

RESEARCH

Evolution through segmental duplications and losses : A Super-Reconciliation approach

Mattéo Delabre¹, Nadia El-Mabrouk^{1*}, Katharina T Huber², Manuel Lafond³, Vincent Moulton², Emmanuel Noutahi¹ and Miguel Sautie Castellanos¹

*Correspondence:

mabrouk@iro.umontreal.ca

¹Département d'informatique (DIRO), Université de Montréal, Québec, Canada

Full list of author information is available at the end of the article

Abstract

Classical gene and species tree reconciliation, used to infer the history of gene gain and loss explaining the evolution of gene families, assumes an independent evolution for each family. While this assumption is reasonable for genes that are far apart in the genome, it is not appropriate for genes grouped into syntenic blocks, which are more plausibly the result of a concerted evolution. Here, we introduce the *Super-Reconciliation* problem which consists in inferring a history of segmental duplication and loss events (involving a set of neighboring genes) leading to a set of present-day syntenies from a single ancestral one. In other words, we extend the traditional Duplication-Loss reconciliation problem of a single gene tree, to a set of trees, accounting for segmental duplications and losses. Existency of a Super-Reconciliation depends on individual gene tree consistency. In addition, ignoring rearrangements implies that existency also depends on gene order consistency. Considering the problem of reconstructing a most parsimonious Super-Reconciliation, if any, we first show that the associated decision problem is NP-hard and give an exact exponential-time algorithm to solve it. Alternatively, we show that accounting for rearrangements in the evolutionary model, but still only minimizing segmental duplication and loss events, leads to an exact polynomial-time algorithm. We finally assess time efficiency of the former exponential time algorithm for the Duplication-Loss model on simulated datasets, and give a proof of concept on the opioid receptor genes.

Keywords: Gene Tree; Reconciliation; Duplication; Loss; Synteny

1 Introduction

Gene gain and loss is known as a major force driving evolution. The classical method used for inferring these events is to reconstruct the tree of the gene family of interest and to embed it into the species phylogeny. Assuming the gene and species trees are known and correspond to the true evolution, incongruence between the two trees can be explained by gain and loss events, and “reconciling” the two trees allows recovering these events.

Tree reconciliation can be performed through different biological models of evolution, the most common being the Duplication-Loss (DL) [1, 2, 3] or Duplication-Loss and Transfer [4, 5, 6] models. Incomplete lineage sorting, i.e. imperfect segregation of alleles, can also be considered [7, 8]. While most reconciliation methods are based on the parsimony principle of minimizing the number or cost of operations, probabilistic models seeking for a reconciliation with maximum likelihood or maximum posterior probability have also been developed [9, 10, 11].

Regardless of the model, current algorithms for reconciliation take each gene family individually, assuming an independent evolution through single gene gain and loss. Although this hypothesis is reasonable for genes that are far apart in the genome, it is clearly too restrictive for those organized in syntenic blocks or paralogons, i.e. sets of homologous chromosomal regions, among one or many genomes, sharing the same genes (e.g. neuropeptide Y-family receptors [12], the Homeobox gene clusters [13, 14, 15], the FGFR fibroblast growth factor receptors [16, 17] or the genes of the opioid system [18, 19, 20]). These genes are more plausibly the result of an evolution from a common ancestral region, rather than from a set of independent gene duplications that would have converged to the same organization in different genomic regions.

The purpose of this paper is to generalize the DL reconciliation of a single gene tree, to a set of gene trees, accounting for segmental duplications and losses. As far as we know, this problem has never been considered before. The closest algorithms are DeCo [21] and DeCoStar [22] which, given a set of gene families, a set of adjacencies between genes, a set of gene trees and a species tree, compute an adjacency forest reflecting the evolution of each adjacency. However, adjacencies are taken independently, and only single duplications and losses are considered. A correction strategy that adjusts the computation of the evolutionary cost to favour co-evolution events, hence grouping seemingly individual events into single segmental ones was latter proposed in [23]. Another related problem asks for the reconciliation of a set of gene trees leading to a minimum number of duplication episodes, referring to possible whole genome duplication events, defined as sets of single duplications mapped to the same node in the species tree [24, 25]. However the considered model does not account for gene orders and duplications involving a set of neighboring genes.

Here, we consider the *DL Super-Reconciliation* problem (or Super-Reconciliation for short when no ambiguity) in which, given a set of gene families, a set of syntenies (chromosomal segments exhibiting a similar gene content), a gene tree for each gene family and a species tree, we seek an evolutionary history of the set of syntenies that is in agreement with the individual gene trees whilst minimizing the number of segmental duplications and losses. Our proposed model is a direct generalization of the reconciliation of a single gene tree. Existency of a Super-Reconciliation depends on individual gene tree consistency. In addition, ignoring rearrangements implies that existency also depends on gene order consistency. Considering the problem of reconstructing a most parsimonious Super-Reconciliation, if any, we first show that the associated decision problem is NP-hard and give an exact exponential-time algorithm to solve it. Alternatively, we show that accounting for rearrangements in the evolutionary model, but still only minimizing segmental duplications and losses, reduces to ignoring gene orders in syntenies, and leads to an exact polynomial-time algorithm.

After defining the new Super-Reconciliation model in the next section, we characterize, in Section 3, the conditions under which a Super-Reconciliation exists for a set of syntenies and a set of gene trees, and exhibit a general framework for inferring a most parsimonious DL Super-Reconciliation. We prove, in Section 4, that the associated decision problem is NP-hard. A dynamic programming algorithm for the main step of the framework is given in Section 5. Section 6 is dedicated

to an extension of the original evolutionary model accounting for rearrangements. We give a polynomial-time algorithm for finding a Super-Reconciliation, under this model, minimizing the number of segmental duplications and losses. An application on simulated datasets and a proof of concept on the genes of the opioid system are then presented in Section 7. We conclude with a discussion in Section 8.

2 Trees, Reconciliation and Problem Statement

A *string* or a *sequence* is an ordered set of characters. Given a string $X = x_1 \cdots x_n$, a *substring* of X is a consecutive set of characters from X in the same order as in X (possibly X itself), and a *subsequence* is a set of characters of X in the same order, but not necessarily consecutive in X (X is a substring and a subsequence of X). We also denote by $\text{Set}(X) = \{x_1, x_2, \dots, x_k\}$ the *range* of X , i.e. the set of all genes contained in X , without any particular order.

All trees are considered rooted. Given a tree T , we denote by $r(T)$ its root, by $V(T)$ its set of nodes and by $\mathcal{L}(T) \subset V(T)$ its leafset. We say that T is a tree for $L = \mathcal{L}(T)$. A node v is an *ancestor* of v' if v is on the path from $r(T)$ to v' ; v is the *father* of v' if it directly precedes v' on this path. In this latter case, v' is called the *child* of v . We denote by $E(T)$ the set of edges of T , where an edge is represented by its two terminal nodes (v, v') , with v being the father of v' . Two nodes v and v' are *separated* in T iff neither one is an ancestor of the other. A node is said to be *unary* if it has a single child and *binary* if it has two children. Given a node v of T , the subtree of T rooted at v is denoted $T[v]$.

A *binary tree* is a tree with all internal (i.e. non-leaf) nodes being binary. If internal nodes have one or two children, then the tree is said *partially binary*.

Creating a unary root consists of creating a new node v , a new edge $(v, r(T))$ and assigning v as the new root of T . *Grafting* a leaf w consists of subdividing an edge (v, v') of T , thereby creating a new node v'' between v and v' , then adding a leaf w with parent v'' . If W is a rooted tree, *grafting W to T* corresponds to grafting a leaf w , then replacing w by the root of W .

The *lowest common ancestor* (LCA) in T of a subset L' of $\mathcal{L}(T)$, denoted $\text{lca}_T(L')$, is the ancestor common to all nodes in L' that is the most distant from the root. The restriction $T|_{L'}$ of T to L' is the tree with leafset L' obtained from the subtree of T rooted at $\text{lca}_T(L')$ by removing all leaves that are not in L' and all unary nodes. Let T' be a tree such that $\mathcal{L}(T') = L' \subseteq \mathcal{L}(T)$. We say that T *displays* T' iff $T|_{L'}$ is label-isomorphic to T' (i.e., isomorphic with preservation of leaf labels). We also say that T is an *extension* of T' .

Species, gene and synteny trees: (See Figure 1) The *species tree* S for a set Σ of species represents an ordered set of speciation events that have led to Σ .

A *gene family* is a set Γ of genes where each gene g belongs to a given species $s(g)$ of Σ . If $\Gamma' \subseteq \Gamma$ is a subset of genes, we denote $s(\Gamma') = \{s(g) : g \in \Gamma'\}$.

A *synteny* X is an ordered sequence of genes belonging to a genome $s(X)$. We consider that genes of a synteny all belong to different gene families (tandem duplications are ignored). More precisely, let $\mathcal{F} = \{\Gamma_1, \Gamma_2, \dots, \Gamma_t\}$ be a set of gene families, and $\lambda_{\mathcal{F}} = \{(g, \Gamma) : g \in \Gamma \wedge \Gamma \in \mathcal{F}\}$ be a function. We say that an ordered sequence of genes $X = g_1 g_2 \dots g_k$ is a synteny on \mathcal{F} iff $\lambda_{\mathcal{F}}$ is well-defined for all genes of X , $\lambda_{\mathcal{F}}$ is injective, and all genes in X belong to the same species.

A *synteny family* is a set \mathcal{X} of syntenies. We say that a set \mathcal{F} of gene families are *organized into a set \mathcal{X} of syntenies* iff there is a bijection between the genes of \mathcal{F} and the genes in \mathcal{X} (each gene of \mathcal{F} belongs to exactly one synteny of \mathcal{X}).

A tree T is a *gene tree* for a gene family Γ (respec. a *synteny tree* for a synteny family \mathcal{X}) if its leafset is in bijection with Γ (respec. \mathcal{X}).

Given a gene tree T , the *corresponding synteny tree* is the tree \tilde{T} obtained from T by replacing each leaf of T by the synteny containing the considered gene.

Given a tree T (either gene tree or synteny tree), we extend the mapping s to internal nodes v of T by defining $s(v) = lca_S(\{s(l) : l \in \mathcal{L}(T[v])\})$.

An evolutionary history is represented by a *labeled tree*, where the label of a node is its corresponding event. In the case of gene families, an event is entirely determined by its type, either a duplication, a speciation or a loss. The labels of a gene tree are obtained through reconciliation, as described below.

2.1 Reconciliation

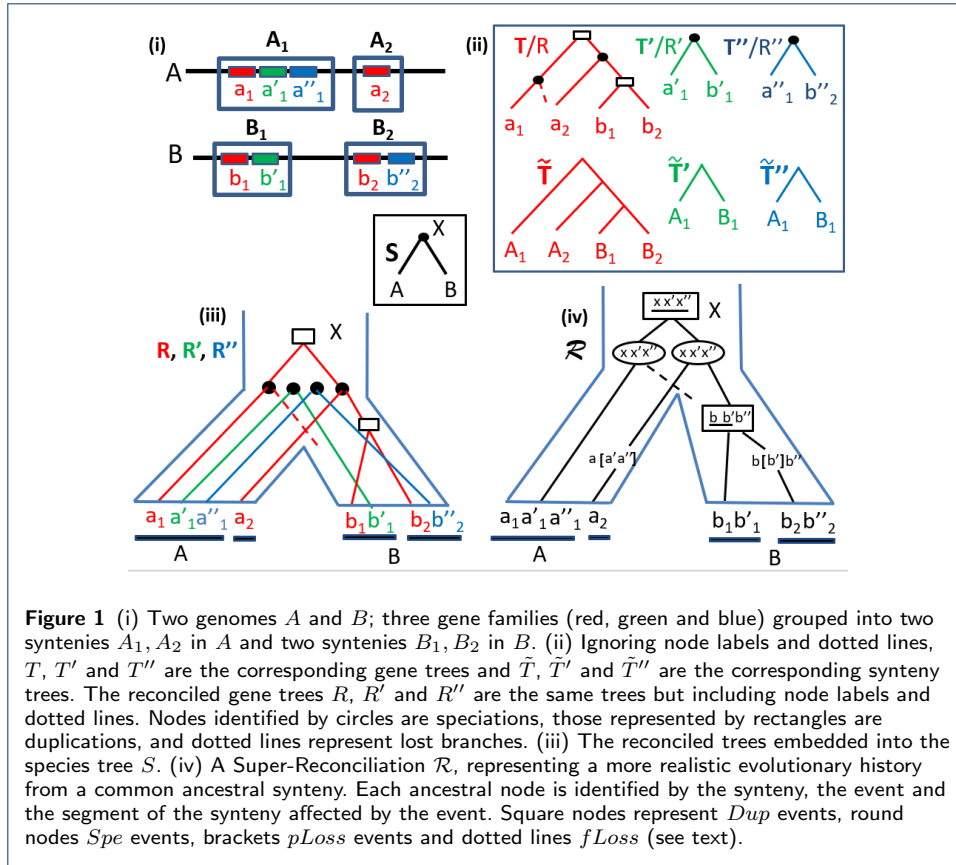
Definition 1 (Reconciled gene tree) *Let T be a binary gene tree and S be a binary species tree. A DL Reconciliation (or simply reconciliation) $R(T, S)$ of T with S is a labeled extension of T obtained by grafting new leaves satisfying: for each internal node v of $R(T, S)$ with two children v_l and v_r , either $s(v_l) = s(v_r) = s(v)$, or $s(v_l)$ and $s(v_r)$ are the two children of $s(v)$. The node v is a duplication in $s(v)$ in the former case and a speciation in the latter case. A grafted leaf on a newly created node v corresponds to a loss in $s(v)$. All other leaves are labeled by the default event “extant”.*

The cost of a reconciliation $R(T, S)$ is the number of induced duplications and losses.

Given a gene tree T and a species tree S , a *minimum reconciliation*, i.e. a reconciliation of minimum cost, is obtained from the LCA-mapping which consists in setting $s(v) = lca_S(s(\mathcal{L}(T[v])))$ for each $v \in V(T)$, and labeling each internal node v of T as a speciation if and only if $s(v_l)$ and $s(v_r)$ are separated in S , and as a duplication otherwise. Observe that in any case, if $s(v_l)$ and $s(v_r)$ are not separated, then it is impossible for v to be a speciation. We denote by *LCA-Reconciliation* the reconciliation labeled by means of the LCA-mapping.

Before extending the reconciliation concept to a set of gene trees, we need to specify an evolutionary model for syntenies. In this paper, syntenies are considered to have evolved from a single ancestral synteny through speciations (defined as for single genes), segmental duplications and segmental losses, where:

- a speciation $Spe(X, [1, l])$ acting on a synteny $X = g_1 \cdots g_l$ belonging to a genome $s(X)$ has the effect of reproducing X in the two genomes s_l and s_r children of $s(X)$ in S .
- a (segmental) duplication $Dup(X, [i, j])$ acting on a synteny X belonging to a genome $s(X)$ is an operation that copies a substring $g_i \cdots g_j$ of size $j - i + 1$ of $X = g_1 g_2 \cdots g_i \cdots g_j \cdots g_l$ somewhere else into the genome $s(X)$, creating a new *copied synteny* $X' = g'_i \cdots g'_j$ where each g'_k , for $i \leq k \leq j$ belongs to the same gene family as g_k ; we say that the copied synteny is *partial* if $[i, j] \neq [1, l]$.



- a (segmental) loss $Loss(X, [i, j])$ acting on a syntenies $X = g_1 \cdots g_i \cdots g_j \cdots g_l$ is an operation that removes a substring $g_i \cdots g_j$ of size $j - i + 1$ of X , leading to the *truncated syntenies* $X' = g_1 \cdots g_{i-1} g_{j+1} \cdots g_l$. A loss is called *full* if X' is the empty string (i.e. all genes of X are removed) and *partial* otherwise. We may denote full loss events as $fLoss$ and partial loss events as $pLoss$.

An evolutionary history of a set of synteny groups can thus be represented as a partially binary tree where leaves correspond to extant synteny groups and lost synteny groups (resulting from full losses), and each internal node v corresponds to an event $\mathcal{E}(X, [i, j])$ with $\mathcal{E} \in \{Spe, Dup, pLoss\}$ (and leaves correspond to either extant genes or $fLoss$ events). Thus, in contrast to a single gene family, a tree representing the evolution of a set of synteny groups is not only labeled by the type of event corresponding to each internal node, but also by the segment of the syntenies affected by the event (see the bottom-right tree in Figure 1). If \mathcal{E} is:

- 1 Spe , then v is a binary node with two children corresponding to synteny groups Y and Z such that $X = Y = Z$ and $s(Y)$ and $s(Z)$ being the two children of $s(X)$ in S .
- 2 Dup , then v is a binary node with two children corresponding to synteny groups X and $X' = X[i, j]$, where $s(X) = s(X')$.
- 3 $pLoss$, then v is a unary node with a child corresponding to the truncated syntenies $X' = X[1, i - 1]X[j + 1, l]$, and $s(X) = s(X')$.

The topology of a tree representing the evolution of a set of syntenies differs from that of a single gene family since the former may contain unary nodes, resulting from partial losses, while the latter only contains binary nodes.

Our goal is to infer an evolutionary history of a set of syntenies which is a reconciliation of a set of individual gene trees, formally defined below.

Definition 2 (Super-Reconciliation) *Let $\mathcal{G} = \{T_1, T_2, \dots, T_t\}$ be a set of binary gene trees for the gene families $\mathcal{F} = \{\Gamma_1, \Gamma_2, \dots, \Gamma_t\}$ organized into a set \mathcal{X} of syntenies belonging to a set Σ of taxa, and let S be a binary species tree for Σ . For each i , $1 \leq i \leq t$, let \tilde{T}_i be the synteny tree corresponding to T_i .*

A Super-Reconciliation $R(\mathcal{G}, S)$ of \mathcal{G} with S is a labeled synteny tree which is an extension of the trees \tilde{T}_i , for $1 \leq i \leq t$, representing a valid history for \mathcal{X} .

The cost of a Super-Reconciliation $R(\mathcal{G}, S)$ is the number of induced Dup, fLoss and pLoss events.

For example, the cost of the Super-Reconciliation in Figure 1 is 5. Notice that, although this cost is higher than that obtained by considering each gene family independently (cost of 3), the induced history is much more realistic as it is unlikely that independent gene duplications would have led to the same gene organization in different genomic regions.

We are now ready to state the optimization problem considered in this paper.

SUPER-RECONCILIATION problem:

Input: A set Σ of species and a species tree S for Σ ; a set of gene families $\mathcal{F} = \{\Gamma_1, \Gamma_2, \dots, \Gamma_t\}$ organized into a set of syntenies \mathcal{X} ; a set of gene trees $\mathcal{G} = \{T_1, T_2, \dots, T_t\}$ one for each family of \mathcal{F} .

Output: A Super-Reconciliation $R(\mathcal{G}, S)$ of minimum cost.

3 Existence conditions

As a synteny is represented by a gene order and can only be modified through losses (duplications create new syntenies but do not modify existing syntenies), an evolutionary history does not always exist for a set of syntenies \mathcal{X} , regardless of the trees linking them. If this holds, the syntenies are said to be *order consistent*.

In addition, in contrast to the reconciliation of a single gene tree which always exists, this is not the case for a Super-Reconciliation as different gene trees may exhibit inconsistent speciation histories for the same syntenies.

The following two subsections are dedicated to characterizing the gene order and gene tree conditions required for the existence of a Super-Reconciliation.

3.1 Consistency of gene orders

Given a set of gene families $\mathcal{F} = \{\Gamma_1, \Gamma_2, \dots, \Gamma_t\}$ organized into a set of syntenies \mathcal{X} , we define the *precedence graph* \mathcal{P} as the directed graph with n vertices, each corresponding to a gene family of \mathcal{F} , such that a directed edge (i, j) between two vertices i and j exists iff there is a synteny $X = x_1 x_2 \dots x_k$ of \mathcal{X} containing a gene in Γ_i preceding a gene in Γ_j , i.e. there is a pair $1 \leq l_1 < l_2 \leq k$ such that $x_{l_1} \in \Gamma_i$ and $x_{l_2} \in \Gamma_j$.

If \mathcal{P} is acyclic, then \mathcal{P} is a Directed Acyclic Graph (DAG). In this case, there is a topological sorting for \mathcal{P} , i.e., a linear ordering X of vertices such that for every directed edge (i, j) in \mathcal{P} , i precedes j in X . Verifying if a directed graph is acyclic and finding a topological sorting of a DAG is a classical problem solvable in linear time.

The following lemma gives necessary and sufficient conditions for a set of syntenies to be order consistent and exhibits the set of possible ancestral syntenies.

Lemma 1 (Order consistency condition) *Let $\mathcal{F} = \{\Gamma_1, \Gamma_2, \dots, \Gamma_t\}$ be a set of gene families organized into a set \mathcal{X} of syntenies. Then \mathcal{X} is order consistent iff the corresponding precedence graph \mathcal{P} is acyclic. In this case, any topological sorting for \mathcal{P} is an order consistent ancestral syntenies for \mathcal{X} .*

Proof The first part of the lemma follows from the fact that a directed graph has a topological sorting if and only if it is acyclic. The second part follows from the fact that, for any topological sorting A for \mathcal{P} and any syntenies X of \mathcal{X} , X is a subsequence of A , and thus X can be obtained from A through losses. \square

The ancestral syntenies A at the root of a Super-Reconciliation $R(\mathcal{G}, S)$ is an order on \mathcal{F} . Moreover, as the syntenies at each internal node of $R(\mathcal{G}, S)$ is obtained from A through losses, a syntenies at each internal node of $R(\mathcal{G}, S)$ should be a subsequence of A . More generally, for any two nodes v and v' of $R(\mathcal{G}, S)$, where v is an ancestor of v' , the syntenies Y at v' is a subsequence of the syntenies X at v .

3.2 Consistency of trees

A set of trees on subsets of \mathcal{X} is said *consistent* iff, for any triplet $Trp = \{X_1, X_2, X_3\}$ of disjoint elements of \mathcal{X} , all trees containing Trp as a sub-leafset exhibit the same topology for Trp .

Lemma 2 (Tree consistency condition) *Let $\mathcal{G} = \{T_1, T_2, \dots, T_t\}$ be a set of gene trees for a set of gene families organized into a set \mathcal{X} of syntenies, and let S be the species tree. If a Super-Reconciliation $R(\mathcal{G}, S)$ exists, then the set of corresponding syntenies trees $\{\tilde{T}_1, \tilde{T}_2, \dots, \tilde{T}_t\}$ is consistent.*

Proof By definition, a Super-Reconciliation $R(\mathcal{G}, S)$ displays \tilde{T}_i , for all $1 \leq i \leq t$, as $R(\mathcal{G}, S)$ is an extension of each tree. Thus, for any triplet $Trp = \{X_1, X_2, X_3\}$ of \mathcal{X} , if \tilde{T}_i and \tilde{T}_j contain the triplet Trp as a sub-leafset, then $R(\mathcal{G}, S)$ displays both $\tilde{T}_i|_{Trp}$ and $\tilde{T}_j|_{Trp}$. In other words, $\tilde{T}_i|_{Trp}$ and $\tilde{T}_j|_{Trp}$ are label-isomorphic. \square

The consistency problem of rooted trees has been widely studied. The BUILD algorithm [26] can be used to test, in polynomial-time, whether a collection of rooted trees is consistent, and if so, construct a compatible, not necessarily fully resolved, supertree, i.e. a tree displaying them all. This algorithm has been generalized to output all compatible minimally resolved supertrees [27, 28, 29], which may be exponential in the number of genes.

The following theorem makes the link between a supertree and a reconciliation.

Theorem 1 *Let $\mathcal{G} = \{T_1, T_2, \dots, T_t\}$ be a set of trees for a set of families organized in an order-consistent set of syntenies \mathcal{X} , and S be the species tree. Let $\tilde{\mathcal{G}} = \{\tilde{T}_1, \tilde{T}_2, \dots, \tilde{T}_t\}$ be the set of synteny trees corresponding to those in \mathcal{G} . If $\tilde{\mathcal{G}}$ is a consistent set of trees then:*

- 1 *A Super-Reconciliation $R(\mathcal{G}, S)$ is an extension of a supertree for $\tilde{\mathcal{G}}$;*
- 2 *Any supertree is the “backbone” of a Super-Reconciliation. Namely, for any supertree \tilde{T} for $\tilde{\mathcal{G}}$, there is a Super-Reconciliation $R(\mathcal{G}, S)$ which is an extension of \tilde{T} .*

The first statement of Theorem 1 follows from Lemma 2. As for the second statement, we will prove it implicitly in Section 5 by providing an algorithm that yields a minimum cost reconciliation on any supertree.

Following Theorem 1, the problem reduces to finding a supertree for the set of synteny trees minimizing the number of segmental duplications and losses. A natural algorithm for the SUPER-RECONCILIATION problem follows:

- 1 Explore the space of all order consistent ancestral syntenies A for \mathcal{X} ;
- 2 Explore the space of all supertrees \tilde{T} for $\tilde{\mathcal{G}}$;
- 3 Find a Super-Reconciliation of minimum cost which is an extension of \tilde{T} with A as an ancestral synteny;
- 4 Select the Super-Reconciliations leading to the minimum cost.

Step 1 and Step 2 have been discussed in this section. Before developing an algorithm for Step 3, which is the purpose of Section 5, we begin by analyzing the theoretical complexity of the SUPER-RECONCILIATION problem.

4 Complexity of the Super-Reconciliation Problem

We have recently considered the problem of finding a supertree of a set of gene trees minimizing the classical single gene duplication and single gene duplication and loss distances. The problem has been shown NP-hard for the duplication distance, and exponential-time algorithms have been developed for both distances. For segmental duplications only, the hardness of SUPER-RECONCILIATION is almost immediate from the results of [30]. For both duplications and losses, the problem remains NP-hard, although the proof is far more technical. Here we give the simpler proof of hardness for minimizing duplications only, and refer the reader to the Supplementary material for the NP-hardness proof for minimizing segmental duplications *and* losses.

Theorem 2 *The SUPER-RECONCILIATION problem is NP-hard for the duplication cost. Furthermore, the minimum number of duplications is hard to approximate within a factor $n^{1-\epsilon}$ for any $0 < \epsilon < 1$, where n is the number of syntenies in the input.*

Proof The hardness follows from that of the MINDUP-SUPERTREE problem, defined as follows. Given a species tree S and a set of gene trees T_1, \dots, T_k , possibly with overlapping leafsets, MINDUP-SUPERTREE asks for a supertree T that displays T_1, \dots, T_k such that the LCA-reconciliation of T and S yields a minimum number d of duplications. It was shown in [30] that it is NP-hard to approximate

d within a factor $n^{1-\epsilon}$ for any $0 < \epsilon < 1$, where here n is the number of genes in $\Gamma = \bigcup_{i=1}^k \mathcal{L}(T_i)$.

To reduce MINDUP-SUPERTREE to the SUPER-RECONCILIATION problem, it essentially suffices to exchange the roles of genes and syntenies. More precisely, given an instance of MINDUP-SUPERTREE consisting of a species tree S and gene trees T_1, \dots, T_k , we compute an instance of SUPER-RECONCILIATION as follows. The species tree is the same as S , and for each gene $g \in \Gamma$, we have a synteney X_g with $s(X_g) = s(g)$. Moreover for each gene tree T_i , we create an identical gene tree T'_i , but in which each gene $g \in \mathcal{L}(T_i)$ is replaced by a unique gene g_{T_i} that belongs to synteney X_g (and hence $s(g) = s(g_{T_i}) = s(X_g)$). Thus the synteney tree \tilde{T}_i for T'_i is obtained by replacing each leaf g of T_i by X_g . In particular, there are n syntenies. The order of the genes on the syntenies is arbitrary (since we are not counting segmental losses).

It only remains to show the correspondence between the solutions for the two problem instances. Suppose that the MINDUP-SUPERTREE instance admits a supertree T with d duplications when reconciled. Let \tilde{T} be the synteney tree obtained from T by replacing each gene $g \in \mathcal{L}(T)$ by X_g . Because $s(g) = s(X_g)$, both T and \tilde{T} have the same duplications under the LCA reconciliation, which is d . Conversely, if our SUPER-RECONCILIATION instance admits a synteney tree \tilde{T} with d duplications, replacing each leaf X_g by g yields a supertree for the MINDUP-SUPERTREE instance with d duplications. Because the value of the solutions are preserved and $n = |\Gamma|$ corresponds to the number of syntenies, this reduction is approximation-preserving and the hardness result follows. \square

We state our second hardness result formally here.

Theorem 3 *The SUPER-RECONCILIATION problem is NP-hard for the Dup, fLoss and pLoss cost.*

5 A Super-Reconciliation for a supertree

In this section, we are given a set $\mathcal{G} = \{T_1, T_2, \dots, T_t\}$ of consistent gene trees for a set of families $\mathcal{F} = \{\Gamma_1, \Gamma_2, \dots, \Gamma_t\}$ organized in an order consistent set of syntenies \mathcal{X} , and a species tree S for the set Σ of taxa containing the genes. In addition, we are given a supertree \tilde{T} for the synteney trees $\tilde{\mathcal{G}} = \{\tilde{T}_1, \tilde{T}_2, \dots, \tilde{T}_t\}$ corresponding to those in \mathcal{G} , and an order consistent ancestral synteney A for \mathcal{X} .

Given a Super-Reconciliation $R(\mathcal{G}, S)$ (R for short), because R is obtained from \tilde{T} by grafting leaves, each node of \tilde{T} is present in R . Hence we say that $v \in V(\tilde{T})$ has a *corresponding node* v' in R . More precisely, if $l \in \mathcal{L}(\tilde{T})$, then $l \in \mathcal{L}(R)$ also and the correspondence is immediate. If v is an internal node of $V(\tilde{T})$, the node v' of R corresponding to v is $lca_R(\{l : l \in \mathcal{L}(\tilde{T}[v])\})$. We show that, as in the traditional reconciliation setting, the nodes of R that are also in \tilde{T} should be mapped to the lowest species possible. To simplify the argument, we will call an internal node a full loss if it is the parent of a *fLoss* event. For later reference, we note that the proofs of Lemma 3 and Lemma 4 do not involve the gene orders in any way.

Lemma 3 *Let $R(\mathcal{G}, S)$ be a Super-Reconciliation of minimum cost which is an extension of \tilde{T} . Let $v \in V(\tilde{T})$ and let v' be the node corresponding to v in $R(\mathcal{G}, S)$. Then $s(v') = lca_S(s(\mathcal{L}(\tilde{T}[v])))$.*

Proof First, observe that the statement is clearly true for the leaves. Assume that the statement is false. Now, let v be a node of \tilde{T} such that its corresponding node v' does not satisfy the statement - moreover, choose v to be a minimal node with this property (meaning that for the children v_l and v_r of v , the corresponding nodes v'_l and v'_r in $R(\mathcal{G}, S)$ satisfy $s(v'_l) = lca_S(s(\mathcal{L}(\tilde{T}[v_l])))$ and $s(v'_r) = lca_S(s(\mathcal{L}(\tilde{T}[v_r])))$). Note that v must exist, since the statement is true for the leaves.

Now, we may assume that $s(v') \neq lca_S(s(v'_l), s(v'_r))$, as otherwise v' satisfies the lemma. Thus in S , there are at least k edges on the path from $s(v')$ to $lca_S(s(v'_l), s(v'_r))$, where here $k > 0$. It is not hard to verify that in this case, v' must be a duplication node, according to the definition of a reconciliation. This implies that there are at least k full losses on the path from v' to v'_l and at least k full losses on the path from v' to v'_r . Consider the Super-Reconciliation R' that is identical to $R(\mathcal{G}, S)$, with the exception that $s(v') = lca_S(s(v'_l), s(v'_r))$. Then the $2k$ losses on the paths between v' and v'_l and between v' and v'_r are not needed anymore, although if v' is not the root, k losses become necessary on the path between v' and w' , where w' is the node corresponding to the parent w of v in \tilde{T} . Remapping v' cannot increase the number of duplications, and so we have saved k losses.

It remains to argue that the number of partial losses remains the same. But this is easy to see. We keep the same synteny assignment at nodes v' , v'_l and v'_r (and w' if v' is not the root) as in $R(\mathcal{G}, S)$. If v' was a segmental duplication in $R(\mathcal{G}, S)$, we set v' to be a segmental duplication in R' as well. The number of partial losses on the paths between v' and v'_l, v'_r (and w') therefore remains the same as in $R(\mathcal{G}, S)$. \square

We now show that speciation and duplication nodes are easy to identify. Essentially, we may set the events of internal nodes as in the classical LCA-mapping reconciliation. In what follows, assume that \tilde{T} is reconciled under the LCA-mapping, and put $s(v) = lca_S(s(\mathcal{L}(\tilde{T}[v])))$ for every $v \in V(\tilde{T})$.

Lemma 4 *Let $R(\mathcal{G}, S)$ be a Super-Reconciliation of minimum cost which is an extension of \tilde{T} . Let $v \in V(\tilde{T})$ be an internal node of \tilde{T} and let v' be its corresponding node in $R(\mathcal{G}, S)$. Moreover let v_l and v_r be the children of v . If $s(v_l)$ and $s(v_r)$ are separated in S , then v' is a speciation, and otherwise v' is a duplication.*

Proof Let v'_l and v'_r be the nodes corresponding to v_l and v_r , respectively, in $R(\mathcal{G}, S)$. First, if $s(v_l)$ and $s(v_r)$ are not separated, then by Lemma 3, $s(v'_l)$ and $s(v'_r)$ are not separated, hence it is not possible for v' to be a speciation. Therefore v' must be a duplication.

Suppose instead that $s(v_l)$ and $s(v_r)$ are separated in S , but that v' is labeled by a duplication event $Dup(X, [i, j])$, where X is the synteny assigned at v' . On the path from v' to v'_l , there may be some $pLoss$ events and some nodes that were grafted owing to full losses. We may assume that all full loss events, if any, have

occurred before the $pLoss$ events on this path (i.e., nodes grafted from full losses are closer to v'). This is without loss of generality, as this does not change the resulting synteny in v'_l . We shall make the same assumption with the path from v' to v'_r . Now, by Lemma 3, $s(v') = lca_S(s(v_l), s(v_r))$. Because v' is a duplication, the two children w_l, w_r of v' in $R(\mathcal{G}, S)$ must satisfy $s(w_l) = s(w_r) = s(v')$. Since $s(v'_l) \neq s(v') \neq s(v'_r)$, we have that $\{w_l, w_r\} \cap \{v'_l, v'_r\} = \emptyset$, and therefore w_l and w_r were grafted on \tilde{T} due to full losses. If we label v' as a speciation $Spe(X, [1, |X|])$, these two full losses are not needed anymore, and by doing so we have one duplication less and two full losses less. Let Y_l and Y_r be the two synteny that were assigned at w_l and w_r in $R(\mathcal{G}, S)$, respectively. Then $Y_l = X$ and $Y_r = X[i, j]$ or vice-versa (assume the former, without loss of generality). Suppose that w_r was an ancestor of v'_r in $R(\mathcal{G}, S)$, again without loss of generality. The substring $X[i, j]$ can be obtained from X by adding at most two partial losses on the path from v' to v'_r . The rest of the reconciliation can remain the same. To sum up, we have removed one duplication and two full losses, and inserted at most two partial losses to reproduce the effect of the segmental duplication. This contradicts that $R(\mathcal{G}, S)$ is a reconciliation of minimum cost. \square

From Lemma 4, it follows that we know the event-type (Dup or Spe) of each internal node of the supertree \tilde{T} . It then remains to extend the tree with losses and infer the actual event at each node (i.e., the corresponding synteny and segment being duplicated or lost). It is easy to see that losses and segments affected by the events are fully determined by gene orders assigned to internal nodes. Therefore, the problem reduces to the classical “small phylogeny problem” generally defined as follows: Given an alphabet Σ (nucleotides or amino-acids or genes), a distance on the set of words of Σ (edit distance for gene sequences or rearrangement distances for gene orders) and a tree T with leaves being words on Σ (extant gene sequences or gene orders), find the labeling of ancestral nodes (ancestral sequences or orders) minimizing the total cost of the tree. This cost is the sum of costs of each branch, which is the distance between the two words connected by the branch.

Here, we are given a synteny tree \tilde{T} for a set \mathcal{X} of synteny on a set of gene families \mathcal{F} , and an ancestral synteny A which is an order of \mathcal{F} . We want to find a *synteny assignment* attributing a partial order on \mathcal{F} to each node of $V(\tilde{T})$. We assume that the root r of \tilde{T} is assigned the synteny A . It follows from the considered evolutionary model that, for two nodes u and v of \tilde{T} with u being an ancestor of v , the synteny X_v assigned to v should be a subsequence of the string X_u assigned to u . A synteny assignment verifying this condition is called a *valid synteny assignment* for \tilde{T} .

For $v \in V(\tilde{T})$, define $d(v, X)$ as the minimum number of segmental duplications and losses induced by a synteny assignment on $\tilde{T}[v]$ with X being the assignment at v . The problem SMALL-PHYLOGENY FOR SYNTENIES is to find an optimal assignment, i.e. an assignment leading to $d(\tilde{T}) = \min_X d(r(\tilde{T}), X)$ for X belonging to the set of synteny that are order consistent with \mathcal{X} .

Solving this problem can be done by dynamic programming by computing $d(v, X)$, for each $v \in V(\tilde{T})$ and each possible synteny X .

Let v be an internal node of \tilde{T} and v_l, v_r be its two children. Let X, X_l, X_r be valid assignments for respectively v, v_l and v_r . Then X_l and X_r are subsequences of

X . If v is a speciation, then all missing genes in X_l and X_r are the result of losses. Otherwise, if v is a duplication, then for at most one of X_l and X_r , the missing prefix or suffix can be due to the partial duplication of a segment of X , and all other missing genes should be the result of losses. This motivates the following two variants of the loss distance between two syntenyies.

Let X and Y be two syntenyies with Y being a subsequence of X . We let $D^T(X, Y)$ denote the minimum number of segmental losses required to transform X to Y and $D^P(X, Y)$ the minimum number of segmental losses required to transform a substring of X to Y .

Theorem 4 *Let v be a node of \tilde{T} , X be a synteny and $\mathcal{S}(X)$ be the set of subsequences of X .*

- *If v is a leaf, then $d(v, X) = 0$ if X is the extant synteny corresponding to leaf v , and $+\infty$ otherwise;*
- *If v is a speciation with children v_l and v_r , then,*

$$d(v, X) = \min_{(X_l \in \mathcal{S}(X))} (D^T(X, X_l) + d(v_l, X_l)) + \min_{(X_r \in \mathcal{S}(X))} (D^T(X, X_r) + d(v_r, X_r));$$

- *If v is a duplication node with children v_l and v_r , then*

$$d(v, X) = 1 + \min \begin{cases} \min_{(X_l \in \mathcal{S}(X))} (D^T(X, X_l) + d(v_l, X_l)) + \\ \min_{(X_r \in \mathcal{S}(X))} (D^P(X, X_r) + d(v_r, X_r)), \\ \min_{(X_l \in \mathcal{S}(X))} (D^P(X, X_l) + d(v_l, X_l)) + \\ \min_{(X_r \in \mathcal{S}(X))} (D^T(X, X_r) + d(v_r, X_r)) \end{cases}$$

The above can be used to solve the SMALL-PHYLOGENY FOR SYNTENIES problem with dynamic programming. To do this, one can simply traverse \tilde{T} in post-order, and apply the recurrences of Theorem 4 at each node encountered. We finish this section by analyzing the complexity of this algorithm. Let $n = |V(\tilde{T})|$ and let t be the number of gene families involved in the SMALL-PHYLOGENY FOR SYNTENIES problem instance. For a node $v \in V(\tilde{T})$ and a synteny X , there are $O(2^t)$ possible subsequences of X . The value of $d(v, X)$ thus depends on the $O(2^t)$ values for its left child v_l and the $O(2^t)$ values for its right child v_r . If these are known, then $d(v, X)$ can be computed in time $O(t2^t)$ (it is straightforward to check that D^T and D^P can be computed in time $O(t)$).

Let us now consider the number of possible entries in our dynamic programming table. The possible syntenyies for X correspond to the subsequences of a topological sorting of an acyclic directed graph with t nodes (see supplementary material). In the worst case, there are $O(2^t \cdot t!) = O(2^{t \log t + t})$ such syntenyies. It follows that there are at most $O(n2^{t \log t + t})$ entries in the dynamic programming table, and each entry takes time $O(t2^t)$. It is known that if there are k possible topological sortings in a directed acyclic graph, then they can be enumerated in time $O(k)$ [31] (it is

worth noting however that counting the number of such topological sortings in #P-complete [32]). Therefore, if t is not too large, then the above recurrences can solve the small phylogeny problem relatively quickly, even if n is large. Put differently, the SMALL-PHYLOGENY FOR SYNTENIES problem is *fixed-parameter tractable* with respect to parameter t .

Corollary 1 *The SMALL-PHYLOGENY FOR SYNTENIES problem can be solved in time $O(t2^{t \log t + 2t}n)$, where t is the number of gene families present in the input and n is the number of syntenies.*

6 Unordered Super-Reconciliation

The strongest and less biologically supported condition for the existency of a DL Super-Reconciliation is probably gene order consistency. In fact, genomes being subject to rearrangements shuffling gene organization, it is hard to expect that a set of homologous chromosomal segments in phylogenetically distant genomes would exhibit the same gene order. In other words, we can hardly ignore the presence of rearrangements in the evolutionary history leading to a set of homologous genomic regions.

The small phylogeny problem, which consists in inferring ancestral gene orders minimizing a given rearrangement distance, has been extensively studied (see for example [33, 34, 35, 36, 37]). Algorithmic developments and results differ depending on the considered rearrangement distance. The most studied one is probably the DCJ distance, accounting for artificial movements implicitly mimicing inversions and transpositions [38, 39].

Almost all versions of the small phylogeny problem with rearrangements have been proven NP-hard, even those accounting for equal gene content for all genomes [40]. Heuristics have also been developed for inferring ancestral gene orders minimizing rearrangements, duplications and loss events (reviews can be found in [41, 42]). Extension of these heuristics to the Super-Reconciliation problem is certainly possible, but can only increase the intractability of the original problem.

Here, we explore a compromise which consists in considering an evolutionary model accounting for segmental duplications, losses and rearrangements, but yet only minimizing duplication and loss events. In other words, gene orders are not important anymore, as we can use as many rearrangements as we want for obtaining the required orders.

Reducing syntenies to their range sets, an *unordered evolutionary history* of a set of syntenies can be represented as a partially binary tree where each internal node v corresponds to an event $\mathcal{E}(\text{Set}(X))$ with $X = \text{synteny}(v)$ being the synteny at v and $\mathcal{E} \in \{\text{Spe}, \text{Dup}, \text{pLoss}\}$ such that, if \mathcal{E} is:

- 1 *Spe*, then v is a binary node with two children corresponding to syntenies Y and Z such that $\text{Set}(X) = \text{Set}(Y) = \text{Set}(Z)$ and $s(Y)$ and $s(Z)$ are the two children of $s(X)$ in S .
- 2 *Dup*, then v is a binary node with two children corresponding to syntenies Y and Z such that $\text{Set}(Y) = \text{Set}(X)$, $\text{Set}(Z) \subseteq \text{Set}(X)$ and $s(X) = s(Y) = s(Z)$.
- 3 *pLoss*, then v is a unary node with a child corresponding to a synteny Y such that $\text{Set}(Y) \subsetneq \text{Set}(X)$ and $s(X) = s(Y)$.

If no ambiguity on the synteny of v , we will denote by $\text{Set}(v)$ the range set of the synteny at node v of tree T (in other words, $\text{Set}(v) = \text{Set}(\text{synteny}(v))$).

An Unordered Super-Reconciliation (USR) $R_u(\mathcal{G}, S)$ of a set $\mathcal{G} = \{T_1, T_2, \dots, T_t\}$ of gene trees with a species tree S is a labeled synteny tree which is an extension of the trees \tilde{T}_i , for $1 \leq i \leq t$, representing a valid unordered evolutionary history for \mathcal{X} . The cost $d(R_u(\mathcal{G}, S))$ of such an unordered Super-Reconciliation is the number of induced *Dup*, *fLoss* and *pLoss*.

The UNORDERED SUPER-RECONCILIATION problem then consists in inferring the USR of minimum cost. Notice that, as gene order is ignored, at most one *pLoss* can separate two binary nodes on a most parsimonious USR.

Regarding existence conditions, Lemma 2 and Theorem 1 clearly apply to the USR problem, as gene order information is not involved in tree consistency. Namely, an USR exists if and only if the trees of $\tilde{\mathcal{G}}$ are consistent, and in this case any supertree for $\tilde{\mathcal{G}}$ is the backbone of an USR. Lemma 3 and Lemma 4 also apply since, as mentioned before, their proofs do not involve gene orders. Therefore, we may assume that, if a supertree \tilde{T} of synteny is given, its nodes can be mapped according to the LCA-mapping and its speciation/duplication nodes identified in this way.

The USR problem thus reduces to a small phylogeny problem which consists in inferring internal node gene contents of the supertree \tilde{T} leading to a minimal duplication and loss cost. As duplications are already determined by the node labeling of \tilde{T} , only loss events remain to be minimized. Notice that the root's gene content is just \mathcal{F} .

We add to \tilde{T} the *fLoss* branches obtained from the LCA-Reconciliation of \tilde{T} with S . In other words, the new tree is an intermediate between \tilde{T} and $R_u(\mathcal{G}, S)$. For practical reasons, we still call it \tilde{T} . Notice that *fLoss* branches can only create speciation nodes. For the requirements of the following algorithms, the empty synteny is assigned to the leaf created by an *fLoss* branch.

We now present a dynamic programming algorithm to find the minimum number of *pLoss* events required for a USR.

A dynamic programming approach for optimal USRs

Given an USR R , we write $\text{Set}_R(v)$ the range set of v in R . For an internal node v of \tilde{T} , denote

$$lca_{\text{Set}}(v) = \bigcup_{l \in \mathcal{L}(\tilde{T}[v])} \text{Set}(l)$$

as the set of all gene families that appear in a synteny under v . Note that v must have a gene in every family in $lca_{\text{Set}}(v)$. For a range set X , denote by $C_X(v)$ the minimum cost of an USR R between $\tilde{T}[v]$ and S in which we assign $\text{Set}_R(v) = X$ (if $lca_{\text{Set}}(v)$ is not a subset of X , put $C_X(v) = \infty$). We denote $C_{lca}(v) := C_{lca_{\text{Set}}(v)}(v)$, i.e. the cost when we assign the smallest possible range set to v . If v is a leaf, we have $C_{lca}(v) = 0$ and $C_X(v) = \infty$ for any $X \neq \text{Set}(v)$. The value we are interested in is $C_{lca}(r)$, where r is the root of \tilde{T} .

We first show that the exact nature of the “extra” content that might be assigned to an internal node v is irrelevant for the computation of the optimal cost.

Lemma 5 *Let $v \in V(\tilde{T})$ be an internal node of \tilde{T} , and let X, Y be any range sets satisfying $lca_{Set}(v) \subsetneq X, Y$. Then $C_X(v) = C_Y(v)$.*

Proof This can be shown by induction on the depth of the nodes of \tilde{T} . The Lemma is true for leaves, as $C_X(v) = C_Y(v) = \infty$ in this case. So assume that v is an internal node. Let R be a minimum USR of $\tilde{T}[v]$ in which $Set_R(v) = X$. We build a reconciliation R' in which $Set_{R'}(v) = Y$. Consider the children v_1 and v_2 of v in \tilde{T} . If $Set_R(v_1) = lca_{Set}(v_1)$, then either there is a *pLoss* node on the vv_1 path in R , or v is a duplication and X was partially duplicated. In any case for our USR R' , we can use this *pLoss* or duplication to lose $Y \setminus lca_{Set}(v_1)$, and assign $Set_{R'}(v_1) = lca_{Set}(v_1)$ without incurring additional cost compared to R . We then use the same reconciliation as R for the subtree $\tilde{T}[v_1]$, and thus R' incurs no extra cost on the v_1 side. If instead $Set_R(v_1)$ strictly contains $lca_{Set}(v_1)$, then in R' we put $Set_{R'}(v_1) = Y$ without requiring any *pLoss* on the vv_1 path. Moreover by induction, $C_{Set_R(v_1)}(v_1) = C_Y(v_1)$. It follows that R' has as many losses as R in the $\tilde{T}[v_1]$ subtree — and since there is no loss on the vv_1 branch, R' has at most as many losses as R on the v_1 side (note that R' could have strictly less losses than R on the v_1 side if R had a loss on the vv_1 branch — this implicitly means that R had no such loss, as otherwise our R' will end up having less losses than R and contradict its optimality). Now, it suffices to observe that the same scheme can be applied to v_2 as well (noting that losses saved by partial duplications cannot happen on both sides of v), showing that there is an R' that is of cost no more than R . \square

Lemma 5 implies that there are two possible minimum reconciliation costs for $\tilde{T}[v]$. Either v is assigned $lca_{Set}(v)$ and its cost is $C_{lca}(v)$, or it is assigned X with extra content and its cost is $C_X(v)$, for any X strictly containing $lca_{Set}(v)$. We will therefore denote by $C^*(v)$ the minimum reconciliation cost of $\tilde{T}[v]$ when v is assigned any X such that $X \neq lca_{Set}(v)$. This leads to a dynamic programming formulation that takes into account the two cases. For an internal node v , define $spec(v) = 1$ if v is a speciation, and $spec(v) = 0$ otherwise. The value of $C^*(v)$ can be computed according to the following Lemma.

Lemma 6 *For an internal node v of \tilde{T} with children v_1 and v_2 , we have*

$$C^*(v) = \min \begin{cases} C_{lca}(v_1) + C_{lca}(v_2) + 1 + spec(v) \\ C^*(v_1) + C_{lca}(v_2) + spec(v) \\ C_{lca}(v_1) + C^*(v_2) + spec(v) \\ C^*(v_1) + C^*(v_2) \end{cases}$$

Proof Let R be an optimal USR in which $Set_R(v) = X \neq lca_{Set}(v)$. It follows that $X \neq lca_{Set}(v_1)$ and $X \neq lca_{Set}(v_2)$. The lemma lists all the possible ways of sending extra content to the children or not. In the first case, we have two losses: we lose $X \setminus lca_{Set}(v_1)$ and $X \setminus lca_{Set}(v_2)$ on both the vv_1 and vv_2 branches, respectively (if v is a duplication, we can save one loss in a partial duplication, hence the $spec(v)$ term). In the second and third cases, we lose content only on one side

(again potentially using a partial duplication), and in the last case, we transfer X to the children without incurring any loss. \square

The computation of $C_{lca}(v)$ has more cases, depending if the children of v have the same lca_{Set} value or not. We show in Algorithm 1 how $C_{lca}(v)$ can be computed. We omit the proof of correctness for this algorithm as it is similar to that of Lemma 6.

Algorithm 1 Computation of $C_{lca}(v)$

```

1: function COMPUTE-CLCA( $v$ )
2:   Let  $v_1$  and  $v_2$  be the children of  $v$ 
3:   if  $lca_{Set}(v) = lca_{Set}(v_1) = lca_{Set}(v_2)$  then
4:      $C_{lca}(v) = C_{lca}(v_1) + C_{lca}(v_2)$  ▷ (no loss required)
5:   else if  $lca_{Set}(v) = lca_{Set}(v_1) \neq lca_{Set}(v_2)$  then
6:      $C_{lca}(v) = \min \begin{cases} C_{lca}(v_1) + C_{lca}(v_2) + spec(v), \\ C_{lca}(v_1) + C^*(v_2) \end{cases}$  ▷ (loss on  $vv_2$  or not)
7:   else if  $lca_{Set}(v) = lca_{Set}(v_2) \neq lca_{Set}(v_1)$  then
8:      $C_{lca}(v) = \min \begin{cases} C_{lca}(v_1) + C_{lca}(v_2) + spec(v), \\ C_{lca}(v_2) + C^*(v_1) \end{cases}$  ▷ (loss on  $vv_1$  or not)
9:   else if  $lca_{Set}(v_1) \neq lca_{Set}(v) \neq lca_{Set}(v_2)$  then
10:     $C_{lca}(v) = \min \begin{cases} C_{lca}(v_1) + C_{lca}(v_2) + 1 + spec(v), \\ C^*(v_1) + C_{lca}(v_2) + spec(v), \\ C_{lca}(v_1) + C^*(v_2) + spec(v), \\ C^*(v_1) + C^*(v_2) \end{cases}$  ▷ (loss on  $vv_1/vv_2$ )
11:   end if
12: end function

```

It is now clear that if the values of $C_{lca}(v_i)$ and $C^*(v_i)$ are known for the children v_1, v_2 of v , then $C_{lca}(v)$ and $C^*(v)$ can be computed in constant time, assuming we have access to $lca_{Set}(v)$ for every $v \in \tilde{T}$. By computing these values in a post-order traversal of \tilde{T} , we can compute $C_{lca}(r)$ for the root of \tilde{T} in time $O(|V(\tilde{T})|)$. It is also straightforward to conceive a backtracking procedure to construct an actual USR. Moreover, every optimal solution can be produced by our dynamic programming paradigm.

This algorithm requires computing $lca_{Set}(v)$ for every vertex, which can be accomplished in time $O(|V(\tilde{T})||\mathcal{F}|)$ (recall that \mathcal{F} is the set of gene families). This actually dominates the running time.

Theorem 5 *The minimum cost of a USR can be obtained in time $O(|V(\tilde{T})||\mathcal{F}|)$.*

7 Application

7.1 Simulated datasets

The dynamic programming algorithm has been implemented in C++^[1] and tested on balanced trees obtained from simulated evolutionary histories. Simulations have been performed according to five parameters: t , the number of gene families in the ancestral synteny; d , the maximum depth of the balanced tree; p_{dupl} , the probability for any given node to be a segmental duplication; p_{loss} , the probability for a loss to occur under any given node; and p_{length} , the probability to remove one gene in a segmental loss, defining the probability for a loss to remove k genes (for $k \in$

^[1]The program and simulations are available at:

<https://github.com/UdeM-LBIT/SuperReconciliation>

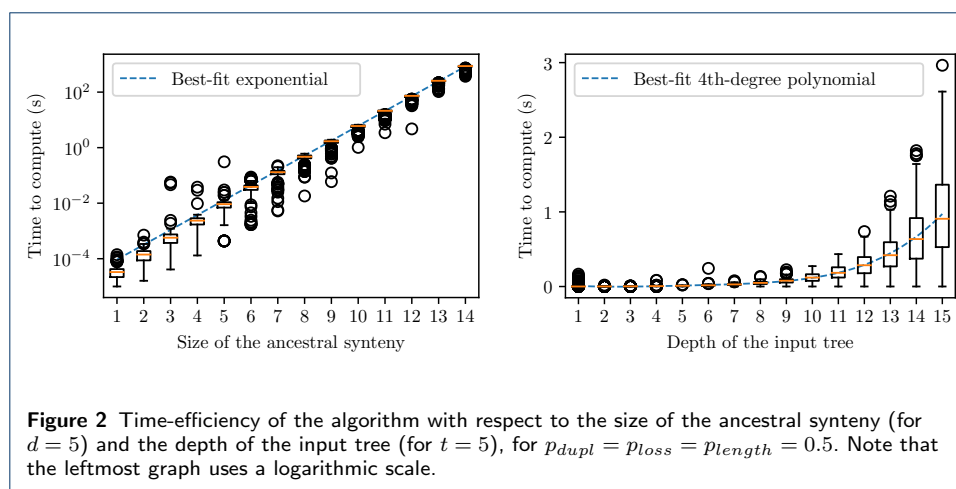
$\{1, 2, 3, \dots\}$): $P(X = k) = (1 - p_{length})^{k-1} p_{length}$, following a shifted geometric distribution.

Simulations yield Super-Reconciliations leading to fully labelled trees. The input of the Super-Reconciliation algorithm is then obtained from those trees by removing loss nodes and synteny information on the internal, non-root nodes.

From an accuracy point of view (results not shown), as expected the larger the density of duplication and loss events, the further is the simulated history from a most parsimonious history, and thus from the inferred tree.

As for time-efficiency, values for inferring the Super-Reconciliation of a single tree, aggregated over 500 simulations per value of t , the size of the ancestral synteny (number of gene families), are given in Figure 2. Computations have been done on the “Cedar” cluster of Compute Canada with 32 *Intel 8160* CPUs operating at 2.10 GHz. As expected, running time exponentially increases with respect to parameter t . This prevented us from extending the simulations beyond an ancestral synteny of size 14, for which the Super-Reconciliation of a single tree of depth 5 required around 15 min. However, if the synteny size remains fixed, running times increase polynomially with the size of the trees. As shown by the right diagram of Figure 2, for an ancestral synteny of size 5, simulations exhibit a running time of no more than few seconds for trees with depth up to 15, representing balanced trees with up to 2^{15} leaves.

Apart from genomic segments related through a recent whole genome duplication event, real biological datasets are more likely to reveal large gene families rather than large sets of gene families evolving in concert. Thus, the increase in running time according to the size of the ancestral synteny is unlikely to be a bottleneck towards applying our Super-Reconciliation algorithm. The particular case of whole genome duplication is however worth exploring in more details.



7.2 The opioid system

The opioid receptors, important regulators of neurotransmission and reward mechanisms in mammals, offer an interesting proof of concept, as these genes are present in clusters with conserved synteny in vertebrate genomes. Three genes for the opioid receptors (OPR) were identified and named OPRD1 (delta), OPRK1 (kappa) and

OPRM1 (mu). A fourth gene was later found (OPRL1) in rodents and human. In human, they are located on the chromosomes 1, 6, 8 and 20.

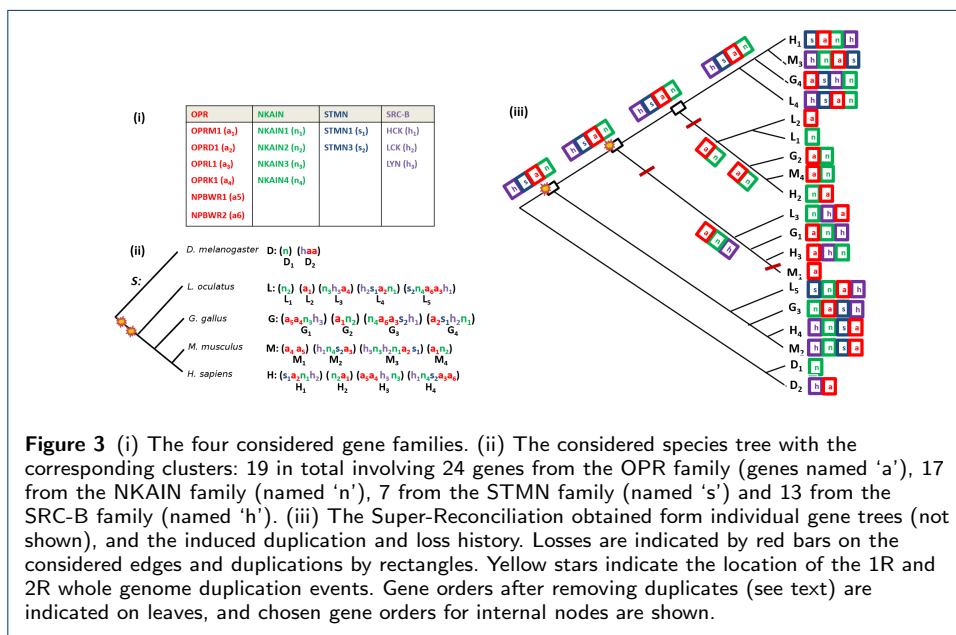


Figure 3 (i) The four considered gene families. (ii) The considered species tree with the corresponding clusters: 19 in total involving 24 genes from the OPR family (genes named 'a'), 17 from the NKAIN family (named 'n'), 7 from the STMN family (named 's') and 13 from the SRC-B family (named 'h'). (iii) The Super-Reconciliation obtained from individual gene trees (not shown), and the induced duplication and loss history. Losses are indicated by red bars on the considered edges and duplications by rectangles. Yellow stars indicate the location of the 1R and 2R whole genome duplication events. Gene orders after removing duplicates (see text) are indicated on leaves, and chosen gene orders for internal nodes are shown.

Previous studies have considered the duplication scenario explaining the evolution of the opioid receptor genes [18, 19, 20]. The main question was whether observed paralogs arose from the two whole genome duplication events, often called 1R and 2R, known to have occurred early in vertebrate evolution. By exploring regions surrounding the OPR genes in human, four syntenic regions, containing genes from three other families (NKAIN, SRC-B and STMN) apparently sharing a common history, were identified. From the analysis of individual gene trees (Neighbor-joining and quartet-puzzling maximum likelihood trees), conclusions associating the evolution of the opioid system related genes to the 1R and 2R events were drawn.

Here, we consider the same four gene families OPR, NKAIN, STMN, and SRC-B, and further extend the OPR family with two neuropeptide NPBWR receptors, known to be closely related to the opioid receptors (Figure 3.(i)). Protein sequences and gene orders were downloaded from the Ensembl database (Release 92) [2] for the following five species: *Homo sapiens*, *Mus musculus*, *Gallus gallus*, *Lepisosteus oculatus* (spotted gar) and *Drosophila melanogaster*. Gene orders are given in Figure 3.(ii).

For each gene family, we built a multiple sequence alignment with ClustalW [43] (Gonnet weight matrix and gap opening and extension penalties respectively set to 10 and 0.2). Maximum likelihood gene trees were subsequently constructed for each family using MEGA7 [44] (Jones-Taylor-Thornton substitution matrix and uniform rates among sites). As some syntenies contained paralogs (multiple copies from the same gene family, for example synteny H_3 contains two 'a'), duplicates were removed in a way maximizing gene tree consistency. Although gene trees were still inconsistent, the overall clustering of gene copies was preserved among gene trees,

[2] <https://useast.ensembl.org/index.html>

and consistency could be attained after some local adjustments, using the species tree as reference.

The obtained Super-Reconciliation is given in Figure 3.(iii). Notice however that gene orders are far from being consistent. In fact, all considered genomes are separated by a considerable evolutionary distance, and therefore, local rearrangements could have occurred along each lineage-specific branch. Choosing the (h, s, a, n) order on every node of the tree and assuming rearrangements to occur at terminal edges, i.e. after duplication and loss events, leads to a history of three duplications and two losses before the speciation of bony fish and tetrapods, with two duplications correlating with the 1R and 2R tetraploidization events. This result is in agreement with previous studies on the opioid receptor genes [18].

Further analysis, using more genes and species, is required to provide a more detailed scenario for the evolution of the opioid receptor genes. Our objective here however, was not to verify a given hypothesis, but rather to provide a proof of concept and explore the applicability and limitations of the proposed reconciliation model on real data.

8 Discussion

We have presented a natural extension of the DL Reconciliation model to handel segmental duplications and losses. This is the first effort towards developing a unifying automated method framework for reconciling a set of gene trees. We provide computational complexity results and a general inference method.

However, as the considered evolutionary model is restricted to losses and transposed duplications, the possibility of application to real datasets remains very limited. In particular, a duplication and loss history does not always exist for a set of synteny if rearrangements are ignored, as the corresponding gene orders may be inconsistent. One solution would be to minimally correct gene orders to ensure consistency, before applying the DL Super-Reconciliation model. In this paper, we have considered an alternative way of working around this problem, which consists in extending the evolutionary model to account for rearrangements, but still only minimize duplication and loss events. The underlying Unordered Super-Reconciliation problem has been shown tractable. However, this way of integrating rearrangements is far from being fully satisfactory as the obtained evolutionary history may lead to a prohibitive number of rearrangements. In other words, the problem of Super-Reconciliation with rearrangements remains open.

Another strong constraint is the fact that tandem duplications, leading to synteny with multiple gene copies, are ignored. In fact, only transposed duplications, i.e. duplications creating new synteny, are allowed. Although the Super-Reconciliation model can easily be extended to tandem duplications by allowing for unary duplication nodes, the inference methodology developed in this paper is hardly applicable in this context. In particular, gene order consistency is a more challenging problem in presence of interleaving tandem duplications. In addition, having many gene paralogs in synteny lead to multi-labeled trees (or mul-trees), i.e. trees with many leaves with the same label. This raises the issue of finding an appropriate definition of mul-tree consistency and of a “mul-supertree”. But more importantly, what would be the meaning of a synteny mul-supertree, with the same synteny labeling more

than one leaf? Clearly such a supertree cannot represent the backbone of a “valid” evolutionary history represented by a Super-Reconciliation.

A way of getting around this problem would be to prune multi-labeled gene trees in a way leading to “optimal” single-labeled trees. The more natural way to state the decision problem is whether there exist a way of choosing a single gene copy from each family represented in a synteny in a way the obtained single-labeled gene trees are consistent. This is the way we implicitly handled the gene families of the opioid system. Alternatively, we can consider the optimization problem of finding the pruning minimizing a Robinson-Foulds distance between trees. Although authors have considered similar problems for mul-trees [45, 46, 47, 48], as far as we know, none have yet handled these particular ones, representing an interesting avenue for future developments.

Competing interests

The authors declare that they have no competing interests.

Author's contributions

Text for this section ...

Acknowledgements

Authors acknowledge the support of the Fonds de Recherche du Québec Nature et Technologie (FRQNT) and of the Natural Sciences and Engineering Research Council (NSERC) (Discovery Grant RGPIN-249834).

Author details

¹Département d'informatique (DIRO), Université de Montréal, Québec, Canada. ²School of Computing Sciences, University of East Anglia, Norwich, UK. ³Department of Computer Science, Université de Sherbrooke, Sherbrooke, Canada.

References

1. Goodman, M., Czelusniak, J., Moore, G.W., Romero-Herrera, A.E., Matsuda, G.: Fitting the gene lineage into its species lineage, a parsimony strategy illustrated by cladograms constructed from globin sequences. *Systematic Zoology* **28**, 132–163 (1979)
2. Zhang, L.X.: On Mirkin-Muchnik-Smith conjecture for comparing molecular phylogenies. *Journal of Computational Biology* **4**, 177–188 (1997)
3. Zmasek, C.M., Eddy, S.R.: A simple algorithm to infer gene duplication and speciation events on a gene tree. *Bioinformatics* **17**, 821–828 (2001)
4. Bansal, M.S., Alm, E.J., Kellis, M.: Efficient algorithms for the reconciliation problem with gene duplication, horizontal transfer and loss. *Bioinformatics* **28**(12), 283–291 (2012). doi:10.1093/bioinformatics/bts225
5. Doyon, J.P., Ranwez, V., Daubin, V., Berry, V.: Models, algorithms and programs for phylogeny reconciliation. *Briefings in bioinformatics* **12**(5), 392–400 (2011)
6. Tofigh, A., Hallett, M., Lagergren, J.: Simultaneous identification of duplications and lateral gene transfers. *IEEE/ACM Trans. Comput. Biol. Bioinform.* **8**(2), 517–535 (2011). doi:10.1109/TCBB.2010.14
7. Vernot, B., Stolzer, M., Goldman, A., Durand, D.: Reconciliation with non-binary species trees. *J Comput Biol* **15**, 981–1006 (2009)
8. Wu, Y.C., Rasmussen, M.D., Bansal, M.S., Kellis, M.: Most parsimonious reconciliation in the presence of gene duplication, loss, and deep coalescence using labeled coalescent trees. *Genome Research* **24**, 475–486 (2014)
9. Akerborg, O., Sennblad, B., Arvestad, L., Lagergren, J.: Simultaneous bayesian gene tree reconstruction and reconciliation analysis. *Proceedings of the National Academy of Sciences USA* **106**(14), 5714–5719 (2009)
10. Sjöstrand, J., Tofigh, A., Daubin, V., Arvestad, L., Sennblad, B., Lagergren, J.: A bayesian method for analyzing lateral gene transfer. *Sys.Biol.* **63**(3), 409–420 (2014)
11. Szöllösi, G.J., E., Tannier, Daubin, V., Boussau, B.: The inference of gene trees with species trees. *Syst.Biol.* **64**(1), 42–62 (2014)
12. Larsson, T.A., Olsson, F., Sundstrom, G., Lundin, L.G., Brenner, S., Venkatesh, B., Larhammar, D.: Early vertebrate chromosome duplications and the evolution of the neuro peptide y receptor gene regions. *BMC Evolutionary Biology* **8**(184) (2008)
13. Abbasi, A.A., Grzeschik, K.H.: An insight into the phylogenetic history of hox linked gene families in vertebrates. *BMC Evolutionary Biology* **7**(239) (2007)
14. Ferrier, D.E.K.: Evolution of homeobox gene clusters in animals: the giga-cluster and primary vs. secondary clustering. *Frontiers in Ecology and Evolution* **4**(34) (2016)
15. Garcia-Fernández, J.: The genesis and evolution of homeobox gene clusters. *Nature Reviews Genetics* **6**, 881–892 (2005)
16. Ajmal, W., Khan, H., Abbasi, A.A.: Phylogenetic investigation of human fgfr-bearing paralogs favors piecemeal duplication theory of vertebrate genome evolution. *Molecular Phylogenetics and Evolution* **81** (2014)
17. Hafeez, M., Shabbir, M., Altaf, F., Abbasi, A.A.: Phylogenomic analysis reveals ancient segmental duplications in the human genome. *Molecular Phylogenetics and Evolution* **94** (2016)

18. Dreborg, S., Sundstrom, G., Larsson, T.A., Larhammar, D.: Evolution of vertebrate opioid receptors. *Proceedings of the National Academy of Sciences USA* **105**(40), 15487–15492 (2008)
19. Stevens, C.W.: The evolution of vertebrate opioid receptors. *Frontiers in bioscience: a journal and virtual library* **14**, 1247–1269 (2009)
20. Sundstrom, G., Dreborg, S., Larhammar, D.: Concomitant duplications of opioid peptide and receptor genes before the origin of jawed vertebrates. *PLoS ONE* **5**(5) (2010)
21. Bérard, S., Gallien, C., Boussau, B., Szollosi, G.J., Daubin, V., Tannier, E.: Evolution of gene neighborhoods within reconciled phylogenies. *Bioinformatics* **28**(18), 382–388 (2012)
22. W, W.D., Anselmetti, Y., Patterson, M., Ponty, Y., Berard, S., Chauve, C., Scornavacca, C., Daubin, V., E, E.T.: DeCoSTAR: Reconstructing the ancestral organization of genes or genomes using reconciled phylogenies. *Genome Biol. Evol.* **9**(5), 1312–1319 (2017)
23. Duchemin, W.: Phylogeny of dependencies and dependencies of phylogenies in genes and genomes. Theses, Université de Lyon (December 2017). <https://tel.archives-ouvertes.fr/tel-01779517>
24. Paszek, J., Gorecki, P.: Efficient algorithms for genomic duplication models. *IEEE/ACM Trans Comput Biol Bioinform.* (2017)
25. Dondi, R., Lafond, M., Scornavacca, C.: Reconciling multiple genes trees via segmental duplications and losses. *WABI*. to appear, ??? (2018)
26. Aho, A.V., Yehoshua, S., Szymanski, T.G., Ullman, J.D.: Inferring a tree from lowest common ancestors with an application to the optimization of relational expressions. *SIAM J. Comput.* **10**(3), 405–421 (1981)
27. Constantinescu, M., Sankoff, D.: An efficient algorithm for supertrees. *J. Classif.* **12**, 101–112 (1995)
28. Ng, M.P., Wormald, N.C.: Reconstruction of rooted trees from subtrees. *Discrete Appl. Math* **69**, 19–31 (1996)
29. Semple, C.: Reconstructing minimal rooted trees. *Discrete Appl. Math.* **127**(3) (2003)
30. Lafond, M., Ouangraoua, A., El-Mabrouk, N.: Reconstructing a supergenetree minimizing reconciliation. *BMC-Genomics* **16**, 4 (2015). Special issue of RECOMB-CG 2015
31. Pruesse, G., Ruskey, F.: Generating linear extensions fast. *SIAM Journal on Computing* **23**(2), 373–386 (1994)
32. Brightwell, G., Winkler, P.: Counting linear extensions. *Order* **8**(3), 225–242 (1991)
33. Moret, B., Wang, L., Warnow, T., Wyman, S.: New approaches for reconstructing phylogenies from gene order data. *Bioinformatics* **17**, 165–173 (2001)
34. Bourque, G., Pevzner, P.A.: Genome-scale evolution: Reconstructing gene orders in the ancestral species. *Genome Research* **12**, 26–36 (2002)
35. Chauve, C., Tannier, E.: A methodological framework for the reconstruction of contiguous regions of ancestral genomes and its application to mammalian genomes. *PLoS Computational Biology* **4**, 1000234 (2008)
36. Ma, J., Zhang, L., Suh, B.B., Raney, B.J., Burhans, R.C., Kent, W.J., Blanchette, M., Haussler, D., Miller, W.: Reconstructing contiguous regions of an ancestral genome. *Genome Research* **16**, 1557–1565 (2007)
37. Sankoff, D., Blanchette, M.: The median problem for breakpoints in comparative genomics. In: Jiang, T., Lee, D.T. (eds.) *Computing and Combinatorics, Proceedings of COCOON '97. Lecture Notes in Computer Science*, pp. 251–263. Springer, Berlin (1997)
38. Bergeron, A., Mixtacki, J., Stoye, J.: Hp distance via double cut and join distance. In: Ferragina, P., Landau, G. (eds.) *Combinatorial Pattern Matching. CPM '08. Lecture Notes in Computer Science*, vol. 5029 (2008)
39. Bergeron, A., Mixtacki, J., Stoye, J.: A new linear time algorithm to compute the genomic distance via the double cut and join distance. *Theoretical Computer Science* **410**, 5300–5316 (2009)
40. Tannier, E., E., C.Z., Sankoff, D.: Multichromosomal median and halving problems under different genomic distances. *BMC Bioinformatics* **10**, 120 (2009)
41. El-Mabrouk, N.: Genome rearrangement with gene families. *Mathematics of Evolution and Phylogeny*, pp. 291–320. Oxford University Press, ??? (2005)
42. Fertin, G., Labarre, A., Rusu, I., Tannier, E., Vialette, S.: *Combinatorics of Genome Rearrangements*. The MIT Press, Cambridge, Massachusetts, London, England (2009)
43. J.D.Thompson, Higgins, D.G., Gibson, T.J.: Clustal w: improving the sensitivity of progressive multiple sequence alignment through sequence weighting, position-specific gap penalties and weight matrix choice. *Nucleic Acids Research* **22**(22), 4673–4680 (1994)
44. Kumar, S., Stecher, G., Tamura, K.: Molecular evolutionary genetics analysis version 7.0 for bigger datasets. *Molecular Biology and Evolution* **33**(7), 1870–1874 (2016)
45. Scornavacca, C., Berry, V., Ranwez, V.: From gene trees to species trees through a supertree approach. In: *LATA* (2009)
46. Chaudhary, R., Burleigh, J.G., Fernández-Baca, D.: Inferring species trees from incongruent multi-copy gene trees using the robinson-foulds distance. *Algorithms for Molecular Biology* **8**(1), 28 (2013). doi:10.1186/1748-7188-8-28
47. Bansal, M.S., Burleigh, J.G., Eulenstein, O., Fernández-Baca, D.: Robinson-foulds supertrees. *Algorithms for Molecular Biology* **5**(1), 18 (2010). doi:10.1186/1748-7188-5-18
48. Deepak, A., Fernández-Baca, D., McMahon, M.M.: Extracting conflict-free information from multi-labeled trees. *Algorithms for Molecular Biology* **8**(1), 18 (2013). doi:10.1186/1748-7188-8-18
49. Holyer, I.: The NP-completeness of edge-coloring. *SIAM Journal on Computing* **10**(4), 718–720 (1981)

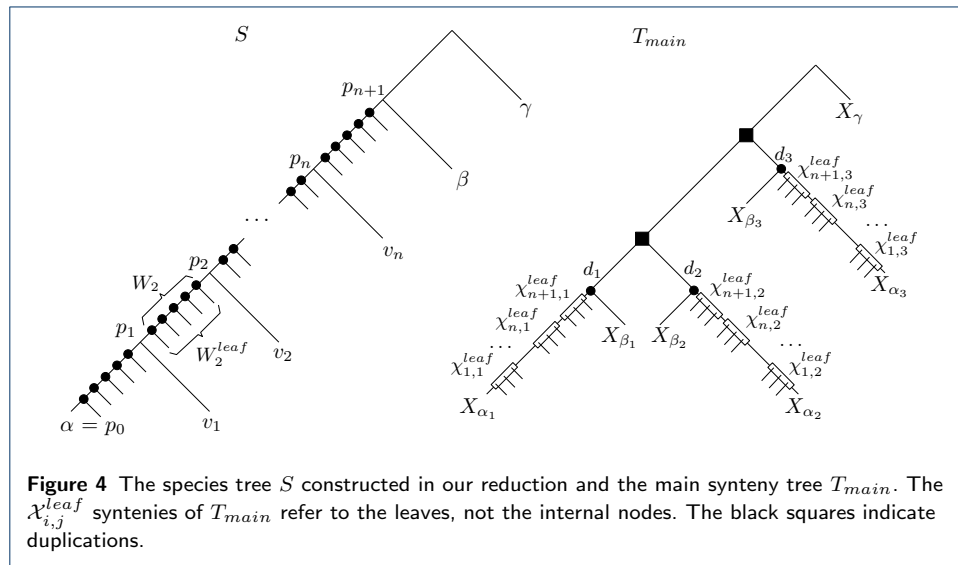
Supplementary material

NP-hardness of the Super-Reconciliation problem

We show here that finding a super-reconciliation that minimizes the number of Dup , $fLoss$ and $pLoss$ events is NP-hard. We reduce the problem of CUBIC 3-EDGE-COLORING to SUPER-RECONCILIATION. Given a graph $G = (V, E)$ in which

each vertex has exactly 3 neighbors, the cubic 3-edge-coloring problem asks whether there exists a *proper coloring* of E with 3 colors, i.e. a partition of E into 3 sets $\{E_1, E_2, E_3\}$ such that for any vertex $v \in V$, the three edges incident to v all belong to a different E_i set. Note that if such a coloring exists, $|E_1| = |E_2| = |E_3| = |E|/3$. This problem was shown to be NP-hard in [49].

In what follows, for an integer k we will denote $[k] = \{1, 2, \dots, k\}$. The father of a node v in a tree will be denoted $p(v)$ (p for ‘parent’). Let $G = (V, E)$ be an instance of CUBIC 3-EDGE-COLORING, and denote $V = \{v_1, \dots, v_n\}$. The ordering of the v_i vertices is not important but must remain fixed for the duration of the proof. To describe our corresponding SUPER-RECONCILIATION instance, we first define the species tree S , which is illustrated in Figure 4. Let S' be a caterpillar on leafset $V \cup \{\alpha, \beta, \gamma\}$ (here a caterpillar is a binary rooted tree in which each internal node has at least one child that is a leaf), where the leaves appear in the order $(\alpha, v_1, v_2, \dots, v_n, \beta, \gamma)$ when traversing from the deepest to the closest leaf to the root. The species α, β and γ are special species, and the v_1, \dots, v_n species are those corresponding to V . For each $i \in [n]$, denote $p_i := p(v_i)$, and $p_0 := \alpha, p_{n+1} := p(\beta)$. To obtain S , for every $i \in [n+1]$, graft a large number of new leaves, say n^{10} , on the branch $p_{i-1}p_i$. Thus there are now n^{10} new internal nodes on the path between p_{i-1} to p_i , and we denote this set of n^{10} internal nodes as W_i , and the set of n^{10} newly inserted leaves as W_i^{leaf} (Figure 4 only shows 5 of the W_i and W_i^{leaf} nodes, with W_2 and W_2^{leaf} shown explicitly). Note that S is a caterpillar with $(n+1)n^{10} + n + 3$ leaves.



Now, denote $E' = \{(v_i, v_j), (v_j, v_i) : \{v_i, v_j\} \in E\}$, where we think of E' as the set of edges E , but where each edge appears in both directions. We define the set of syntenyies (the content of these syntenyies will be defined later — we are only listing them here)

$$\mathcal{X} = \{X_\gamma\} \cup \{X_{\alpha_i}, X_{\beta_i} : 1 \leq i \leq 3\} \cup \mathcal{X}_{E'} \cup \bigcup_{i \in [n+1]} \bigcup_{j \in [3]} \mathcal{X}_{i,j}^{leaf}$$

where $\mathcal{X}_{E'} = \{X_{ij} : (v_i, v_j) \in E'\}$ and for each $i \in [n+1], j \in [3]$, $\mathcal{X}_{i,j}^{leaf}$ is a set of n^{10} synteny trees that has exactly one synteny for each member of W_i^{leaf} . We put $s(X_{\alpha_i}) = \alpha, s(X_{\beta_i}) = \beta$ for each $1 \leq i \leq 3, s(X_\gamma) = \gamma$ and $s(X_{ij}) = v_i$ for each $(v_i, v_j) \in E'$. Hence the first subscript of a X_{ij} synteny indicates its species. Thus each species has 3 synteny trees, except γ which has one. Observe that $|\mathcal{X}_{E'}| = 3n$.

Instead of describing the set of input gene trees, we describe them as synteny trees directly. We thus omit the usual tilde symbol, e.g., we write T instead of \tilde{T} with the understanding that T is a synteny tree whose corresponding gene tree has a unique gene at each of its leaves. There is one main synteny tree T_{main} , illustrated on Figure 4. To obtain it, for each $j \in [3]$, define a tree T_{main}^j as the unique synteny tree on leafset $\{X_{\alpha_j}, X_{\beta_j}\} \cup \bigcup_{i \in [n+1]} \mathcal{X}_{i,j}^{leaf}$ that has only speciations. Another way to view T_{main}^j is that it is obtained by taking a copy of S , by removing the v_1, \dots, v_n species and relabeling the species at the leaves by the appropriate synteny trees for the index j . The tree T_{main} is then obtained by joining the T_{main}^j trees and X_γ as follows (using Newick notation):

$$T_{main} = (((T_{main}^1, T_{main}^2), T_{main}^3), X_\gamma)$$

Note that T_{main} has 2 duplication nodes. We will call d_1 and d_2 the two children of the lower duplication, and d_3 the other duplication child (see Figure 4). The rough idea behind our reduction is that T_{main} has 3 subtrees that each contain a synteny from each species, except v_1, \dots, v_n . These appear as losses in each T_{main}^j subtree. However, in \mathcal{X} we have three synteny trees for each species v_1, \dots, v_n , just enough to “fill-up” these losses. The main goal in our construction is to make this complete “fill-up” of the losses possible if and only if G is 3-colorable by making each T_{main}^j subtree represents a color. We will build additional input trees to realize this idea. Our goal is to enforce that X_{ij} and X_{ji} can only fill-up the losses under the same subtree.

We have another synteny tree T'_{main} defined as:

$$T'_{main} = (((X_{\alpha_1}, X_{\alpha_2}), X_{\alpha_3}), X_\gamma)$$

Note that T'_{main} does not provide any new information on the structure of a supertree, but will play a role later in the gene ordering in the synteny trees.

Then for each $v_i \in V$, let v_j, v_k, v_l be the neighbors of v_i and let

$$T_{ijk} = (((X_{ij}, X_{ji}), (X_{ik}, X_{ki})), X_\gamma)$$

$$T_{ijl} = (((X_{ij}, X_{ji}), (X_{il}, X_{li})), X_\gamma)$$

$$T_{ikl} = (((X_{ik}, X_{ki}), (X_{il}, X_{li})), X_\gamma)$$

For notational convenience, for all i, j, k we will say that $T_{ijk} = T_{ikj}$, i.e. both refer to the same tree. Note that for every edge $v_i v_j \in E$, X_{ij} occurs in exactly 4 trees (in 2 trees of the form T_{ijk} and 2 trees of the form T_{jik}). The set of synteny trees/families \mathcal{G} contains T_{main}, T'_{main} and all the T_{ijk} trees.

It only remains to define the given gene orders for the extant syntenies \mathcal{X} . We will define these orders as strings on alphabet \mathcal{G} directly. For $X \in \mathcal{X}$, let $t(X)$ be the set of synteny trees of \mathcal{G} that X appears in. Observe that X_γ appears in every synteny tree, and thus there is only one possible ancestral sequence of genes families. In other words, the X_γ synteny is a string Z over alphabet \mathcal{G} , and the order of every other synteny X is the subsequence of Z for the characters $t(X)$. We let Z be any string that begins with T_{main} and ends with T'_{main} . This completes the construction of our SUPER-RECONCILIATION instance.

We show that G admits a proper 3-edge coloring if and only if our instance formed by S, \mathcal{G} and order Z admits a Super-Reconciliation of total cost at most $3(n+1)n^{10} + 15n + 8$.

(\Rightarrow) For the first direction, suppose that G admits a proper 3-edge coloring. Let E_1, E_2, E_3 be the underlying partition of the edges into 3 color classes. Let D_1, D_2 and D_3 be the subtrees of T_{main} rooted at d_1, d_2 and d_3 , respectively. Notice that for each $l \in [3]$, the set $\{s(X) : X \in \mathcal{L}(D_l)\}$ is equal to $\mathcal{L}(S) \setminus V$. We “fill-up” each D_l subtree with a leaf from a synteny that belongs to v_j for every $v_j \in V$. More precisely, for each $l \in [3]$ and each $v_i \in V$, let $v_i v_j$ be the edge of E_l that is incident to v_i . We graft the corresponding synteny X_{ij} (from species v_i) onto D_l on the branch that makes its parent a speciation (this location is unique). Call T the synteny tree resulting from these graftings. In this manner, every D_l subtree has exactly one synteny from every v_i species. Moreover, every synteny gets grafted onto T , no new duplications are created and T has no full losses. It is not difficult to verify that the resulting tree displays all input trees (the T_{ijk} trees are displayed because two X_{ij} and X_{ji} syntenies will be grafted under the same D_l subtree, and this subtree is different from where X_{ik} and X_{ki} get inserted into, since $v_i v_j$ and $v_i v_k$ belong to a different E_l set).

Our tree T has 2 duplications and no full losses. To count partial losses, we assign the string Z to every internal node. Let X be any leaf of T that belongs to a species in W_i^{leaf} for some $i \in \{1, \dots, n+1\}$. The X synteny only appears in T_{main} and T'_{main} , so because Z starts with T_{main} and ends with T'_{main} , we may add a single partial loss on the branch from X to its parent. This amounts to $3(n+1)n^{10}$ segmental losses. The syntenies X_{α_i} and X_{β_i} can be handled similarly with a single loss on the branch to their parent. As $i \in \{1, 2, 3\}$, this amounts to 6 losses. The X_γ leaf does not incur any losses. As for a synteny $X_{ij} \in \mathcal{X}_{E'}$, recall that X_{ij} appears in at most 4 gene families. This can be handled by using at most 5 segmental losses on the path between X_{ij} and its parent. As $|\mathcal{X}_{E'}| = 3n$, this adds at most $15n$ losses. In total, we the total cost is $3(n+1)n^{10} + 15n + 8$, which is conveniently the number that we predicted.

(\Leftarrow) For the converse direction, let T be a supertree for \mathcal{G} on leafset \mathcal{X} that yields a reconciliation of cost at most $3(n+1)n^{10} + 15n + 8$. We assume that each internal node is assigned a gene family sequence that is a subsequence of Z .

Note that since T is compatible with T_{main} , T can be seen as a tree obtained by starting with T_{main} , then grafting some subtrees T_1, \dots, T_r successively onto some branches of T_{main} . To see this, let T' be the tree obtained from T by deleting every

node that does not have a descendant in $\mathcal{L}(T_{main})$. Then T' is the same tree as T_{main} , but with some nodes of degree 2 denoted t_1, \dots, t_r (excluding the root) that consist of the locations where the T_1, \dots, T_r trees were grafted (note that the roots of T_{main} and T must be the same due to X_γ). For each $i \in [r]$, we will assume that the root of T_i was grafted onto T_{main} under t_i , i.e. $r(T_i)$ is a child of the node t_i . Note that $\{\mathcal{L}(T_1), \dots, \mathcal{L}(T_r)\}$ forms a partition of $\mathcal{X} \setminus \mathcal{L}(T_{main})$. Under the above view of T , we will think of a node x of T_{main} as also a node of T , whether x is internal or leaf.

We show how to obtain a proper edge coloring of G . The proof is divided into a series of Claims, the first one showing that every leaf in a W_i^{leaf} must incur a segmental loss.

Claim 1 *Let $X \in \mathcal{X}_{i,j}^{leaf}$ be a synteny for which $s(X) \in W_i^{leaf}$ for some $i \in [n+1]$ and $j \in [3]$. Moreover let p_X be the parent of X in T_{main} . Then in T , there is either a duplication or a partial loss on the path from p_X to X .*

Proof Let Z_p be the gene family string assigned at p_X . Observe that in T , the p_X node has a descendant X_{α_i} for some $i \in [3]$. The family X_{α_i} appears in the trees T_{main} and T'_{main} , implying that both these families are in Z_p . Since X only appears in T_{main} , the T'_{main} character from Z_p must be lost on the path from p_X to X , either by a duplication or partial segmental loss. \square

As a consequence of Claim 1, there are at least $3(n+1)n^{10}$ duplications and/or partial segmental losses in T . Our strategy is the following: we will show that if the X_{ij} syntenies are not setup to “fill up” every hole in the d_1, d_2 and d_3 subtrees of T as in our solution for the converse direction, then there must be at least n^{10} full losses in T (as opposed to partial losses). As these were not counted in Claim 1, this would imply that T has $3(n+1)n^{10} + n^{10} > 3(n+1)n^{10} + 15n + 8$ losses, a contradiction. Recall that we view T as a obtained from T_{main} by grafting subtrees T_1, \dots, T_r on nodes t_1, \dots, t_r that were inserted on branches of T_{main} . We next show that the trees that get grafted onto T_{main} to obtain T all consist of syntenies from a single species.

Claim 2 *Let $X_{ij} \in \mathcal{X}_{E'}$, and let T_h be the tree grafted onto T_{main} that contains X_{ij} . If T_h has another leaf $X_{kl} \in \mathcal{X}_{E'}$, then $k = i$.*

Proof Suppose instead that $k \neq i$. Then $v_i = s(X_{ij}) \neq s(X_{kl}) = v_k$. We then have $s(r(T_h)) \geq lca_S(v_i, v_k)$. Assume without loss of generality that $k > i$. Observe that T_h cannot contain any leaf with a species in W_{i+1}^{leaf} , because all syntenies with a species in W_{i+1}^{leaf} are already in T_{main} . However, the path from $s(r(T_h))$ to v_i in S contains the set of nodes W_{i+1} (because $k > i$), where $|W_{i+1}| = n^{10}$. By the definition of reconciliation, each node in W_{i+1} must have at least one corresponding node in T_h on the path between $r(T_h)$ and X_{ij} , all of which must have a child that is a full loss in a node in W_{i+1}^{leaf} . It follows that T has at least n^{10} additional losses, a contradiction. \square

Recall that T_{main} (and hence T) has two duplication nodes with children d_1, d_2 and d_3 . Ignoring X_γ , these three nodes partition the leaves of T_{main} into 3 subsets. These will correspond to our edge colors. Towards this goal, for a synteny $X_{ij} \in \mathcal{X}_{E'}$, we will say that X_{ij} is of color 1 (respectively color 2 and 3) if it is a descendant of d_1 (respectively of d_2 and d_3). Note that X_{ij} has at most one color, but may have none - which we prove to not be the case.

Claim 3 *For each synteny $X_{ij} \in \mathcal{X}_{E'}$, X_{ij} has a unique color.*

Proof Suppose otherwise that there is some X_{ij} that has no color. Let T_h be the subtree grafted onto T_{main} that contains X_{ij} , with t_h the parent of $r(T_h)$. Since X_{ij} has no color, it follows that t_h must be an ancestor of d_1, d_2, d_3 or X_γ (all or some of these cases can hold simultaneously). In all cases, it is easy to see that $s(t_h) \geq lca_S(\alpha, \beta)$. Moreover by Claim 2, T_h has only leaves from the v_i species. It follows that on the path from t_h to X_{ij} , there is a loss for each node in W_{i+1} . Once again, this incurs n^{10} additional losses, a contradiction. \square

We then show that synteny from the same species get distinct colors, owing to the T_{ijk} trees.

Claim 4 *Let $X_{ij}, X_{ik} \in \mathcal{X}_{E'}$ be two distinct synteny from the same species v_i . Then X_{ij} and X_{ik} do not have the same color.*

Proof Suppose that X_{ij} and X_{ik} have color 1, without loss of generality. Let $T_{i'}, T_{j'}$ and $T_{k'}$ be the subtrees grafted onto T_{main} that contain X_{ij}, X_{ji} and X_{ik} , respectively (where $t_{i'}, t_{j'}, t_{k'}$ are the parents of $r(T_{i'}), r(T_{j'}), r(T_{k'})$, respectively). By Claim 2, we know that $T_{i'} \neq T_{j'} \neq T_{k'}$, although $T_{i'} = T_{k'}$ is possible. Recall that we have the tree $T_{ijk} = (((X_{ij}, X_{ji}), (X_{ik}, X_{ki})), X_\gamma)$ in the input. Since T displays T_{ijk} , this implies that $t_{k'}$ cannot be a descendant of $lca_T(t_{i'}, t_{j'})$, and therefore $T_{i'} \neq T_{k'}$. Also, because X_{ik} is of color 1, $t_{k'}$ must be a descendant of d_1 . Thus $t_{k'}$ is either (1) an ancestor of $lca_T(t_{i'}, t_{j'})$, or (2) $t_{k'}$ is on the path between a leaf $w \in \mathcal{X}_{l,1}^{leaf}$ and its parent $p(w)$ in T_{main} , where $l \in \{1, \dots, n+1\}$. In case (1), the only way that T can display T_{ijk} is if X_{ki} belongs to $T_{k'}$, along with X_{ik} . This contradicts Claim 2. In case (2), let $T_{k''}$ be the subtree grafted on T_{main} that contains X_{ki} . Due to the T_{ijk} tree, $t_{k''}$ must also be on the path between w and $p(w)$. Thus in the subtree of T rooted at $lca_T(X_{ik}, X_{ki})$, there is at most one leaf other than X_{ik} and X_{ki} (namely w). This subtree must contain at least $n^{10} - 1$ losses, either for the W_{k+1} nodes if $i > k$, or the W_{i+1} nodes if $k > i$. We reach the same contradiction. \square

It only remains to show that edge colors are consistent between their two directions.

Claim 5 *Let $X_{ij}, X_{ji} \in \mathcal{X}_{E'}$. Then X_{ij}, X_{ji} have the same color.*

Proof Let v_k, v_l be the neighbors of v_i other than v_j . Suppose that X_{ij} and X_{ji} do not have the same color. If one of X_{ij} or X_{ji} is of color 3 and the other of color 1

or 2, then because of the T_{ijk} tree in the input, X_{ik} and X_{ki} cannot be a descendant of any of d_1, d_2 or d_3 . Hence they have no color, contradicting Claim 3. So we may assume that X_{ij} and X_{ji} are of color 1 and 2 (not necessarily respectively). Again because of the T_{ijk} tree, X_{ik} and X_{ki} must be of color 3. And because of the T_{ijl} tree, X_{il} and X_{li} must also be of color 3. But then, X_{ik} and X_{il} are both of color 3, contradicting Claim 4. \square

We can now color the edges of E as follows: color $v_i v_j$ with color $c \in \{1, 2, 3\}$ if and only if X_{ij} and X_{ji} have color c . By Claim 3 and Claim 5, each edge gets assigned a unique color. Two adjacent edges $v_i v_j$ and $v_i v_k$ get assigned the colors of X_{ij} and X_{ik} . By Claim 4, X_{ij} and X_{ik} have different colors. It follows that the edge coloring is proper, concluding the proof.