path for every node easily. The node set of D is $X \cup Y \cup \{t\}$, where t is a new node. The edge set of D is defined as follows; see Figure 1 for an example.

- If $x \in X$ is not matched by M, D has an edge from x to t with weight zero.
- If $y \in Y$ is matched by M, D has an edge from y to t with weight zero.
- If M has an edge (x, y) where $x \in X$ and $y \in Y$, D has an edge from x to y with weight $-w(x, y)$.
- If $E - M$ has an edge (x, y) where $x \in X$ and $y \in Y$, D has an edge from y to x with weight $w(x, y)$.

Note that D has $n + 1$ nodes and at most $n + m$ edges. The weight of each edge

in D is an integer in $[-N, N]$.

LEMMA 2.2.

1. *D contains no positive-weight cycle.*
2. *Each alternating path P in G that starts from $u \in X$ corresponds to a simple path Q in D from u to t, and vice versa. Also, $\Delta(P) = w(Q)$.*
3. *For each $u \in X$ matched by M, $\mathrm{mwm}(G - \{u\})$ is the sum of $\mathrm{mwm}(G)$ and the weight of the longest path in D from u to t.*

Proof. Statement 1. Consider a simple cycle $C = u_1, u_2, \ldots, u_k, u_1$ in D. Since t has no outgoing edges, no u_i equals t. By the definition of D, C is also an alternating cycle in G. Therefore, $w(C)$ is the net change induced by transforming M with C. Since M is a maximum weight matching in G, this net change is nonpositive.

Statement 2. Consider an alternating path $P = u, u_1, u_2, \ldots, u_k$ in G starting from u. In D, P is also a simple path. If $u_k \in X$, then u_k is not matched by M, and D contains the edge (u_k, t). If $u_k \in Y$, then u_k is matched by M, and D again contains the edge (u_k, t). Therefore, D contains the simple path $Q = u, u_1, u_2, \ldots, u_k, t$. The weight of Q is $\Delta(P)$. The reverse direction of the statement is straightforward.

Statement 3. This statement follows from Lemma 2.1(2) and Statement 2.   □

THEOREM 2.3. *Given G, we can compute $\mathrm{mwm}(G - \{u\})$ for all nodes $u \in G$ in $O(\sqrt{n}m \log(nN))$ time.*

Proof. By symmetry and Lemmas 2.1(1) and 2.2(3), we compute $\mathrm{mwm}(G - \{u\})$ for all $u \in X$ as follows.

1. Compute a maximum weight matching M of G.
2. Construct D as above and find the weights of its longest paths to t.
3. For each $u \in X$, if u is matched by M, then $\mathrm{mwm}(G - \{u\})$ is the sum of $\mathrm{mwm}(G)$ and the weight of the longest path from u to t in D; otherwise, $\mathrm{mwm}(G - \{u\}) = \mathrm{mwm}(G)$.

Step 1 takes $O(\sqrt{n}m \log(nN))$ time. At step 2, constructing D takes $O(n + m)$ time, and the single-destination longest paths problem takes $O(\sqrt{n}m \log N)$ time [14]. Step 3 takes $O(n)$ time. Thus, the total time is $O(\sqrt{n}m \log(nN))$.   □

3. The main result. This section gives a formal definition of maximum agreement subtrees and an overview of our new subtree algorithm.

3.1. Basics. Throughout the remainder of this paper, unrooted trees are denoted by U or X, and rooted trees by T, W, or R. A node of degree 0 or 1 is a *leaf*; otherwise, it is *internal*. Adopted to avoid technical trivialities, this definition is somewhat nonstandard in that if the root of a rooted tree is of degree 1, it is also a leaf.

For an unrooted tree U and a node $u \in U$, let U^u denote the rooted tree constructed by rooting U at u. For a rooted tree T and a node $v \in T$, let T^v denote the *rooted subtree* of T that comprises v and its descendents. Similarly, for a node $v \in U^u$, U^{uv} is the rooted subtree of U^u rooted at v, which is also called a *rooted subtree* of U.

An *evolutionary* tree is a tree whose leaves are labeled with distinct symbols. Let T be a rooted evolutionary tree with leaves labeled over a set L. A label subset $L' \subseteq L$ *induces* a subtree of T, denoted by $T|L'$, whose nodes are the leaves of T labeled over L' as well as the least common ancestors of such leaves in T, and whose edges preserve the ancestor-descendent relationship of T. Consider two rooted evolutionary trees T_1 and T_2 labeled over L. Let T_1' be a subtree of T_1 induced by some subset of L. We similarly define T_2' for T_2. If there exists an isomorphism between T_1' and T_2' mapping each leaf in T_1' to one in T_2' with the same label, then T_1'

and T_2' are each called *agreement subtrees* of T_1 and T_2. Note that this isomorphism is unique. Consider any nodes $u \in T_1$ and $v \in T_2$. We say that u *is mapped to v in* T_1' *and* T_2' if this isomorphism maps u to v. A *maximum* agreement subtree of T_1 and T_2 is one containing the largest possible number of labels. Let $\text{mast}(T_1, T_2)$ denote the number of labels in such a tree. A *maximum agreement subtree* of two unrooted evolutionary trees U_1 and U_2 is one with the largest number of labels among the maximum agreement subtrees of U_1^u and U_2^v over all nodes $u \in U_1$ and $v \in U_2$. Let

$$(3.1) \qquad \text{mast}(U_1, U_2) = \max\{\text{mast}(U_1^u, U_2^v) \mid u \in U_1, v \in U_2\}.$$

Remark. The nodes u (or v) can be restricted to internal nodes when the trees have at least three nodes. We can also generalize the above definition to handle a pair of rooted tree and unrooted tree (T, U). That is, $\text{mast}(T, U)$ is defined to be $\max\{\text{mast}(T, U^v) \mid v \in U\}$.

3.2. Our subtree algorithm. The next theorem is our main result. The *size* $|U|$ (or $|T|$) of an unrooted tree U (or a rooted tree T) is its node count.

THEOREM 3.1. *Let U_1 and U_2 be two unrooted evolutionary trees. We can compute* $\text{mast}(U_1, U_2)$ *in* $O(N^{1.5} \log N)$ *time, where* $N = \max\{|U_1|, |U_2|\}$.

We prove Theorem 3.1 by presenting our algorithm in a top-down manner with an outline here. As in previous work, our algorithm only computes $\text{mast}(U_1, U_2)$ and can be augmented to report a corresponding subtree. It uses graph separators. A *separator* of a tree is an internal node whose removal divides the tree into connected components each containing at most half of the tree's nodes. Every tree that contains at least three nodes has a separator, which can be found in linear time.

If U_1 or U_2 has at most two nodes, $\text{mast}(U_1, U_2)$ as defined in (3.1) can easily be computed in $O(N)$ time. Otherwise, both trees have at least three nodes each, and we can find a separator x of U_1. We then consider three cases.

Case 1. In some maximum agreement subtree of U_1 and U_2, the node x is mapped to a node $y \in U_2$. In this case, $\text{mast}(U_1, U_2) = \text{mast}(U_1^x, U_2)$. To compute $\text{mast}(U_1^x, U_2)$, we might simply evaluate $\text{mast}(U_1^x, U_2^y)$ for different y in U_2. This approach involves solving the mast problem for $\Theta(N)$ different pairs of rooted trees and introduces much redundant computation. For example, consider a rooted subtree R of U_2. For all $y \in U_2 - R$, R is a common subtree of U_2^y. Hence, R is examined repeatedly in the computation of $\text{mast}(U_1^x, U_2^y)$ for these y. To speed up the computation, we devise the technique of label compression in section 4 to elicit sufficient information between U_1^x and R so that we can compute $\text{mast}(U_1^x, U_2^y)$ for all $y \in U_2 - R$ without examining R. This leads to an efficient algorithm for handling Case 1; the time complexity is stated in the following lemma.

LEMMA 3.2. *Assume that U_1 and U_2 have at least three nodes each. Given an internal node $x \in U_1$, we can compute* $\text{mast}(U_1^x, U_2)$ *in* $O(N^{1.5} \log N)$ *time.*

Proof. See section 4 to section 7. $\square$

Case 2. In some maximum agreement subtree of U_1 and U_2, two certain nodes v_1 and v_2 of U_1 are mapped to nodes in U_2, and x is on the path in U_1 between v_1 and v_2. This case is similar to Case 1. Let $\tilde{U}_2$ be the tree constructed by adding a dummy node in the middle of every edge in U_2. Then, $\text{mast}(U_1, U_2) = \text{mast}(U_1^x, \tilde{U}_2^y)$ for some dummy node y in $\tilde{U}_2$. Thus, $\text{mast}(U_1, U_2) = \text{mast}(U_1^x, \tilde{U}_2)$. As in Case 1, $\text{mast}(U_1^x, \tilde{U}_2)$ can be computed in $O(N^{1.5} \log N)$ time.

Case 3. Neither Case 1 nor Case 2 holds. Let $U_{1,1}, U_{1,2}, \ldots, U_{1,b}$ be the evolutionary trees formed by the connected components of $U_1 - \{x\}$. Let $J_1, \ldots, J_b$ be the sets of labels in these components, respectively. Then, a maximum agreement subtree of

```
/* U₁ and U₂ are unrooted trees. */
mast(U₁, U₂)
    find a separator x of U₁;
    construct Ũ₂ by adding a dummy node w at the middle of each edge (u, v) in
    U₂;
    val = mast(U₁ˣ, U₂);
    val' = mast(U₁ˣ, Ũ₂);
    let U₁,₁, U₁,₂, ..., U₁,ᵦ be the connected components of U₁ − {x};
    for all i ∈ [1, b], let Jᵢ be the set of labels of U₁,ᵢ;
    for all i ∈ [1, b], set valᵢ = mast(U₁,ᵢ, U₂|Jᵢ)};
    return max{val, val', max₁≤ᵢ≤ᵦ valᵢ};
```

FIG. 2. *Algorithm for computing* $\mathrm{mast}(U_1, U_2)$.

U_1 and U_2 is labeled over some J_i. Therefore, $\mathrm{mast}(U_1, U_2) = \max\{\mathrm{mast}(U_{1,i}, U_2|J_i) \mid i \in [1, b]\}$, and we compute each $\mathrm{mast}(U_{1,i}, U_2|J_i)$ recursively.

Figure 2 summarizes the steps for computing $\mathrm{mast}(U_1, U_2)$. Here we analyze the time complexity $T(N)$ based on Lemma 3.2. Cases 1 and 2 each take $O(N^{1.5} \log N)$ time. Let $N_i = |U_{1,i}|$. Then Case 3 takes $\sum_{i \in [1,b]} T(N_i)$ time. By recursion,

$$T(N) = O(N^{1.5} \log N) + \sum_{i \in [1,b]} T(N_i).$$

Since x is a separator of U_1, $N_i \leq \frac{N}{2}$. Then, since $\sum_{i \in [1,b]} N_i \leq N$, $T(N) = O(N^{1.5} \log N)$ [5, 19] and the time bound in Theorem 3.1 follows. To complete the proof of Theorem 3.1, we devote section 4 through section 7 to proving Lemma 3.2.

4. Label compressions. To compute a maximum agreement subtree, our algorithm recursively processes overlapping subtrees of the input trees. The technique of label compression compresses overlapping parts of such subtrees to reduce their total size. We define label compressions with respect to a rooted subtree in section 4.1 and with respect to two label-disjoint rooted subtrees in section 4.2. We do not use label compression with respect to three or more trees.

As a warm-up, let us define a concept called *subtree shrinking*, which is a primitive form of label compression. Let T be a rooted tree. Let R be a rooted subtree of T. Let $T \ominus R$ denote the rooted tree obtained by replacing R with a leaf γ. We say that γ is a *shrunk* leaf. The other leaves are *atomic* leaves. Similarly, for two label-disjoint rooted subtrees R_1 and R_2 of T, let $T \ominus (R_1, R_2)$ denote the rooted tree obtained by replacing R_1 and R_2 with *shrunk* leaves γ_1 and γ_2, respectively. We extend these notions to an unrooted tree U and define $U \ominus R$ and $U \ominus (R_1, R_2)$ similarly.

4.1. Label compression with respect to one rooted subtree. Let T be a rooted tree. Let v be a node in T and u an ancestor of v. Let P be the path of T from u to v. A node *lies* between u and v if it is in P but differs from u and v. A subtree of T is *attached* to u if it is some T^w where w is a child of u. A subtree of T *hangs* between u and v if it is attached to some node lying between u and v, but its root is not in P and is not v.

We are now ready to define the concept of label compression. Let T and R be rooted evolutionary trees labeled over L and K, respectively. The *compression* of T with respect to R, denoted by $T \otimes R$, is a tree constructed by affixing extra nodes to

FIG. 3. *An example of label compression.*

$T|(L - K)$ with the following steps; see Figure 3 for an example. Consider each node y in $T|(L - K)$, let x be its parent in $T|(L - K)$.

- Let $\mathcal{A}(T, K, y)$ denote the set of subtrees of T that are attached to y and whose leaves are all labeled over K. If $\mathcal{A}(T, K, y)$ is nonempty, compress all the trees in $\mathcal{A}(T, K, y)$ into a single node z_1 and attach it to y.
- Let $\mathcal{H}(T, K, y)$ denote the set of subtrees of T that hang between x and y (by definition of $T|(L - K)$, these subtrees are all labeled over K). If $\mathcal{H}(T, K, y)$ is nonempty, compress the parents $p_1, \ldots, p_m$ of the roots of the trees in $\mathcal{H}(T, K, y)$ into a single node p_1, and insert it between x and y; also compress all the trees in $\mathcal{H}(T, K, y)$ into a single node z_2 and attach it to p_1.

The nodes z_1, z_2, and p_1 are called *compressed* nodes, and the leaves in $T \otimes R$ that are not compressed are *atomic* leaves.

We further store in $T \otimes R$ some auxiliary information about the relationship between T and R. For an internal node v in $T \otimes R$, let $\alpha(v) = \text{mast}(T^v, R)$. For a compressed leaf v in $T \otimes R$, if it is compressed from a set of subtrees $T^{v_1}, \ldots, T^{v_s}$, let $\alpha(v) = \max\{\text{mast}(T^{v_1}, R), \ldots, \text{mast}(T^{v_s}, R)\}$.

Let T_1 and T_2 be two rooted evolutionary trees. Assume T_2 contains a rooted subtree R. Given $T_1 \otimes R$, we can compute $\text{mast}(T_1, T_2)$ without examining R. We first construct $T_1 \ominus R$ by replacing R of T_2 with a shrunk leaf and then compute $\text{mast}(T_1, T_2)$ from $T_1 \otimes R$ and $T_2 \ominus R$. To further our discussion, we next generalize the definition of maximum agreement subtree for a pair of trees that contain compressed leaves and a shrunk leaf, respectively.

Let $W_1 = T_1 \otimes R$ and $W_2 = T_2 \ominus R$. Let γ be the shrunk leaf in W_2. We define an agreement subtree of W_1 and W_2 similar to that of ordinary evolutionary trees. An atomic leaf must still be mapped to an atomic leaf with the same label. However, the shrunk leaf γ of W_2 can be mapped to any internal node or compressed leaf v of W_1 as long as $\alpha(v) > 0$. The size of an agreement subtree is the number of its atomic leaves, plus $\alpha(v)$ if γ is mapped to a node $v \in W_1$. A *maximum* agreement subtree of W_1 and W_2 is one with the largest size. Let $\text{mast}(W_1, W_2)$ denote the size of such a subtree. The following lemma is the cornerstone of label compression.

LEMMA 4.1. $\text{mast}(T_1, T_2) = \text{mast}(W_1, W_2)$.

Proof. It follows directly from the definition. $\square$

We can compute $\text{mast}(W_1, W_2)$ as if W_1 and W_2 were ordinary rooted evolutionary trees [9, 11, 20] with a special procedure on handling the shrunk leaf. The time complexity is stated in the following lemma. Let $n = \max\{|W_1|, |W_2|\}$ and $N = \max\{|T_1|, |T_2|\}$.

LEMMA 4.2. *Suppose that all the auxiliary information of W_1 has been given.*

Then $\mathrm{mast}(W_1, W_2)$ *can be computed in* $O(n^{1.5} \log N)$ *time, and afterwards we can retrieve* $\mathrm{mast}(W_1^v, W_2)$ *for any node* $v \in W_1$ *in* $O(1)$ *time.*

Proof. We adapt Farach and Thorup's rooted subtree algorithm [11] to compute $\mathrm{mast}(W_1, W_2)$. Details are given in Appendix A. $\square$

We demonstrate a scenario where label compression speeds up the computation of $\mathrm{mast}(U_1^x, U_2)$ for Lemma 3.2. Suppose that we can identify a rooted subtree R of U_2 such that x is mapped to a node outside R, i.e., we can reduce (3.1) to

$$(4.1) \quad \mathrm{mast}(U_1^x, U_2) = \max\{\mathrm{mast}(U_1^x, U_2^y) \mid y \text{ is an internal node not in } R\}.$$

Note that every U_2^y contains R as a common subtree. To avoid overlapping computation on R, we construct $W = U_1^x \otimes R$ and $X = U_2 \ominus R$. Then $X^y = U_2^y \ominus R$ and from Lemma 4.1, $\mathrm{mast}(U_1^x, U_2^y) = \mathrm{mast}(W, X^y)$. We rewrite (4.1) as

$$(4.2) \quad \mathrm{mast}(U_1^x, U_2) = \max\{\mathrm{mast}(W, X^y) \mid y \text{ is an internal node of } X\}.$$

If R is large, then W and X are much smaller than U_1^x and U_2. Consequently, it is beneficial to compress U_1^x and compute $\mathrm{mast}(U_1^x, U_2)$ according to (4.2).

4.2. Label compression with respect to two rooted subtrees. Let T, R_1, R_2 be rooted evolutionary trees labeled over L, K_1, K_2, respectively, where $K_1 \cap K_2 = \phi$. Let $K = K_1 \cup K_2$. The *compression* of T with respect to R_1 and R_2, denoted by $T \otimes (R_1, R_2)$, is a tree constructed from $T|(L - K)$ by the following two steps. For each node y and its parent x in $T|(L - K)$,

1. if $\mathcal{A}(T, K, y)$ is nonempty, compress all the trees in $\mathcal{A}(T, K, y)$ into a single leaf z and attach it to y; create and attach an auxiliary node $\bar{z}$ to y;
2. if $\mathcal{H}(T, K, y)$ is nonempty, compress the parents $p_1, \ldots, p_m$ of the roots of the subtrees in $\mathcal{H}(T, K, y)$ into a single node p_1 and insert it between x and y; compress the subtrees in $\mathcal{H}(T, K, y)$ into a single node z and attach it to p_1; create and insert an auxiliary node $\bar{p}_1$ between p_1 and y; create auxiliary nodes $\bar{z}$ and $\bar{\bar{z}}$ and attach them to p_1 and $\bar{p}_1$, respectively.

The nodes p_1 and z are *compressed* nodes of $T \otimes (R_1, R_2)$. The nodes $\bar{p}_1, \bar{z}$, and $\bar{\bar{z}}$ are *auxiliary* nodes. These nodes are added to capture the topology of T that is isomorphic with the subtrees R_1 and R_2 of T'.

We also store auxiliary information in $T \otimes (R_1, R_2)$. Let R^+ be the tree obtained by connecting R_1 and R_2 together with a node, which becomes the root of R^+.

Consider the internal nodes of $T \otimes (R_1, R_2)$. If v is an internal node inherited from $T|(L - K)$, then let $\alpha_1(v) = \mathrm{mast}(T^v, R_1)$ and $\alpha_2(v) = \mathrm{mast}(T^v, R_2)$. If p_1 and $\bar{p}_1$ are internal nodes compressed from some path $p_1, \ldots, p_m$ of T, then only p_1 stores the values $\alpha_1(p_1) = \mathrm{mast}(T^{p_1}, R_1)$, $\alpha_2(p_1) = \mathrm{mast}(T^{p_1}, R_2)$, and $\alpha_+(p_1) = \mathrm{mast}(T^{p_1}, R^+)$.

We do not store any auxiliary information at the atomic leaves in $T \otimes (R_1, R_2)$. Consider the other leaves in $T \otimes (R_1, R_2)$ based on how they are created.

Case 1. Nodes $z, \bar{z}$ are leaves created with respect to $\mathcal{A}(T, K, y)$ for some node y in $T|(L - K)$. Let $\mathcal{A}(T, K, y) = \{T^{v_1}, \ldots, T^{v_k}\}$. We store the following values at z.
- $\alpha_1(z) = \max\{\mathrm{mast}(T^{v_i}, R_1) \mid i \in [1, k]\}$, $\alpha_2(z) = \max\{\mathrm{mast}(T^{v_i}, R_2) \mid i \in [1, k]\}$, $\alpha_+(z) = \max\{\mathrm{mast}(T^{v_i}, R^+) \mid i \in [1, k]\}$;
- $\beta(z) = \max\{\mathrm{mast}(T^{v_i}, R_1) + \mathrm{mast}(T^{v_{i'}}, R_2) \mid T^{v_i} \text{ and } T^{v_{i'}} \text{ are distinct subtrees in } \mathcal{A}(T, K, y)\}$.

Case 2. Nodes $z, \bar{z}$, and $\bar{\bar{z}}$ are leaves created with respect to the subtrees in $\mathcal{H}(T, K, y) = \{T^{v_1}, \ldots, T^{v_k}\}$ for some node y in $T|(L - K)$. We store the following values at z:

- $\alpha_1(z)$, $\alpha_2(z)$, and $\alpha_+(z)$ as in Case 1;
- $\beta(z) = \max\{\text{mast}(T^{v_i}, R_1) + \text{mast}(T^{v_j}, R_2) \mid T^{v_i} \text{ and } T^{v_j} \text{ are distinct subtrees}$ in $\mathcal{H}(T, K, y)$ that are attached to the same node in $T\}$;
- $\beta_{1 \succ 2}(z) = \max\{\text{mast}(T^{v_j}, R_1) + \text{mast}(T^{v_{j'}}, R_2) \mid (j, j') \in Z\}$ and
  $\beta_{2 \succ 1}(z) = \max\{\text{mast}(T^{v_j}, R_2) + \text{mast}(T^{v_{j'}}, R_1) \mid (j, j') \in Z\}$,
  where $Z = \{(j, j') \mid T^{v_j}, T^{v_{j'}} \in \mathcal{H}(T, K, y) \text{ and the parent of } v_j \text{ in } T \text{ is a}$ proper ancestor of the parent of $v_{j'}\}$.

Let T_1 and T_2 be rooted evolutionary trees. Let R_1 and R_2 be label-disjoint rooted subtrees of T_2. Let $W_1 = T \otimes (R_1, R_2)$ and $W_2 = T' \ominus (R_1, R_2)$. Below, we give the definition of a maximum agreement subtree of W_1 and W_2.

Let γ_1 and γ_2 be the two shrunk leaves in W_2 representing R_1 and R_2, respectively. Let y_c be the least common ancestor of γ_1 and γ_2 in W_2. Intuitively, in a pair of agreement subtrees (W_1', W_2') of W_1 and W_2, atomic leaves are mapped to atomic leaves, and shrunk leaves are mapped to internal nodes or leaves. Moreover, we allow W_2' to contain y_c as a leaf, which can be mapped to an internal node or leaf of W_1'. More formally, we require that there is an isomorphism between W_1' and W_2' satisfying the following conditions:

1. Every atomic leaf is mapped to an atomic leaf with the same label.
2. If W_2' contains y_c as a leaf and thus neither γ_1 nor γ_2 is found in W_2', then y_c is mapped to a node v with $\alpha_+(v) > 0$.
3. If only one of γ_1 and γ_2 exists in W_2', say γ_1, then it is mapped to a node v with $\alpha_1(v) > 0$.
4. If both γ_1 and γ_2 exist in W_2', then any of the following cases is permitted:
   - γ_1 and γ_2, respectively, are mapped to a compressed leaf z and its sibling $\bar{z}$ in W_1' with $\beta(z) > 0$.
   - γ_1 and γ_2, respectively, are mapped to a compressed leaf z and the accompanying auxiliary leaf $\bar{\bar{z}}$ in W_1' with $\beta_{1 \succ 2}(z) > 0$, or the leaves $\bar{\bar{z}}$ and z in W_1' with $\beta_{2 \succ 1}(z) > 0$.
   - γ_1 and γ_2, respectively, are mapped to two leaves or internal nodes v and w with $\alpha_1(v)$, $\alpha_2(w) > 0$.

The way we measure the size of W_1' and W_2' depends on their isomorphism. For example, if y_c is mapped to some node v in W_1', then the size is the total number of atomic leaves in W_1' plus $\alpha_+(v)$. More precisely, the size of W_1' and W_2' is defined to be the total number of atomic leaves in W_1' plus the corresponding α or β values depending on the isomorphism between W_1' and W_2'. A *maximum* agreement subtree of W_1 and W_2 is one with the largest possible size. Let $\text{mast}(W_1, W_2)$ denote the size of such a subtree. The following lemma, like Lemma 4.1, is also the cornerstone of label compression.

LEMMA 4.3. $\text{mast}(T_1, T_2) = \text{mast}(W_1, W_2)$.

Proof. It follows directly from the definition of $\text{mast}(W_1, W_2)$. □

Again, $\text{mast}(W_1, W_2)$ can be computed by adapting Farach and Thorup's rooted subtree algorithm [11]. The time complexity is stated in the following lemma. Let $n = \max\{|W_1|, |W_2|\}$ and $N = \max\{|T_1|, |T_2|\}$.

LEMMA 4.4. *Suppose that all the auxiliary information of W_1 has been given. Then we can compute $\text{mast}(W_1, W_2)$ in $O(n^{1.5} \log N)$ time. Afterwards we can retrieve $\text{mast}(W_1^v, W_2)$ for any $v \in W$ in $O(1)$ time.*

Proof. See Appendix A. □

5. Computing $\text{mast}(U_1^x, U_2)$—proof of Lemma 3.2. At a high level, we first apply label compression to the input instance (U_1^x, U_2). We then reduce the problem

```
/* W is a rooted tree with compressed leaves. X is unrooted with shrunk leaves. */

mast(W, X)

    let y be a separator of X;
    val = mast(W, X^y);
    if (X has at most one shrunk leaf) or (y lies between the two shrunk leaves)
    then
        new_subproblem(W, X, y);
        for each (W_i, X_i), val_i = mast(W_i, X_i);
    else
        let y' be the node on the path between the two shrunk leaves that is the closest to
        y;
        val = mast(W, X^{y'});
        new_subproblem(W, X, y');
        for each (W_i, X_i), set val_i = mast(W_i, X_i);
    return max{val, max_{i=1}^b val_i};


/* Generate new subproblems {(W_1, X_1), ..., (W_b, X_b)}. */

new_subproblem(W, X, y)

    let v_1, ..., v_b be the neighbors of y in X;
    for all i ∈ [1, b]
        let X_i be the unrooted tree formed by shrinking the subtree X^{v_i y} into a shrunk leaf;
        let W_i be the rooted tree formed by compressing W with respect to X^{v_i y};
    compute and store the auxiliary information in W_i for all i ∈ [1, b];
```

FIG. 4. *Algorithm for computing* mast(W, X).

to a number of smaller subproblems (W, X), each of which is similar to (U_1^x, U_2) and is solved recursively. For each (W, X) generated, X is a subtree of U_2 with at most two shrunk leaves, and W is a label compression of U_1^x with respect to some rooted subtrees of U_2 that are represented by the shrunk leaves of X. Also, W and X contain the same number of atomic leaves.

5.1. Recursive computation of mast(W, X). Our subtree algorithm initially sets $W = U_1^x$ and $X = U_2$. In general, $W = U_1^x \otimes R$ and $X = U_2 \ominus R$, or $W = U_1^x \otimes (R, R')$ and $X = U_2 \ominus (R, R')$ for some rooted subtrees R and R' of U_2. If W or X has at most two nodes, then mast(W, X) can easily be computed in linear time. Otherwise, both W and X each have at least three nodes. Let $N = \max\{|U_1|, |U_2|\}$ and $n = \max\{|W|, |X|\}$. Our algorithm first finds a separator y of X and computes mast(W, X) for the following two cases. The output is the larger of the two cases. Figure 4 outlines our algorithm.

Case 1. mast$(W, X) = $ mast(W, X^y). We root X at y and evaluate mast(W, X^y). By Lemma 4.4, this takes $O(n^{1.5} \log N)$ time.

Case 2. mast$(W, X) = $ mast(W, X^z) for some internal node $z \neq y$. We compute $\max\{\text{mast}(W, X^z) \mid z \text{ is an internal node and } z \neq y\}$ by solving a set of subproblems $\{\text{mast}(W_1, X_1), ..., \text{mast}(W_b, X_b)\}$ such that their total size is n and $\max\{\text{mast}(W, X^z) \mid z \text{ is an internal node and } z \neq y\} = \max\{\text{mast}(W_i, X_i) \mid i \in [1, b]\}$. Moreover, our algorithm enforces the following properties:

- If X contains at most one shrunk leaf, every subproblem generated has size at most half that of X.
- If X has two shrunk leaves, at most one subproblem (W_{i_o}, X_{i_o}) has size

greater than half that of X, but X_{i_o} contains only one shrunk leaf. Thus, in the next recursion level, every subproblem spawned by (W_{i_o}, X_{i_o}) has size at most half that of X.

To summarize, whenever the recursion goes down by two levels, the size of a subproblem reduces by half.

The subproblems $\mathrm{mast}(W_1, X_1), \ldots, \mathrm{mast}(W_b, X_b)$ are formally defined as follows. Assume that the separator y has b neighbors in X, namely, $v_1, \ldots, v_b$. For each $i \in [1, b]$, let C_i be the connected component in $X - \{y\}$ that contains v_i. The size of C_i is at most half that of X. Intuitively, we would like to shrink the subtree $X^{v_i y}$ into a leaf, producing a smaller unrooted tree X_i. We first consider the simple case where X has at most one shrunk leaf. Then no C_i contains more than one shrunk leaf.

If C_i contains no shrunk leaf, then X_i contains only one shrunk leaf representing the subtree $X^{v_i y}$. Note that $X^{v_i y}$ corresponds to the subtree $U_2^{v_i y}$ in U_2 and $X_i = U_2 \ominus U_2^{v_i y}$. Let $W_i = U_1^x \otimes U_2^{v_i y}$.

If C_i contains one shrunk leaf γ_1, then X_i contains γ_1 as well as a new shrunk leaf representing the subtrees $X^{v_i y}$. The two subtrees are label-disjoint. Again, $X^{v_i y}$ corresponds to the subtree $U_2^{v_i y}$ in U_2. Assume that γ_1 corresponds to a subtree $U_2^{v' y'}$ in U_2. Then $X_i = U_2 \ominus (U_2^{v' y'}, U_2^{v_i y})$. Let $W_i = U_1^x \otimes (U_2^{v' y'}, U_2^{v_i y})$.

We now consider the case where X itself already has two shrunk leaves γ_1 and γ_2. If y lies on the path between γ_1 and γ_2, then no C_i contains more than one shrunk leaf and we define the smaller problem instances (W_i, X_i) as above. Otherwise, there is a C_i containing both γ_1 and γ_2. X_i as defined contains three compressed leaves, violating our requirement. In this case, we replace y with the node y' on the path between γ_1 and γ_2, which is the closest to y. Now, to compute $\mathrm{mast}(W, X)$, we consider the two cases depending on whether the root of W is mapped to y' or not. Again, we first compute $\mathrm{mast}(W, X^{y'})$. Then, we define the connected components C_i and the smaller problem instances (W_i, X_i) with respect to y'. Every X_i has at most two compressed leaves, but y' may not be a separator and we cannot guarantee that the size of every subproblem is reduced by half. However, there can exist only one connected component C_{i_o} with size larger than half that of X. Indeed, C_{i_o} is the component containing y. In this case, both γ_1 and γ_2 are not inside C_{i_o}, and X_{i_o} as defined contains only one compressed leaf. Thus, the subproblems that $\mathrm{mast}(W_{i_0}, X_{i_0})$ spawns in the next recursion level each have size of at most half that of (W, X).

With respect to y or y', computing the topology of all X_i and W_i from X and W is straightforward; see section 5.2. Computing the auxiliary information in all W_i efficiently requires some intricate techniques, which are detailed in sections 6 and 7.

5.2. Computing the topology of compressed trees. The topology of X_i can be constructed from X by replacing the subtree $X^{v_i y}$ of X with a shrunk leaf. Let J and J_i be the sets of labels in X and X_i, respectively. For the trees W_i, recall that the definitions of W and the trees W_i are based on affixing some nodes to the trees $U_1^x | J$ and $U_1^x | J_i$, respectively. Observe that $W | J$ and $U_1^x | J$ have the same topology. Moreover, $W | J_i = (W | J) | J_i$ and $U_1^x | J_i = (U_1^x | J) | J_i$. Thus, $W | J_i$ and $U_1^x | J_i$ have the same topology. We can obtain $U_1^x | J_i$ by constructing $W | J_i$. Note that $J = \bigcup_{1 \le i \le b} J_i$ and all the label sets J_i are disjoint. We can construct all the trees $W | J_i$ from W in $O(n)$ time [7, 10]. Next, we show how to construct W_i from $W | J_i$ in time linear in the size of $W | J_i$. We only detail the case where X_i consists of two shrunk leaves. The case for one shrunk leaf is similar. The following procedure is derived directly from the definition of the compression of U_1^x with respect to two subtrees.

X^y R_i

FIG. 5. *The structures of X^y and R_i.*

Let v be any node of $W|J_i$. If v is not the root, let u be the parent of v in $W|J_i$.
- If $\mathcal{A}(U_1^x, L-J_i, v)$ is nonempty or equivalently the degree of v in U_1^x is different from its degree in $W|J_i$, then attach auxiliary leaves z and $\bar{z}$ to v.
- If $\mathcal{H}(U_1^x, L-J_i, v)$ is nonempty or equivalently u is not the parent of v in U_1^x, then create a path between u and v consisting of two nodes p and $\bar{p}$, attach auxiliary leaves z and $\bar{z}$ to p, and attach $\bar{\bar{z}}$ to $\bar{p}$.

5.3. Time complexity of computing $\mathrm{mast}(W, X)$.

LEMMA 5.1. *We can compute $\mathrm{mast}(W, X)$ in $O(n^{1.5} \log N)$ time.*

Proof. Let $T(n)$ be the computation time of $\mathrm{mast}(W, X)$. The computation is divided into two cases. Case 1 of section 5.1 takes $O(n^{1.5} \log N)$ time. For Case 2, a set of subproblems $\{\mathrm{mast}(W_i, X_i) \mid i \in [1, b]\}$ is generated. As to be shown in sections 6 and 7, the time to prepare all these subproblems is also $O(n^{1.5} \log N)$. These subproblems, except possibly one, are each of size less than $n/2$. For the exceptional subproblem, say, $\mathrm{mast}(W_l, X_l)$, its computation is again divided into two cases. One case takes $O(n^{1.5} \log N)$ time. For the other case, another set of subproblems is generated in $O(n^{1.5} \log N)$ time. This time every such subproblem has size less than $n/2$. Let Σ be the set of all the subproblems generated in both steps. The total size of the subproblems in Σ is at most n, and

$$T(n) = O(n^{1.5} \log N) + \sum_{\mathrm{mast}(W', X') \in \Sigma} T(|X'|).$$

It follows that $T(n) = O(n^{1.5} \log N)$. $\square$

By letting $W = U_1^x$ and $X = U_2$, we have proved Lemma 3.2. What remains is to show how to compute the auxiliary information stored in all W_i from (W, X) in $O(n^{1.5} \log N)$ time. Note that X contains at most two shrunk leaves. Depending on the number of shrunk leaves in X, we divide our discussion into sections 6 and 7.

6. Auxiliary information for X with no shrunk leaf. The case of X containing no shrunk leaf occurs only when the algorithm starts, i.e., $W = U_1^x$, $X = U_2$, and $N = n$. The subproblems $\mathrm{mast}(W_1, X_1), \ldots, \mathrm{mast}(W_b, X_b)$ spawned from (W, X) are defined by an internal node y in X, which is adjacent to the nodes $v_1, \ldots, v_b$. Let R_i and $\bar{R}_i$ denote the rooted subtrees $X^{v_i y}$ and $X^{y v_i}$, respectively. Note that the rooted tree X^y is composed of the subtrees $\bar{R}_1, \ldots, \bar{R}_b$. Also, $W_i = W \otimes R_i$ and $X_i = W \ominus R_i$. The total size of all $\bar{R}_i$ is at most n. Furthermore, each R_i is X^y with $\bar{R}_i$ removed; see Figure 5. This section discusses how to compute the auxiliary information required by each W_i in $O(n^{1.5} \log N)$ time.

6.1. Auxiliary information in the compressed leaves of W_i. Consider any compressed leaf v in W_i. Let S_v denote the set of subtrees from which v is compressed. Then, the auxiliary information to be stored in v is

(6.1) $$\alpha(v) = \max\{\text{mast}(W^z, R_i) \mid W^z \in S_v\}.$$

Observe that for any $W^z \in S_v$, W^z contains no labels outside R_i. Thus, $\text{mast}(W^z, R_i) = \text{mast}(W^z, X^y)$ and we can rewrite (6.1) as

$$\alpha(v) = \max\{\text{mast}(W^z, X^y) \mid W^z \in S_v\}.$$

We use the rooted subtree algorithm of [11] to compute $\text{mast}(W, X^y)$ in $O(n^{1.5} \log N)$ time. Then, we can retrieve the value of $\text{mast}(W^z, X^y)$ for any node $z \in W$ in $O(1)$ time. To compute $\max\{\text{mast}(W^z, X^y) \mid W^z \in S_v\}$ efficiently, we assume that for any node $u \in W$, the subtrees attached to u are numbered consecutively, starting from 1. We consider a preprocessing for efficient retrieval of the following types of values:

- for some node $u \in W$ and some interval $[a, b]$, $\max\{\text{mast}(W^z, X^y) \mid W^z$ is a subtree attached to u and its number falls in $[a, b]\}$;
- for some path P of W, $\max\{\text{mast}(W^z, X^y) \mid W^z$ is a subtree attached to some node in $P\}$.

LEMMA 6.1. *Assume that we can retrieve* $\text{mast}(W^z, X^y)$ *for any* $z \in W$ *in* $O(1)$ *time. Then we can preprocess* W *and* X *and construct additional data structures in* $O(n \log^* n)$ *time so that any value of the above types can be retrieved in* $O(1)$ *time.*

Proof. We adapt preprocessing techniques for on-line product queries in [3]. □

With the preprocessing stated in Lemma 6.1, we can determine $\alpha(v)$ as follows. Note that S_v is either a subset of the subtrees attached to a node u in W or the set of subtrees attached to nodes on a particular path in W. In the former case, u is also a parent of v and S_v is partitioned into at most $d_u + 1$ intervals where d_u is the degree of u in W_i. From Lemma 6.1, $\alpha(v)$ can be found in $O(d_u + 1)$ time. Similarly, for the latter case, $\alpha(v)$ can be found in $O(1)$ time. Thus, the compressed leaves in W_i are processed in $O(|W_i|)$ time. Summing over all W_i, the time complexity is $O(n)$. Therefore, the overall computation time for preprocessing and finding auxiliary information in the leaves of all W_i is $O(n^{1.5} \log N)$.

6.2. Auxiliary information in the internal nodes of W_i. Consider any internal node v in W_i with $i \in [1, b]$. Our goal is to compute the auxiliary information $\alpha(v) = \text{mast}(W^v, R_i)$. Note that R_i may be of size $\Theta(n)$, and even computing one particular $\text{mast}(W^v, R_i)$ already takes $O(n^{1.5} \log N)$ time. Fortunately, these R_i are very similar. Each R_i is X^y with $\overline{R}_i$ removed. Exploiting this similarity and using the algorithm in section 2 for all-cavity matchings, we can perform an $O(n^{1.5} \log N)$-time preprocessing so that we can retrieve $\text{mast}(W^v, R_i)$ for any internal node v in W and $i \in [1, b]$ in $O(\log^2 n)$ time. Therefore, it takes $O(|W_i| \log^2 n)$ time to compute $\alpha(v)$ for all internal nodes v of one particular W_i, and $O(n \log^2 n)$ time for all W_i. The $O(n^{1.5} \log N)$-time preprocessing is detailed as follows.

First, note that if we remove y from X^y, the tree would decompose into the subtrees $\overline{R}_1, \ldots, \overline{R}_b$. Thus, the total size of all $\overline{R}_i$ is at most n. The next lemma suggests a way to retrieve efficiently $\text{mast}(W^v, \overline{R}_i)$ and $\max\{\text{mast}(W^v, \overline{R}_j) \mid j \in I\}$ for any $v \in W$ and and $I \subseteq [1, b]$.

LEMMA 6.2. *We can compute* $\text{mast}(W, \overline{R}_i)$ *for all* $i \in [1, b]$ *in* $O(n^{1.5} \log N)$ *time. Then, we can retrieve* $\text{mast}(W^v, \overline{R}_i)$ *for any node* v *in* W *and* $i \in [1, b]$ *in* $O(\log n)$ *time. Furthermore, we can build a data structure to retrieve* $\max\{\text{mast}(W^v, \overline{R}_j) \mid j \in I\}$ *for any* $v \in W$ *and* $I \subseteq [1, b]$ *in* $O(\log^2 n)$ *time.*

Proof. This lemma follows from the rooted subtree algorithm and related data structures in [11]. □

Below, we give a formula to compute $\text{mast}(W^v, R_i)$ efficiently. For any $z \in W$ and $i \in [1, b]$, let $\text{r-mast}(W^z, R_i)$ denote the maximum size among all the agreement subtrees of W^z and R_i in which z is mapped to the root of R_i.

LEMMA 6.3.

$$\text{mast}(W^v, R_i) = \max \left\{ \begin{array}{l} \max\{\text{mast}(W^v, \overline{R}_j) \mid j \in [1, b], j \neq i\}; \\ \max\{\text{r-mast}(W^z, R_i) \mid z \in W^v\}. \end{array} \right.$$

Proof. Observe that $\text{mast}(W^v, R_i) = \text{mast}(W^z, R_i) = \text{r-mast}(W^z, R_i)$ if in some maximum agreement subtree of W^v and R_i, the root of R_i is mapped to some node z in W^v. On the other hand, $\text{mast}(W^v, R_i) = \text{mast}(W^v, \overline{R}_j)$ for some $j \neq i$ if in some maximum agreement subtree of W^v and R_i, the root of R_i is not mapped to any node z in W^v. □

By Lemma 6.3, we decompose the computation of $\text{mast}(W^v, R_i)$ into two parts. The value $\max\{\text{mast}(W^v, \overline{R}_j) \mid j \in [1, b], j \neq i\}$ is determined by answering two queries $\max\{\text{mast}(W^v, \overline{R}_j) \mid j \in [1, i-1]\}$ and $\max\{\text{mast}(W^v, \overline{R}_j) \mid j \in [i+1, b]\}$ in $O(\log^2 n)$ time by Lemma 6.2. The computation of $\max\{\text{r-mast}(W^z, R_i) \mid z \in W^v\}$ makes use of a maximum weight matching of some bipartite graph as follows.

Let $\text{Ch}(z)$ denote the set of children of a node z in a tree. Let $G_{z,i} \subseteq \text{Ch}(z) \times \{\overline{R}_1, \ldots, \overline{R}_{i-1}, \overline{R}_{i+1}, \ldots, \overline{R}_b\}$ be a bipartite graph where $w \in \text{Ch}(z)$ is connected to $\overline{R}_j$ if and only if $\text{mast}(W^w, \overline{R}_j) > 0$. Such an edge has weight $\text{mast}(W^w, \overline{R}_j) \leq N$.

FACT 6.4 (see [11]). *If the root of R_i is mapped to z in some maximum agreement subtree of W^z and R_i, then a maximum weight matching of $G_{z,i}$ consists of at least two edges, and $\text{mwm}(G_{z,i}) = \text{r-mast}(W^z, R_i)$.*

Note that if a maximum weight matching of $G_{z,i}$ consists of one edge, it corresponds to an agreement subtree of W^z and R_i in which the root of R_i is not mapped to any node in W^z. Thus, it is possible that $\text{mwm}(G_{z,i}) > \text{r-mast}(W^z, R_i)$. Nevertheless, in this case we are no longer interested in the exact value of $\text{r-mast}(W^z, R_i)$ since in a maximum agreement subtree of W^z and R_i, the root of R_i is not mapped to any node in W^z. In fact, Lemma 6.3 can be rewritten with the $\text{r-mast}(W^z, R_i)$ replaced by $\text{mwm}(G_{z,i})$. Furthermore, since $G_{z,1}, G_{z,2}, \ldots, G_{z,b}$ are very similar, the weights of a maximum weight matching cannot be all distinct.

LEMMA 6.5. *At least $b - d_z$ of $\text{mwm}(G_{z,1}), \text{mwm}(G_{z,2}), \ldots, \text{mwm}(G_{z,b})$ have the same value, where d_z denotes the degree of z in W.*

Proof. Consider the bipartite graph $K \subseteq \text{Ch}(z) \times \{\overline{R}_1, \ldots, \overline{R}_b\}$ in which a node $w \in \text{Ch}(z)$ is connected to $\overline{R}_i$ if and only if $\text{mast}(W^w, \overline{R}_i) > 0$. This edge is given a weight of $\text{mast}(W^w, \overline{R}_i)$. Then, every $G_{z,i}$ is a subgraph of K. Let M be a maximum weight matching of K. Observe that if an R_i is not adjacent to any edge in M, then M is also a maximum weight matching of $G_{z,i}$. Since M contains at most d_z edges, there are at least $b - d_z$ trees R_i not adjacent to any edge in M and the corresponding $\text{mwm}(G_{z,i})$ have the same value. □

We next use $O(n^{1.5} \log N)$ time to find for all z in W, $\text{mwm}(G_{z,1}), \ldots, \text{mwm}(G_{z,b})$. The results are to be stored in an array A_z of dimension b for each node z, i.e., $A_z[i] = \text{mwm}(G_{z,i})$. Note that if we represent each A_z as an ordinary array, then filling these arrays entry by entry for all $z \in W$ would cost $\Omega(bn)$ time. Nevertheless, by Lemma 6.5, most of the weights $\text{mwm}(G_{z,i})$ have the same value. Thus, we store these values in sparse arrays. Like an ordinary array, any entry in a *sparse array A* can be read and modified in $O(1)$ time. In addition, we require that all the entries

in A can be initialized to a fixed value in $O(1)$ time and that all the distinct values stored in A can be retrieved in $O(m)$ time, where m denotes the number of distinct values in A. For an implementation of sparse array, see Exercise 2.12, page 71 of [2].

Before showing how to build these sparse arrays, we illustrate how they support the computation of

(6.2) $$\max\{\mathrm{mwm}(G_{z,i}) \mid z \in W^v\} = \max\{A_z[i] \mid z \in W^v\}.$$

An efficient data structure for answering such a query is given in Appendix B. Let m_z be the number of distinct values in A_z, and $m = \sum_{z \in W}(m_z + 1)$. Let $\alpha(n)$ denote the inverse Ackermann function. Appendix B shows how to construct a data structure on top of the sparse arrays A_z in $O(m\alpha(|W|))$ time such that we can retrieve for any $v \in W$ and $i \in [1, b]$ the value of $\max\{A_z[i] \mid z \in W^v\}$ in $O(\log|W|)$ time. From Lemma 6.5, $m_z \leq d_z + 1$ for all $z \in W$; thus, $m = O(|W|)$. Therefore, the data structure can be built in time $O(m\alpha(|W|)) = O(|W|\alpha(|W|)) = O(n \log n)$ and the retrieval time of (6.2) is $O(\log|W|) = O(\log n)$.

To summarize, after building all the necessary data structures, we can retrieve $\max\{\mathrm{mast}(W^v, \overline{R}_j) \mid j \in [1, b], j \neq i\}$ in $O(\log^2 n)$ time and $\max\{\mathrm{r\text{-}mast}(W^z, R_i) \mid z \in W^v\}$ in $O(\log n)$ time. Hence, for any $v \in W$ and $i \in [1, b]$, $\mathrm{mast}(W^v, R_i)$ can be computed in $O(\log^2 n)$ time.

To complete our discussion, we show below how to construct a sparse array A_z or equivalently compute the weights $\{\mathrm{mwm}(G_{z,i}) \mid i \in [1, b]\}$ efficiently. We cannot afford to examine every $G_{z,i}$ and compute $\mathrm{mwm}(G_{z,i})$ separately. Instead we build only one weighted graph $G_z \subseteq \mathrm{Ch}(z) \times \{\overline{R}_1, \ldots, \overline{R}_b\}$ as follows.

For a node z in W, the *max-child* z' of z is a child of z such that the subtree rooted at z' contains the maximum number of atomic leaves among all the subtrees attached to z. Let $\kappa(z)$ denote the total number of atomic leaves that are in W^z but not in $W^{z'}$. The edges of G_z are specified as follows.

- For any non-max-child u of z, G_z contains an edge between u and some $\overline{R}_i$ if and only if $\mathrm{mast}(W^u, \overline{R}_i) > 0$. There are at most $\kappa(z)$ such edges.
- Regarding the max-child z' of z, we put into G_z only a limited number of edges between z' and $\{\overline{R}_1, \ldots, \overline{R}_b\}$. For each $\overline{R}_i$ already connected to some non-max-child of z, G_z has an edge between z' and $\overline{R}_i$ if $\mathrm{mast}(W^{z'}, \overline{R}_i) > 0$. Among all other $\overline{R}_i$, we pick $\overline{R}_{i'}$ and $\overline{R}_{i''}$ such that $\mathrm{mast}(W^{z'}, \overline{R}_{i'})$ and $\mathrm{mast}(W^{z'}, \overline{R}_{i''})$ are the first and second largest.
- Every edge $(u, \overline{R}_i)$ in G_z is given a weight of $\mathrm{mast}(W^u, \overline{R}_i)$.

LEMMA 6.6. *For all $i \in [1, b]$, $\mathrm{mwm}(G_z - \{\overline{R}_i\}) = \mathrm{mwm}(G_{z,i})$. Furthermore, G_z can be built in $O((\kappa(z) + 1)\log^2 n)$ time.*

Proof. The fact that $\mathrm{mwm}(G_z - \{\overline{R}_i\}) = \mathrm{mwm}(G_{z,i})$ follows from the construction of G_z. Note that G_z contains $O(\kappa(z) + 1)$ edges. All edges in G_z, except $(z', \overline{R}_{i'})$ and $(z', \overline{R}_{i''})$, can be found using $O(\kappa(z))$ time. The weight of these edges can be found in $O(\kappa(z) \log n)$ time using Lemma 6.2. To identify $(z', \overline{R}_{i'})$ and $(z', \overline{R}_{i''})$, note that at most $\kappa(z)$ instances of $\overline{R}_i$ are connected to some non-max-child of z. All other $\overline{R}_i$ are partitioned into at most $\kappa(z) + 1$ intervals. For each interval, say $I \subseteq [1, b]$, by Lemma 6.2, the corresponding $\mathrm{mast}(W^{z'}, \overline{R}_i)$ which attains the maximum in the set $\{\mathrm{mast}(W^{z'}, \overline{R}_j) \mid j \in I\}$ can be found in $O(\log^2 n)$ time. Thus, by scanning all the $\kappa(z) + 1$ intervals, $\overline{R}_{i'}$ can be found in $O((\kappa(z) + 1)\log^2 n)$ time. $\overline{R}_{i''}$ can be found similarly. $\square$

Since G_z contains $O(\kappa(z) + 1)$ edges, and each edge has weight at most N, we use the Gabow–Tarjan algorithm [13] to compute $\mathrm{mwm}(G_z)$ in $O(\sqrt{\kappa(z) + 1}(\kappa(z) +$

1) $\log N$) time. Then, using our algorithm for all-cavity maximum weight matching, we can compute $\mathrm{mwm}(G_z - \{\overline{R}_i\})$ for all $i \in [1, b]$, and store the results in a sparse array A_z in the same amount of time.

Thus, all G_z with $z \in W$ can be constructed in time $\sum_{z \in W} O((\kappa(z) + 1) \log^2 n)$, which is $O(n^{1.5} \log N)$ as $\sum_{z \in W} \kappa(z) = O(n \log n)$ [9]. Given all G_z, the time for computing A_z for all $z \in W$ is $O(\sum_{z \in W}(\kappa(z) + 1)^{1.5} \log N)$.

LEMMA 6.7. $\sum_{z \in W}(\kappa(z) + 1)^{1.5} \log N = O(n^{1.5} \log N)$.

Proof. Let $T(W) = \sum_{z \in W}(\kappa(z) + 1)^{1.5} \log N$. Let P be a path starting from the root of W such that every next node is the max-child of its predecessor. Then $\sum_{z \in P} \kappa(z) \le |W| \le n$. Let $\chi(P)$ denote the set of subtrees attached to some node on P. The subtrees in $\chi(P)$ are label-disjoint and each has size at most $n/2$. Thus,

$$
\begin{aligned}
T(W) &\le \sum_{z \in P}(\kappa(z) + 1)^{1.5} \log N + \sum_{W' \in \chi(P)} T(W') \\
&\le n^{1.5} \log N + \sum_{W' \in \chi(P)} T(W') \\
&= O(n^{1.5} \log N). \quad \square
\end{aligned}
$$

7. Auxiliary information for X with one or two shrunk leaves.

7.1. X has one shrunk leaf. Consider the computation of $\mathrm{mast}(W, X)$. According to the algorithm, $\mathrm{mast}(W, X)$ will spawn b subproblems $\mathrm{mast}(W_1, X_1), \ldots, \mathrm{mast}(W_b, X_b)$, which are defined by an internal node y in X adjacent to the nodes $v_1, \ldots, v_b$. Also, for every $i \in [1, b]$, R_i and $\overline{R}_i$ denote the subtrees $X^{v_i y}$ and $X^{y v_i}$, respectively. Suppose that X has one shrunk leaf and without loss of generality, assume that the shrunk leaf of X is in $\overline{R}_b$, i.e., X_b has two shrunk leaves and all the other X_i have one shrunk leaf each. This section shows how to find the auxiliary information required by $W_1, \ldots, W_b$ in $O(n^{1.5} \log N)$ time.

LEMMA 7.1. *The auxiliary information required by $W_1, \ldots, W_{b-1}$ can be computed in $O(n^{1.5} \log N)$ time.*

Proof. Note that $\mathrm{mast}(W_1, X_1), \ldots, \mathrm{mast}(W_{b-1}, X_{b-1})$ are almost identical to the subproblems considered in section 6 in that all the X_i have exactly one shrunk leaf each. Using exactly the same approach, we can compute the auxiliary information in $W_1, \ldots, W_{b-1}$. $\square$

The remaining section focuses on computing the auxiliary information in W_b. Let γ_1 and γ_2 be the two shrunk leaves of X_b. Assume that γ_1 is also a shrunk leaf in X, and γ_2 represents R_b. Let Q^+ be the subtree obtained by connecting γ_1 and R_b together with a node. To compute the auxiliary information in W_b, we require the values $\mathrm{mast}(W^v, \gamma_1)$, $\mathrm{mast}(W^v, R_b)$, and $\mathrm{mast}(W^v, Q^+)$ for all nodes $v \in W$. These values are computed based on the following lemma.

LEMMA 7.2. $\mathrm{mast}(W^v, \gamma_1)$, $\mathrm{mast}(W^v, R_b)$, *and* $\mathrm{mast}(W^v, Q^+)$ *for all nodes $v \in W$ can be computed in $O(n^{1.5} \log N)$ time.*

Proof. By Lemma 4.2, $\mathrm{mast}(W, R_b)$ and $\mathrm{mast}(W, Q^+)$ can be computed in time $O(n^{1.5} \log N)$ and afterwards, for each node $v \in W$, $\mathrm{mast}(W^v, R_b)$ and $\mathrm{mast}(W^v, Q^+)$ can be retrieved in $O(1)$ time. For each node $v \in W$, $\mathrm{mast}(W^v, \gamma_1)$ is the auxiliary information stored at v in W and can be retrieved in $O(1)$ time. $\square$

Now, we are ready to compute the auxiliary information stored at each node $v \in W_b$. No auxiliary information is required for atomic leaves. Below, Lemmas 7.3 and 7.4 show that using $O(n)$ additional time, we can compute the auxiliary information

in internal nodes and in compressed leaves, respectively. In summary, the auxiliary information in $W_1, \ldots, W_b$ can be computed in $O(n^{1.5} \log N)$ time.

LEMMA 7.3.   *Given* $\mathrm{mast}(W^v, \gamma_1)$, $\mathrm{mast}(W^v, R_b)$, *and* $\mathrm{mast}(W^v, Q^+)$ *for all nodes* $v \in W$, *the auxiliary information stored at the internal nodes in* W_b *can be found in* $O(n)$ *time.*

Proof. Let J_b be the set of labels of the atomic leaves of W_b. An internal node v can be either an auxiliary node, a compressed node, or a node of $W|J_b$. If $v \in W|J_b$, then $v \in W$. Thus, $\alpha_1(v) = \mathrm{mast}(W^v, \gamma_1)$ and $\alpha_2(v) = \mathrm{mast}(W^v, R_b)$.

If v is a compressed node, then we need to compute $\alpha_1(v), \alpha_2(v)$ and $\alpha_+(v)$. Recall that v represents some tree path $\sigma = v_1, \ldots, v_k$ of W, where v_1 is the closest to the root, i.e., $v = v_1$. Thus, $\alpha_1(v) = \mathrm{mast}(W^{v_1}, \gamma_1)$, $\alpha_2(v) = \mathrm{mast}(W^{v_1}, R_b)$, and $\alpha_+(v) = \mathrm{mast}(W^{v_1}, Q^+)$.

Thus, $O(n)$ time is sufficient for finding the auxiliary information stored at every internal node of W_b.   ☐

LEMMA 7.4.   *Given* $\mathrm{mast}(W^v, \gamma_1)$, $\mathrm{mast}(W^v, R_b)$, *and* $\mathrm{mast}(W^v, Q^+)$ *for all nodes* $v \in W$, *the auxiliary information stored at the compressed leaves in* W_b *can be found in* $O(n)$ *time.*

Proof. If v is a compressed leaf in W, v's parent u must not be an auxiliary node. Depending on whether u is a compressed node, we have two cases.

Case A. u is not a compressed node. We need to compute $\alpha_1(v), \alpha_2(v), \alpha_+(v), \beta(v)$. Note that u is also in W. When W_b is constructed from W, some of the subtrees of W attached to u are replaced by v and no longer exist in W_b. Let $W^{p_1}, \ldots, W^{p_k}$ be these subtrees. Observe that both v and $W^{p_1}, \ldots, W^{p_k}$ represent the same set of subtrees in T_1. Thus,

- $\alpha_1(v) = \max\{\mathrm{mast}(W^{p_i}, \gamma_1) \mid 1 \leq i \leq k\}$;
- $\alpha_2(v) = \max\{\mathrm{mast}(W^{p_i}, R_b) \mid 1 \leq i \leq k\}$;
- $\alpha_+(v) = \max\{\mathrm{mast}(W^{p_i}, Q^+) \mid 1 \leq i \leq k\}$;
- $\beta(v) = \max\{\mathrm{mast}(W^{p_i}, \gamma_1) + \mathrm{mast}(W^{p_j}, R_b) \mid 1 \leq i \neq j \leq k\}$.

These four values can be found in $O(k)$ time. Since $W^{p_1}, \ldots, W^{p_k}$ are subtrees attached to u in W, k is at most the degree of u in W. Moreover, the sum of the degrees of all internal nodes of W is $O(n)$. Therefore, $O(n)$ time suffices to compute the auxiliary information for all the compressed leaves in W_b whose parents are not compressed node.

Case B. u is a compressed node. We need to compute $\alpha_1(v), \alpha_2(v), \alpha_+(v), \beta(v)$, $\beta_{1 \succ 2}(v)$, and $\beta_{1 \succ 2}(v)$. Note that u is compressed from a tree path $p_1, \ldots, p_k$ in W, where p_1 is the closest to the root. Moreover, v is compressed from the subtrees hanging between p_1 and p_k. For every $i \in [1, k]$, let $\mathcal{T}_i$ be the set of subtrees of W attached to p_i that are compressed into v. Both v and the subtrees in $\cup_{1 \leq i \leq k} \mathcal{T}_i$ represent the same set of subtrees in T_1. The auxiliary information stored at v can be expressed as follows.

- $\alpha_1(v) = \max\{\mathrm{mast}(W^q, \gamma_1) \mid W^q \in \mathcal{T}_i \text{ for some } i \in [1, k]\}$.
- $\alpha_2(v) = \max\{\mathrm{mast}(W^q, R_b) \mid W^q \in \mathcal{T}_i \text{ for some } i \in [1, k]\}$.
- $\alpha_+(v) = \max\{\mathrm{mast}(W^q, Q^+) \mid W^q \in \mathcal{T}_i \text{ for some } i \in [1, k]\}$.
- $\beta(v) = \max_{1 \leq i \leq k}[\max\{\mathrm{mast}(W^q, \gamma_1) + \mathrm{mast}(W^{q'}, R_b) \mid W^q, W^{q'} \in \mathcal{T}_i\}]$.
- $\beta_{1 \succ 2}(v) = \max_{1 \leq j < i \leq k}[\max\{\mathrm{mast}(W^q, \gamma_1) \mid W^q \in \mathcal{T}_i\} + \max\{\mathrm{mast}(W^{q'}, R_b) \mid W^{q'} \in \mathcal{T}_j\}]$.
- $\beta_{2 \succ 1}(v) = \max_{1 \leq j < i \leq k}[\max\{\mathrm{mast}(W^q, R_b) \mid W^q \in \mathcal{T}_i\} + \max\{\mathrm{mast}(W^{q'}, \gamma_1) \mid W^{q'} \in \mathcal{T}_j\}]$.

These values can be found in $O(\sum_{1 \leq i \leq k} d_{p_i})$ time, where d_{p_i} is the degree of p_i in W.

Thus, the auxiliary information for every compressed leaf of W_b, whose parents are compressed nodes, can be computed in $O(n)$ time.   ☐

7.2. X has two shrunk leaves. Recall that the subproblems $\mathrm{mast}(W_1, X_1)$, $\ldots, \mathrm{mast}(W_b, X_b)$ are spawned from $\mathrm{mast}(W, X)$. This section considers the case where X has two shrunk leaves. Without loss of generality, assume that the two shrunk leaves are in $\overline{R}_{b-1}$ and $\overline{R}_b$, respectively. Then, $X_1, \ldots, X_{b-2}$ each have one shrunk leaf. X_{b-1} and X_b each have two shrunk leaves. Below, we show how to compute the auxiliary informations of $W_1, \ldots, W_b$ in $O(n^{1.5} \log N)$ time.

LEMMA 7.5. *The auxiliary information required by $W_1, \ldots, W_{b-2}$ can be computed in $O(n^{1.5} \log N)$ time.*

Proof. The proof of this lemma is the same as that of Lemma 7.1.   ☐

For the remaining subproblems $\mathrm{mast}(W_{b-1}, X_{b-1})$ and $\mathrm{mast}(W_b, X_b)$, both X_{b-1} and X_b have two shrunk leaves. By symmetry, it suffices to discuss the computation of $\mathrm{mast}(W_b, X_b)$ only. Lemma 7.6 shows that the auxiliary information in W_b can be computed in $O(n^{1.5} \log N)$ time. Therefore, the auxiliary information in $W_1, \ldots, W_b$ can be computed in $O(n^{1.5} \log N)$ time.

LEMMA 7.6. *The auxiliary information in W_b can be computed in $O(n^{1.5} \log N)$ time.*

Proof. Let γ_1 and γ_2 be the two shrunk leaves of X_b. Assume that γ_1 is also a shrunk leaf in X and γ_2 represents R_b, i.e., γ_2 represents the subtree $U_2^{v_b y}$ of T_1. Let Q^+ be the subtree obtained by connecting γ_1 and R_b. By the same argument as in Lemmas 7.3 and 7.4, the auxiliary information in W_b can be computed based on the values $\mathrm{mast}(W^v, \gamma_1)$, $\mathrm{mast}(W^v, R_b)$, and $\mathrm{mast}(W^v, Q^+)$ for all $v \in W$. The value $\mathrm{mast}(W^v, \gamma_1)$ can be found in W. The values $\mathrm{mast}(W^v, R_b)$ and $\mathrm{mast}(W^v, Q^+)$ for all $v \in W$ can be retrieved in $O(1)$ time after $\mathrm{mast}(W, R_b)$ and $\mathrm{mast}(W, Q^+)$ are computed in $O(n^{1.5} \log N)$ time based on Lemma 4.4. Then the auxiliary information in W_b can be computed in $O(n)$ time.   ☐

8. Extension. We have presented an $O(N^{1.5} \log N)$-time algorithm for computing a maximum agreement subtree of two unrooted evolutionary trees of at most N nodes each. This algorithm can be modified slightly to compute a maximum agreement subtree for two mixed trees M_1 and M_2.

For a mixed tree M, a node ℓ is *consistent* with a node u if the directed edges on the path between u and ℓ all point away from u. Let M^u be the rooted tree constructed by assigning u in M as the root and removing the nodes of M inconsistent with u. Given two mixed tree M_1 and M_2, we define a maximum agreement subtree of M_1 and M_2 to be the one with the largest number of labels among the maximum agreement subtree of M_1^u and M_2^v over all nodes $u \in M_1$ and $v \in M_2$. That is,

$$\mathrm{mast}(M_1, M_2) = \max\{\mathrm{mast}(M_1^u, M_2^v) \mid u \in M_1, v \in M_2\}.$$

As in the unrooted case, to compute $\mathrm{mast}(M_1, M_2)$, we find a separator y of M_1 and compute $\mathrm{mast}(M_1^y, M_2)$. However, we need to delete the nodes of M_1 not in M_1^y. When computing $\mathrm{mast}(M_1^y, M_2)$, we construct some rooted subtrees of M_2. Again, we delete the nodes of M_2 not in these rooted subtrees. Such deletions are straightforward and do not increase the time complexity of computing $\mathrm{mast}(M_1, M_2)$. Thus, $\mathrm{mast}(M_1, M_2)$ can be computed in $O(N^{1.5} \log N)$ time.

Appendix A. Computing $\mathrm{mast}(W_1, W_2)$. Let T_1 and T_2 be rooted evolutionary trees. Let R_1 and R_2 be two label-disjoint rooted subtrees of T_2. Let $W_1 = T_1 \otimes (R_1, R_2)$ and $W_2 = T_2 \ominus (R_1, R_2)$. This section shows that $\mathrm{mast}(W_1, W_2)$

can be computed as if W_1 and W_2 were ordinary rooted evolutionary trees [9, 11, 20] with some special procedures on handling compressed and shrunk leaves. Note that the case where W_1 and W_2 are compressed and shrunk with respect to a subtree can be treated as the special case where R_1 is empty.

LEMMA A.1. *We can compute* $\mathrm{mast}(W_1, W_2)$ *in* $O(n^{1.5} \log N)$ *time, where* $n = \max\{|W_1|, |W_2|\}$ *and* $N = \max\{|T_1|, |T_2|\}$. *Afterwards, we can retrieve* $\mathrm{mast}(W_1^u, W_2)$ *for any node* u *of* W_1 *in* $O(1)$ *time.*

Proof. We adopt the framework of Farach and Thorup's algorithm [11], which is essentially a sparsified dynamic programming based on the following formula. For any internal nodes u of W_1 and v of W_2,

$$
(A.1) \quad \mathrm{mast}(W_1^u, W_2^v) = \max \begin{cases} \max\{\mathrm{mast}(W_1^x, W_2^v) \mid x \text{ is a child of } u\}; \\ \max\{\mathrm{mast}(W_1^u, W_2^y) \mid y \text{ is a child of } v\}; \\ \text{r-mast}(W_1^u, W_2^v), \end{cases}
$$

where r-$\mathrm{mast}(W_1^u, W_2^v)$ denotes the maximum size of all the agreement subtrees of W_1^u and W_2^v in which u is mapped to v.

Our algorithm differs from Farach and Thorup's algorithm in the way how each individual $\mathrm{mast}(W_1^u, W_2^v)$ is computed. When W_1 and W_2 are ordinary evolutionary trees, each $\mathrm{mast}(W_1^u, W_2^v)$ is found by computing a maximum weight matching of some bipartite graph, and it takes $O(n^{1.5} \log n)$ time to compute $\mathrm{mast}(W_1, W_2)$. Below, we show that when W_1 and W_2 have compressed and shrunk leaves, each $\mathrm{mast}(W_1^u, W_2^v)$ can be found either in constant time or by computing at most two maximum weight bipartite matchings of similar graphs but with edge weights bounded by N instead of n. Thus, we can compute $\mathrm{mast}(W_1, W_2)$ using the same sparsified dynamic programming in [11]; as a by-product, we can afterwards retrieve $\mathrm{mast}(W_1^u, W_2)$ for any node u of W_1 in $O(1)$ time. The enlarged upper bound of edge weights increases the time complexity to $O(n^{1.5} \log N)$, though.

In the rest of this section, we show how each $\mathrm{mast}(W_1^u, W_2^v)$ is computed. First, we consider the case when u is a leaf. The following case analysis shows that $O(1)$ time suffices to compute $\mathrm{mast}(W_1^u, W_2^v)$.

Case 1. u is an atomic leaf. If W_2^v contains a leaf with the same label as that of u, then $\mathrm{mast}(W_1^u, W_2^v) = 1$; otherwise it equals zero.

Case 2. u is an auxiliary leaf. Then, $\mathrm{mast}(W_1^u, W_2^v) = 0$.

Case 3. u is a compressed leaf. By definition, u can only be mapped to γ_1, γ_2, or the least common ancestor y_c of γ_1 and γ_2. If W_2^v has no shrunk leaves, then $\mathrm{mast}(W_1^u, W_2^v) = 0$. If W_2^v has only one shrunk leaf, say γ_1, then $\mathrm{mast}(W_1^u, W_2^v) = \alpha_1(u)$. If W_2^v has two shrunk leaves, then W_2^v must also contain y_c and $\mathrm{mast}(W_1^u, W_2^v) = \max\{\alpha_1(u), \alpha_2(u), \alpha_+(u)\}$.

Next, we consider the case when u is an internal node. Assume that v is an atomic leaf. Then $\mathrm{mast}(W_1^u, W_2^v) = 1$ if W_1^u contains a leaf with the same label as that of v, and zero otherwise. If v is a shrunk leaf, say γ_1, then $\mathrm{mast}(W_1^u, W_2^v) = \alpha_1(u)$. It remains to consider the case when v is an internal node. Due to the nature of dynamic programming, we only need to compute r-$\mathrm{mast}(W_1^u, W_2^v)$, then we can apply (A.1) to compute $\mathrm{mast}(W_1^u, W_2^v)$. We further divide our discussion into the following three cases.

Case 1. u is an auxiliary internal node. In such case, u has only two children; one of them is an auxiliary leaf. By definition, an auxiliary leaf will not be mapped to any node in any agreement subtree of W_1^u and W_2^v; thus, there is no agreement subtree in which u is mapped to v and r-$\mathrm{mast}(W_1^u, W_2^v) = 0$.

FIG. 6. *Structure of W_1^u.*

Case 2. u is an ordinary internal node. As in [11], we first construct the bipartite graph defined as follows: Let A and B be the set of children of u and v, respectively; define $G[A, B]$ to be the bipartite graph formed by the edges $(x, y) \in A \times B$ with $\text{mast}(W_1^x, W_2^y) > 0$, and (x, y) is given a weight $\text{mast}(W_1^x, W_2^y)$.

If none of u's children is an auxiliary leaf, then r-mast$(W_1^u, W_2^v) = \text{mwm}(G[A, B])$. Otherwise, let $\bar{z}$ be the child of u which is an auxiliary. In this case, u also has a compressed child z. Other than z and $\bar{z}$, no other child of u is a compressed and auxiliary leaf. On the other hand, consider the rooted subtrees W_2^x rooted at the children x of v. If the shrunk leaves appear together in one of such subtrees, then by definition, $\bar{z}$ cannot be mapped to any shrunk leaf in any agreement subtree of W_1^u and W_2^v, and r-mast$(W_1^u, W_2^v) = \text{mwm}(G[A, B] - \{\bar{z}\})$. If the shrunk leaves appear in two different subtrees rooted at two children y_1 and y_2 of v, then

$$\text{r-mast}(W_1^u, W_2^v) = \max\left\{\text{mwm}(G[A, B] - \{\bar{z}\}), \text{mwm}(G[A, B] - \{z, \bar{z}, y_1, y_2\}) + \beta(z)\right\}.$$

Case 3. u is a compressed internal node. By definition of a compressed node, the structure of W_1^u is very restrictive—u has exactly three children z, $\bar{z}$, and an auxiliary internal node $\bar{u}$; $\bar{u}$ has two children, an auxiliary leaf $\bar{\bar{z}}$ and an uncompressed internal node w; see Figure 6. To find r-mast(W_1^u, W_2^v), we note that there are only six possible ways for how $z, \bar{z}, \bar{\bar{z}}$ are mapped to γ_1 and γ_2. We consider each of these cases and r-mast(W_1^u, W_2^v) is the maximum of the values found. We discuss only the case where γ_1 and γ_2 are mapped to z and $\bar{\bar{z}}$, respectively. The other cases can be handled similarly. Let P be the path between γ_2 and y_c. Let $S(P)$ denote the set of subtrees hanged on P. The size of the largest agreement subtrees of W_1^u and W_2^v in which γ_1 and γ_2 are mapped to z and $\bar{z}$, respectively, equals

(A.2) $$\beta_{1 \succ 2}(z) + \max\{\text{mast}(W_1^w, \tau) \mid \tau \in S(P)\}.$$

Note that using the technique in [11], we can precompute $\max\{\text{mast}(W_1^x, \tau) \mid \tau \in S(P)\}$ for all $x \in W_1$ in $O(n^{1.5} \log N)$ time. Afterwards, (A.2) can be found in constant time. $\square$

Appendix B. Preprocessing for finding $\max\{A_z[i] | z \in W^v\}$.

Let h be the number of nodes in W. Consider the h arrays A_z of dimension b where $z \in W$. Recall that m_z is the number of distinct values in A_z, and $m = \sum_{z \in W}(m_z + 1)$. This section describes an $O(m\alpha(h))$-time preprocessing, which supports finding $\max\{A_z[i] \mid z \in W^v\}$, for any $i \in [1, b]$ and any node v of W, in $O(\log h)$ time.

By definition, each A_z has at least $b - m_z$ entries storing some common value c_z. For every $i \in [1, b]$, let Γ_i be the set of nodes z where $A_z[i]$ stores a value different from c_z. Note that $\sum m_z = \sum_{1 \leq i \leq b} |\Gamma_i|$. We assume that each node of W is identified

uniquely by an integer in $[1, h]$ assigned by a preorder tree traversal [8]. For any node $v \in W$, let $\beta(v)$ be the number of proper descendents of v in W.

Based on Lemma B.1, $\max\{A_z[i] \mid z \in W^v\} = \max\{A_z[i] \mid z \in [v, v + \beta(v)]\}$ for any $v \in W$, $i \in [1, b]$. Therefore, to solve our problem, it is sufficient to give an $O(m\alpha(h))$-time preprocessing to support finding $\max\{A_z[i] \mid z \in H\}$ for any $i \in [1, b]$ and any interval $H \subseteq [1, h]$ in $O(\log h)$ time.

LEMMA B.1. *For any $v \in W$ and $i \in [1, b]$, $\max\{A_z[i] \mid z \in W^v\} = \max\{A_z[i] \mid z \in [v, v + \beta(v)]\}$.*

Proof. It is straightforward. $\square$

Our preprocessing does not work on each sequence $A_1[i], A_2[i], \ldots, A_h[i]$ directly. Instead, it first draws out useful information about the common values c_z stored in the sequences and applies a contraction technique to shorten each sequence. Then, it executes Fact B.2 on these shortened sequences.

FACT B.2 (see [3]). *Given any sequence $a_1, \ldots, a_h$ of real numbers, we can preprocess these h numbers in $O(h)$ time so that we can find the maximum of any subsequence $a_x, a_{x+1}, \ldots, a_y$ in $O(\alpha(h))$ time.*

Our preprocessing is detailed as follows. Its time complexity is $O(m\alpha(h))$ as shown in Lemma B.3.

1. For each $i \in [1, b]$, find Γ_i and arrange the integers in Γ_i in ascending order.
2. Apply Fact B.2 to the sequence $c_1, \ldots, c_h$.
3. For every nonempty $\Gamma_i = \{x_1 < \cdots < x_d\}$, compute $\beta_\ell = \max\{A_x[i] \mid x \in (x_\ell, x_{\ell+1})\}$ for every $\ell \in [1, d-1]$, and then apply Fact B.2 to the sequence $A_{x_1}[i], \beta_1, A_{x_2}[i], \ldots, \beta_{d-1}, A_{x_d}[i]$.

LEMMA B.3. *The preprocessing requires $O(m\alpha(h))$ time.*

Proof. We can examine all the entries of A_z whose values differ from c_z in $O(m_z)$ time. By examining all such entries of $A_1, \ldots, A_h$, we can construct Γ_i and arrange the integers in Γ_i in ascending order. Thus, step 1 takes $O(m)$ time. Step 2 takes $O(h) = O(m)$ time. For step 3, we first analyze the time required to process one nonempty $\Gamma_i = \{x_1 < \cdots < x_d\}$. Note that $(x_\ell, x_{\ell+1}) \cap \Gamma_i = \phi$ for every $\ell \in [1, d-1]$. Thus, $\beta_\ell = \max\{c_x \mid x \in (x_\ell, x_{\ell+1})\}$ can be computed in $O(\alpha(n))$ time using the result of step 2. Summing over all $\ell \in [1, d-1]$, computing all β_ℓ takes $O(|\Gamma_i|\alpha(h))$ time. Applying Fact B.2 to the sequence $A_{x_1}[i], \beta_1, A_{x_2}[i], \ldots, \beta_{d-1}, A_{x_d}[i]$ takes $O(|\Gamma_i|)$ time. In total, it takes $O(|\Gamma_i|\alpha(h))$ time to process one Γ_i, and step 3 takes $O(\sum |\Gamma_i|\alpha(h)) = O(m\alpha(h))$ time. Thus, the total time of our preprocessing is $O(m\alpha(h))$. $\square$

After the preprocessing, each query can be answered in $O(\log h)$ time as stated in the following lemma.

LEMMA B.4. *After the preprocessing, $\max\{A_z[i] \mid z \in H\}$ can be found in $O(\log n)$ time for any $i \in [1, b]$ and any interval $H \subseteq [1, h]$.*

Proof. Let $H = [p, q]$. A crucial step is to find $[p, q] \cap \Gamma_i$. Without loss of generality, assume $\Gamma_i \neq \phi$. To find $[p, q] \cap \Gamma_i$, we first find the smallest integer x_s in Γ_i that is greater than p, and the largest integer x_t in Γ_i that is smaller than q. Since Γ_i is sorted, we can find x_s and x_t in $O(\log |\Gamma_i|) = O(\log h)$ time. If $x_s > x_t$, then $[p, q] \cap \Gamma_i = \phi$; otherwise, $[p, q] \cap \Gamma_i$ is the set of integers between x_s and x_t in Γ_j.

If $[p, q] \cap \Gamma_i = \phi$, then $\max\{A_x[i] \mid x \in [p, q]\} = \max\{c_x \mid x \in [p, q]\}$. Because of step 1 of our preprocessing, we can find $\max\{c_x \mid x \in [p, q]\}$ in $O(\alpha(h))$ time.

If $[p, q] \cap \Gamma_i = \{x_s < x_{s+1} < \cdots < x_t\}$, then $[p, q] = [p, x_s - 1] \cup \{x_s\} \cup (x_s, x_{s+1}) \cup \cdots \cup \{x_t\} \cup [x_t + 1, q]$ and $\max\{A_z[i] \mid z \in [p, q]\}$ equals the maximum of

1. $\max\{A_x[i] \mid x \in [p, x_s - 1]\}$,

2. $\max\{A_x[i] \mid x \in \{x_s\} \cup (x_s, x_{s+1}) \cup \cdots \cup (x_{t-1}, x_t) \cup \{x_t\}\}$,
3. $\max\{A_x[i] \mid x \in [x_t + 1, q]\}$.

Note that item 2 equals the maximum of $A_{x_s}[i], \beta_s, \ldots, \beta_{t-1}, A_{x_t}[i]$, which can be computed in $O(\alpha(h))$ time after step 3 of our preprocessing. Since $\Gamma_i \cap [p, x_s - 1] = \phi$ and $\Gamma_i \cap [x_t + 1, q] = \phi$, step 2 enables us to compute items 1 and 3 in $O(\alpha(h))$ time. As a result, $\max\{A_z[i] \mid z \in [p, q]\}$ can be answered in $O(\log h)$ time. ☐

Acknowledgments. The authors thank the referees for helpful comments.

REFERENCES

[1] R. AGARWALA AND D. FERNÁNDEZ-BACA, *A polynomial-time algorithm for the perfect phylogeny problem when the number of character states is fixed*, SIAM J. Comput., 23 (1994), pp. 1216–1224.

[2] A. V. AHO, J. E. HOPCROFT, AND J. D. ULLMAN, *The Design and Analysis of Computer Algorithms*, Addison-Wesley, Reading, MA, 1974.

[3] N. ALON AND B. SCHIEBER, *Optimal Preprocessing for Answering On-Line Product Queries*, Tech. Rep. 71, The Moise and Frida Eskenasy Institute of Computer Science, Tel Aviv University, Tel Aviv, Israel, 1987.

[4] A. AMIR AND D. KESELMAN, *Maximum agreement subtree in a set of evolutionary trees: Metrics and efficient algorithms*, SIAM J. Comput., 26 (1997), pp. 1656–1669.

[5] J. L. BENTLEY, D. HAKEN, AND J. B. SAXE, *A general method for solving divide-and-conquer recurrences*, SIGACT News, 12 (1980), pp. 36–44.

[6] H. L. BODLAENDER, M. R. FELLOWS, AND T. J. WARNOW, *Two strikes against perfect phylogeny*, in Proceedings of the 19th International Colloquium on Automata, Languages, and Programming, Lecture Notes in Comput. Sci. 623, Springer-Verlag, New York, 1992, pp. 273–283.

[7] R. COLE AND R. HARIHARAN, *An $O(n \log n)$ algorithm for the maximum agreement subtree problem for binary trees*, in Proceedings of the 7th Annual ACM-SIAM Symposium on Discrete Algorithms, Atlanta, GA, 1996, pp. 323–332.

[8] T. H. CORMEN, C. L. LEISERSON, AND R. L. RIVEST, *Introduction to Algorithms*, MIT Press, Cambridge, MA, 1990.

[9] M. FARACH, T. PRZYTYCKA, AND M. THORUP, *Computing the agreement of trees with bounded degrees*, in Proceedings of the 3rd Annual European Symposium on Algorithms, Lecture Notes in Comput. Sci. 979, P. Spirakis, ed., Springer-Verlag, New York, 1995, pp. 381–393.

[10] M. FARACH AND M. THORUP, *Fast comparison of evolutionary trees*, Inform. and Computation, 123 (1995), pp. 29–37.

[11] M. FARACH AND M. THORUP, *Sparse dynamic programming for evolutionary-tree comparison*, SIAM J. Comput., 26 (1997), pp. 210–230.

[12] C. R. FINDEN AND A. D. GORDON, *Obtaining common pruned trees*, J. Classification, 2 (1985), pp. 255–276.

[13] H. N. GABOW AND R. E. TARJAN, *Faster scaling algorithms for network problems*, SIAM J. Comput., 18 (1989), pp. 1013–1036.

[14] A. V. GOLDBERG, *Scaling algorithms for the shortest paths problem*, SIAM J. Comput., 24 (1995), pp. 494–504.

[15] D. GUSFIELD, *Optimal mixed graph augmentation*, SIAM J. Comput., 16 (1987), pp. 599–612.

[16] D. GUSFIELD, *Efficient algorithms for inferring evolutionary trees*, Networks, 21 (1991), pp. 19–28.

[17] S. K. KANNAN, E. L. LAWLER, AND T. J. WARNOW, *Determining the evolutionary tree using experiments*, J. Algorithms, 21 (1996), pp. 26–50.

[18] M. Y. KAO, *Data security equals graph connectivity*, SIAM J. Discrete Math., 9 (1996), pp. 87–100.

[19] M. Y. KAO, *Multiple-size divide-and-conquer recurrences*, in Proceedings of the International Conference on Algorithms, the 1996 International Computer Symposium, National Sun Yat-Sen University, Kaohsiung, Taiwan, Republic of China, 1996, pp. 159–161. Reprinted in SIGACT News, 28 (1997), pp. 67–69.

[20] M. Y. KAO, *Tree contractions and evolutionary trees*, SIAM J. Comput., 27 (1998), pp. 1592–1616.

[21] M. Y. Kao, T. W. Lam, W. K. Sung, and H. F. Ting, *All-cavity maximum matchings*, in Proceedings of the 8th Annual International Symposium on Algorithms and Computation, Lecture Notes in Comput. Sci. 1350, H. Imai and H. W. Leong, eds., Springer-Verlag, New York, 1997, pp. 364–373.

[22] T. W. Lam, W. K. Sung, and H. F. Ting, *Computing the unrooted maximum agreement subtree in sub-quadratic time*, Nordic J. Comput. 3 (1996), pp. 295–322.

[23] T. M. Przytycka, *Transforming rooted agreement into unrooted agreement*, J. Comput. Bio., 5 (1998), pp. 333–348.

[24] M. Steel and T. Warnow, *Kaikoura tree theorems: Computing the maximum agreement subtree*, Informat. Process. Lett., 48 (1994), pp. 77–82.

[25] L. Wang, T. Jiang, and E. Lawler, *Approximation algorithms for tree alignment with a given phylogeny*, Algorithmica, 16 (1996), pp. 302–315.