Merging partially labelled trees: hardness and a declarative programming solution

Anthony Labarre Sicco Verwer

***Abstract*—Intraspecific studies often make use of haplotype networks instead of gene genealogies to represent the evolution of a set of genes. Cassens et al. [4] proposed one such network reconstruction method, based on the global maximum parsimony principle, which was later recast by the first author of the present work as the problem of finding a minimum common supergraph of a set of t partially labelled trees. Although algorithms were proposed for solving the problem on two graphs, the complexity of the general problem remains unknown. In this paper, we show that the corresponding decision problem is NP-complete for $t = 3$. We then propose a declarative programming approach to solving the problem to optimality in practice, as well as a heuristic approach, both based on the IDP system, and assess the performance of both methods on randomly generated data.**

I. Introduction

Phylogenetic trees are the traditional tool for representing the evolution of a given set of species [7]. The last two decades, however, have witnessed the emergence of a new way of reconstructing and representing evolution, which has become widespread in phylogenetic studies: *phylogenetic networks*, which generalise phylogenetic trees by allowing multiple paths between species. The main reason for using networks rather than trees is that evolution is not always tree-like: genes may be duplicated, transferred or lost, and recombination events (i.e. the breaking of a DNA strand followed by its reinsertion into a different DNA molecule) as well as hybridisation events (i.e. the combination of genetic material from several species) are known to occur. Moreover, even when evolution is tree-like, situations exist in which a relatively large number of tree topologies might be "equally good", and not enough information is available to discriminate between those trees. One proposed solution to the latter issue is the use of *consensus trees*, where the idea is to find a tree that represents a compromise between the given topologies; another approach, on which we focus in this paper, is to build a *network* [8, 10] that is compatible with all topologies of interest.

Haplotype networks are used in the context of *intraspecific* studies, which focus on relations between genes rather than between species. Cassens et al. [4] proposed a new method for reconstructing such networks, based on a given set of trees rather than on the input sequences. Note that the trees studied in that context, namely, *gene genealogies*, differ from the typical phylogenetic trees studied in comparative genomics: whereas phylogenetic trees are usually binary (i.e. internal nodes have degree three), have labels attached only to their leaves, and contain branches of arbitrary real length, gene genealogies allow internal nodes of arbitrary degree, as well as labelled nodes that are not leaves, and their branches have length exactly one. Cassens et al.'s

approach comprises two steps: *most parsimonious trees* are built from the sequences, and a subset of these trees is then merged into a graph. Their approach, which they refer to as "Union of Most Parsimonious trees", does not aim at building a smallest graph that contains *all* most parsimonious trees, as Bandelt et al. [1] did using *median networks*, but rather to summarise the information contained in a selected portion of those most parsimonious trees in a graph that is as "succinct" as possible.

The results produced by UMP on simulated data seemed promising, in comparison with a few traditional algorithms [4]. However, the algorithm and the overall approach proposed by the authors lacked a proper formalisation, and were later recast by the first author of the present work as a *minimum common supergraph problem*: given a set of partially labelled trees on the same label set, find a graph on the same vertex set which contains all input trees as subgraphs and which has as few edges as possible [11]. In the same work, two exact algorithms were given for the same problem on two partially labelled *graphs*, running in polynomial time under some assumptions and in exponential time in the general case. To the best of our knowledge, the complexity of the problem has since remained open.

In this work, we settle the complexity of the above optimisation problem, by showing that the associated decision problem is NP-complete for three trees. We make up for this bad news by proposing a practical approach to solving the problem to optimality in practice, using the IDP system [15]. This allows us to model our minimum common supergraph problem as a constraint satisfaction problem that is automatically translated into a SAT instance and then solved quickly by a SAT solver. We give an exact and a greedy method for UMP, both based on this declarative programming approach, and assess the performances of both approaches on random instances of various sizes.

II. Background

We recall here a few definitions and notation that will be needed in the study of our problem, formally stated at the end of this section. Any graph-theoretical concept the reader might lack familiarity with can be found in any textbook on the topic, e.g. Diestel [6].

Definition II.1. [11] An (n,k)-*graph* $G = (V, E, \mathcal{L})$, where $V(G) = V_l(G) \cup V_u(G)$ and $|V_l(G)| = k$, is a graph on n vertices, k of which are labelled.

We distinguish between the set $V_l(G)$ of labelled vertices and the set $V_u(G)$ of unlabelled vertices. Unless explicitly stated, all (n,k)-graphs will use the label set $\{1, 2, \ldots, k\}$ for labelled vertices. The $\mathcal{L}$ in the above definition corresponds to the following concept.

Definition II.2. [11] The *labelling* $\mathcal{L}$ assigns a distinct label to each vertex in $V_l(G)$; it is called a *partial* labelling if $k < n$ (in which case we say that G is *partially labelled*), and a *complete* labelling if $k = n$ (in which case we say that G is *completely labelled*).

Definition II.3. [11] An (n,k)-*tree* is a connected (n,k)-graph with $n-1$ edges and whose labelled vertex set includes all vertices of degree 1.

The following function, which (possibly) returns the label of vertex v in the (n,k)-graph G, allows us to adapt classical concepts from graph theory to our needs:

$$lab : V(G) \to \{1, 2, \ldots, k\} \cup \{\varnothing\}$$
$$: v \mapsto lab(v) = \begin{cases} i & \text{if } v \text{ has label } i, \\ \varnothing & \text{otherwise.} \end{cases}$$

This is not to be confused with the labellings introduced in Definitions II.1 and II.2: labelling $\mathcal{L}$ *assigns* labels to

vertices, while function lab (possibly) *returns* labels. We will also use lab on edges, in order to obtain the pairs of labels that correspond to the endpoints of interest: if v, $w \in V(G)$, then $lab(\{v, w\}) = \{lab(v), lab(w)\}$. Therefore, we have:

$$lab(E(G)) = \{\{i, j\} \mid i, j \in \{1, 2, \ldots, k\} \cup \{\varnothing\} \text{ and } \exists\, v, w \in V(G) : lab(v) = i, lab(w) = j\}.$$

We intentionally refrain from providing an explicit definition of the image of $lab(\{v,\ w\})$ for unlabelled vertices, because we will only use that function on pairs of labelled vertices.

Definition II.4. An (n, k)-graph G is a *subgraph* of an (n, k)-graph H if the labellings of G and H can be completed in such a way that the resulting (n, n)-graphs G' and H' satisfy $lab(E(G')) \subseteq lab(E(H'))$. In that case, we also say that H is a *supergraph* of G.

By "completing a labelling", we mean assigning distinct labels to the remaining unlabelled vertices; already labelled vertices must not be altered. We adapt the following definition from Bunke et al. [3] to our purposes.

Definition II.5. [11] A *common supergraph* of a set $\{G_1, G_2, \ldots, G_t\}$ of (n, k)-graphs is an (n, k)-graph G that is a supergraph of each G_i (for $1 \leq i \leq t$). It is *minimum* if there is no other graph G' with $|E(G')| < |E(G)|$ that shares this property.

Figure 1 shows two (n, k)-trees, along with two supergraphs G_1 and G_2 of those trees. G_2 is not minimum, since G_1 has fewer edges; and G_1 is indeed minimum, since $|E(G_1)| = |E(T_1)|+1 = |E(T_2)|+1$ and, although both trees are isomorphic, their labellings cannot be completed in such a way that $lab(E(T_1)) = lab(E(T_2))$.

We now have everything we need to formally state our problem as a decision problem:

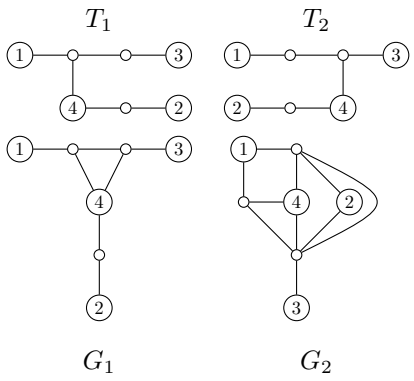

Fig. 1. Two $(7, 4)$-trees T_1 and T_2, and common supergraphs G_1 and G_2 of T_1 and T_2; G_1 is minimum, but G_2 is not.

COMMON SUPERGRAPH OF PARTIALLY LABELLED TREES (CS-PLT)

- **Instance:** (n, k)-trees T_1, T_2, ..., T_t on the same label set, a natural upper bound K.
- **Question:** can the labellings of T_1, T_2, ..., T_t be completed in such a way that $\cup_{i=1}^{t} lab(E(T_i)) \leq K$?

Note that a common supergraph of the input trees is defined exactly by the above union.

III. The complexity of CS-PLT

In this section, we prove the hardness of CS-PLT.

Theorem III.1. CS-PLT *is* NP*-complete for three trees.*

Proof: We present a reduction from MONOTONE 1-IN-3 SATISFIABILITY (see Schaefer [14]):

MONOTONE 1-IN-3 SATISFIABILITY

- **Instance:** a Boolean formula $\phi = C_1 \wedge C_2 \wedge \cdots \wedge C_n$ without negations over a set $\Sigma = \{\ell^1, \ell^2, \ldots, \ell^m\}$, with exactly three distinct literals per clause.
- **Question:** does there exist an assignment of truth values $f : \Sigma \to \{\text{TRUE}, \text{FALSE}\}$ such that exactly one literal is TRUE in every clause of ϕ?

a) The transformation: We encode instances of MONOTONE 1-IN-3 SATISFIABILITY using three trees, whose construction and purpose are explained below, and we illustrate the transformation on an example in Figure 2.

1) The first tree T_1 encodes the occurrences of literals in the MONOTONE 1-IN-3 SATISFIABILITY instance ϕ. It is constructed using a matrix indexed by the literals and clauses from ϕ. Every occurrence of a literal ℓ^j in a clause C_i is mapped onto a pair of nodes connected by an edge, where one node is a leaf labelled with L_i^j, which we call a *literal node*, and the other node is unlabelled. After creating these nodes for all literal occurrences, we connect the unlabelled nodes that share an edge with occurrences of the same literal by adding edges vertically in the matrix, i.e., in order of occurrence. The first occurrence of every literal is then connected to a root node R, which is itself connected to a TRUE node T and a FALSE node F (all three nodes are labelled).
2) In tree T_2, R is connected to three paths:
   a) a first path that consists of all $3n$ literal nodes;
   b) a second path, called the TRUE CHAIN, that contains n unlabelled nodes and the node labelled T at one end;
   c) a third path, called the FALSE CHAIN, that contains $2n$ unlabelled nodes and the node labelled F at one end.

   The first path is connected to node R, while the unlabelled extremities of the TRUE CHAIN and of the FALSE CHAIN are both connected to R. The TRUE CHAIN and the FALSE CHAIN represent a truth assignment to the literals in ϕ. This assignment is determined by labelling T_1 and T_2 in the CS-PLT instance: a literal ℓ^j in C_i represented by an unlabelled node u connected to L_i^j in T_1 is set to TRUE (resp. FALSE) if u is assigned the same label as a node from the TRUE CHAIN (resp. FALSE CHAIN) from T_2.
3) Tree T_3 overlaps for a large part with T_2. The only difference being that the TRUE CHAIN is split up and every unlabelled node from this chain is connected to T and three literal nodes from a unique clause. These edges thus encode the different clauses in ϕ. In addition, by limiting the number of allowed edges in a CS-PLT solution by a value K (see below), they encode the constraint that every clause contains exactly one TRUE literal.

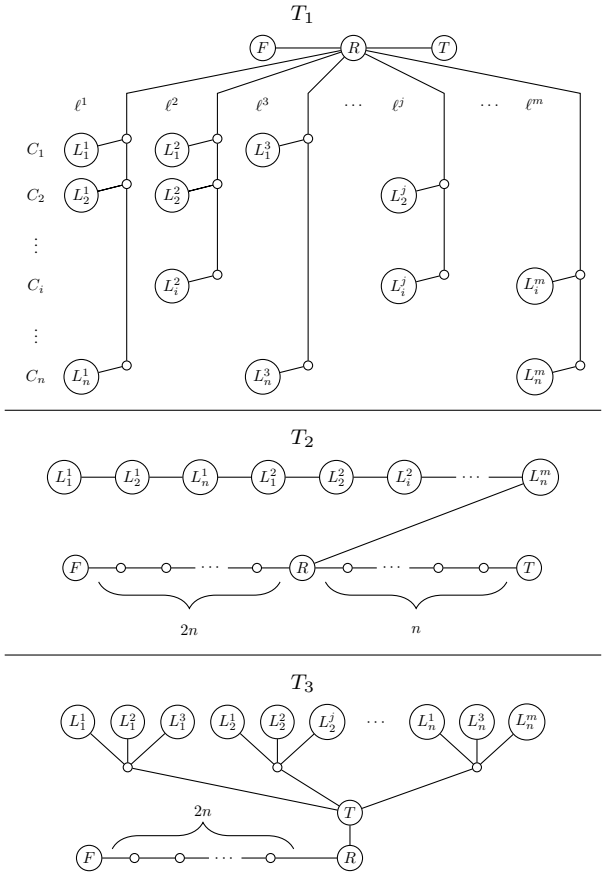

Fig. 2. The three trees built in our transformation.

Figure 3 shows an example of the construction applied to a small example instance. In addition to these trees, the CS-PLT decision problem requires an upper bound which is given as (derived later in the proof):

$$K = 12n + m + 1.$$

We now show that ϕ is satisfiable under the monotone 1-in-3 restrictions if and only if the labellings of these three trees can be completed in such a way that the union of the resulting labelled edge sets has size at most K.

($\Rightarrow$): Let f be a solution to ϕ. We use f to construct a solution to the CS-PLT instance of size at most K, which consists of three respective labellings for the unlabelled nodes of T_1, T_2 and T_3, as follows.

1) To every unlabelled node U_i^j connected to a literal node L_i^j in T_1, we assign the label $a(U_i^j)$ defined below and which corresponds to the number of literal nodes $L_{i'}^{j'}$ connected to unlabelled nodes $U_{i'}^{j'}$ representing either literals with smaller labels alphabetically (i.e. $l^{j'} < l^j$) or the same literal but occurring in an earlier clause (i.e. $l^{j'} = l^j$ and $i' < i$):

$$a(U_i^j) = |\{L_{i'}^{j'} \mid (l^{j'} < l^j) \vee (l^{j'} = l^j \wedge i' < i)\}|.$$

2) The k^{th} unlabelled node from the TRUE CHAIN $U_{k,T}$ in T_2 (ordered from R to T) receives the label assigned to the k^{th} unlabelled node U_i^j in T_1 (in ascending label order) that represents a TRUE literal:

$$(a(U_{1,T}), \ldots, a(U_{n,T})) = \text{SORT}(\{a(U_i^j) \text{ such that } f(l^j) = \text{TRUE}).$$

Since f is a 1-in-3 solution, it is guaranteed that this assigns a unique label to every unlabelled node from the TRUE CHAIN.

3) Similarly, the i^{th} node from the FALSE CHAIN in T_2 and T_3, namely, $U_{i,F}$, receives the label of the U_i^j nodes representing FALSE literals:

$$(a(U_{1,F}), \ldots, a(U_{2n,F})) = \text{SORT}(\{a(U_i^j) \text{ such that } f(l^j) = \text{FALSE}).$$

4) The k^{th} split up TRUE CHAIN nodes $U_{k,s}$ from T_3 are all connected to all three literal nodes from clause C_k. We label these nodes with the label of the U_k^j node representing the TRUE literal l^j in C_k:

$$a(U_{k,s}) = a(U_k^j) \text{ such that } f(l^j) = \text{TRUE}$$

This labelling is uniquely defined since f assigns the TRUE value to exactly one literal in every clause C_i. Figure 3 shows the completely labelled trees that result from applying the aforementioned steps to the trees shown in Figure 2.

We now show that these labellings yield a graph that contains exactly $K = 12n + m + 1$ edges. Every tree potentially adds all its $6n + 3$ edges to the resulting graph, so we derive K by counting the overlapping edges between the different trees starting with those from T_2:

in T_3: The $2n + 1$ edges connecting the FALSE CHAIN nodes to R overlap with those from T_2 since the unlabelled nodes are assigned exactly the same label by $a(\cdot)$, and 1 of the edges between T and the split up TRUE CHAIN overlaps with a TRUE CHAIN edge;

in T_1: Tree T_1 contains a lot of overlapping edges due to the $a(\cdot)$ labelling:

- n of the edges to literal nodes overlap with those from T_3 because every $U_{k,s}$ is assigned the same label as some U_k^j.
- All of the $3n - m$ edges connecting the unlabelled nodes U_i^j are shared by nodes representing literals that are assigned the same truth value by f and consecutive labels by a. Hence

these nodes are already connected by either the TRUE CHAIN or the FALSE CHAIN from T_2.

- 2 edges between R and newly labelled nodes overlap with those in T_2 since the TRUE CHAIN and FALSE CHAIN start with the smallest label assigned to a TRUE and a FALSE literal by a, corresponding to the first occurrence of these literals.
- 1 edge connecting R with T.

This sums up to $3(6n+2)-(2n+1+1)-(n+3n-m+2+1) = 18n+6-6n+m-5 = 12n+m+1$ edges, which equals K.

($\Leftarrow$)*:* If the constructed CS-PLT instance is true, then the original MONOTONE 1-IN-3 SATISFIABILITY instance is true. We first observe that the K edges determined above (see the ($\Rightarrow$) part) is the minimum number of edges resulting that can be obtained by labelling our three trees and taking the union of the resulting edge sets since we counted the maximum number of overlapping edges, making in total:

- 2 edges connecting R with T and F, which is minimal due to T_1;
- m edges connecting R with unlabelled nodes, which is minimal due to T_1;
- $5n$ edges between literal nodes and unlabelled nodes, minimal due to T_1 and T_3;
- $3n$ edges connecting literal nodes, minimal due to T_2;
- $2n-1$ edges between unlabelled nodes in the FALSE CHAIN, minimal due to T_2;
- $n-1$ edges between unlabelled nodes in the TRUE CHAIN, minimal due to T_2;
- n edges between T and unlabelled nodes in the split up TRUE CHAIN, minimal due to T_3;
- and 1 edge between F and an unlabelled node in the FALSE CHAIN, minimal due to T_2.

This sums up to $12n+m+1 = K$. We note that after assigning labels to the unlabelled nodes from T_1, and T_3, every literal node shares at least one and at most two edges with unlabelled nodes, independent of the labelling. Since the number of edges between literal nodes and unlabelled nodes does not influence any of the other counts, any additional edges between unlabelled and literal nodes will therefore result in strictly more than K edges.

In a solution of size K, we thus have exactly n literal nodes that share exactly a single edge with an unlabelled node. Consequently, for every split up TRUE CHAIN node from T_3, exactly one is assigned a label such that one if its three edges with literal nodes overlaps with an edge from T_1. The literal nodes connected to these overlapping edges determine the TRUE literals in the original satisfiability problem, all the other literals are set to FALSE. Since every split up TRUE CHAIN node in T_3 is connected to nodes representing exactly the literal occurrences of a single clause, this makes exactly one literal true in every clause. In addition to this property, we require that if a literal is true, then every instance of that literal is true. We show this by making the following observation that is key to our translation:

> A CS-PLT solution of size K has no edges between the TRUE CHAIN and the FALSE CHAIN.

To see why this holds, one only has to observe that in the above edge counts, we counted exactly $3n-2$ edges between unlabelled nodes. Since the TRUE CHAIN and the FALSE CHAIN in T_2 already contribute this amount of edges, any additional edge will yield a solution of size $K+1$. Thus, in a solution of size K, the same labels are assigned to the FALSE CHAIN nodes from T_2 and T_3. Furthermore, all of the edges between unlabelled nodes

from T_1 have to overlap with those from T_3. Since these edges connect the different occurrences of literals, these occurrences are all labelled with either TRUE CHAIN or FALSE CHAIN labels, but not both. Consequently, if the CS-PLT problem is true (has a solution of size K), then using our construction, every literal occurrence of the same literal is assigned the same truth value, and exactly one literal is set to TRUE in every clause, making the MONOTONE 1-IN-3 SATISFIABILITY instance satisfied.

b) Time complexity: The transformation clearly runs in time polynomial in the size of the MONOTONE 1-IN-3 SATISFIABILITY instance, and a solution to CS-PLT can easily be verified in polynomial time. The CS-PLT problem is therefore NP-complete. ■

IV. FINDING A MINIMUM COMMON SUPERGRAPH IN PRACTICE

Theorem III.1 deprives us of any hope of obtaining a polynomial-time algorithm for solving CS-PLT, but does not imply that we should abandon our search for efficient and exact solutions. We decided to use the efficient approach that consists in *translating* our problem into a constraint satisfaction problem, and to rely on an efficient SAT *solver* to obtain an exact solution to it, which allows us to make use of advanced solving techniques such as conflict analysis, intelligent back-jumping, and clause learning [2].

Figure 4 shows the typical workflow of a SAT solver based approach. We circumvent the difficulties pointed out in that workflow by relying on the IDP model expansion system [15], which is particularly convenient since it only requires us to provide a logical description of our problem. IDP translates the description into a constraint satisfaction problem, runs a solver, and translates the result back into a solution to our problem. Another attractive feature of IDP is that it can be used as an *anytime algorithm*: if one chooses to terminate the solving process before its completion, the system returns the best solution found so far.

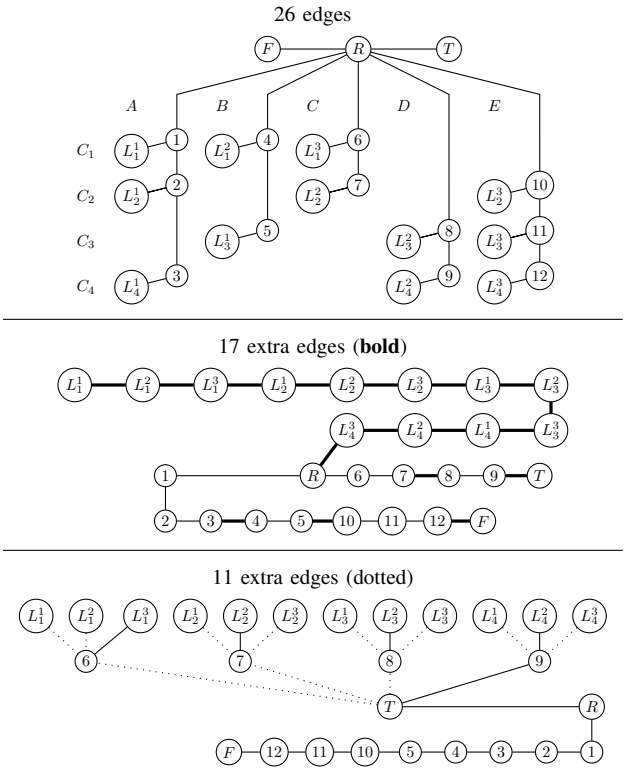

Fig. 3. A solution to the CS-PLT instance constructed from the MONOTONE 1-IN-3 SATISFIABILITY instance $(A \vee B \vee C) \wedge (A \vee C \vee E) \wedge (B \vee D \vee E) \wedge (A \vee D \vee E)$, which has as satisfying assignment $f(C) = f(D) = \text{TRUE}$. The union of the labelled edge sets has size $26 + 17 + 11 = 54 = 12n + m + 1$, with $n = 4$ and $m = 5$.

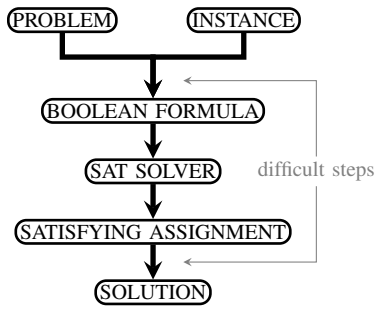

Fig. 4. The typical workflow of a SAT solver based approach.

We describe our solution in more detail in the following sections, starting with an introduction to SAT solvers in Section IV-A. We then describe IDP, its input and two models in Sections IV-B to IV-D, and explore their efficiency in practice on artificial data in Section IV-F.

A. Satisfiability and SAT solvers

The NP-complete *satisfiability problem*, which we recall below for completeness, is central to the field of computational complexity theory [5].

SATISFIABILITY (SAT)

- **Instance:** a Boolean formula ϕ in conjunctive normal form.
- **Question:** is there a satisfying assignment for ϕ?

SAT and its variants have spawned tremendous interest among researchers, who developed a number of practical and efficient algorithms, generally referred to as SAT *solvers*, for solving instances of those problems in practice (see e.g. Gomes et al. [9] for a recent account). A number of highly-optimised implementations exist, which make it possible to solve several well-known hard problems to optimality in a reasonable amount of time in many cases. One of the difficulties lies in formulating the problem as a satisfiability problem [13]. Luckily, the IDP system, described in the next section, makes that step a lot easier.

B. The IDP system

The IDP system [15] consists of two parts: a grounder [16] and a solver [12]. The grounder (GIDL) transforms a search or optimisation problem specified in IDP into a propositional theory that can be solved using the solver. The solver (MINISATID) then produces a solution to this theory, if one exists. This provides an easy method for declarative problem solving: all we have to do is provide a high-level specification of our problem; the IDP system then determines, using searches and heuristics, a good (efficient) formulation of this problem in propositional logic (i.e. as a satisfiability problem), and finally runs the solver, translating upon completion any solution it finds back to the high-level specification.

The IDP language is straightforward and easy to use, thanks to a multitude of logical operators, the ability to perform some arithmetic operations, and the possibility of providing inductive definitions. The latter in particular make it possible to define complex constraints or optimisation parameters in a neat and succinct way. Although such definitions would normally result in a blow-up of the propositional specification of the problem, the IDP solver contains specialised propagation mechanisms suitable for reasoning directly on such inductive definitions. These mechanisms are built on top of the popular MINISAT SAT solver without sacrificing much performance. The ability to write complex problem descriptions in just a few lines of code makes it an ideal tool for testing different problem specifications. In our opinion, this ease of use is the main strength of the IDP system.

C. A basic model

Figure 5 shows an IDP model we designed to represent the optimisation version of CS-PLT. This model is basic, but we show it nonetheless for clarity, and will improve it in Section IV-D. It consists of four sections:

1) the "Given:" section specifies the format in which the data should be given (in our case, a list of edges for each tree, along with some labels that are already assigned to a few vertices in each tree);
2) the "Find:" section describes the format of a solution (in our case, a set of labelled edges);
3) the "Satisfying:" section specifies the constraints edges and labels are subject to, and finally,

4) the "Minimize:" section describes which function should be optimised when searching for a solution (in our case, the size of the union of the completely labelled edge sets).

```
Given:
  type int Tree
  type int Node
  type int Color
  partial PreColor(Tree, Node) : Color    // some nodes are already labelled
  TEdge(Tree, Node, Node)

Find:
  Color(Tree, Node) : Color          // label the remaining nodes in each tree
  Edges(Color,Color)

Satisfying:
  { Edges(n,m) <- TEdge(t,x,y)           // once labelled, edges are assembled
                & Color(t,x) = n          // to build the common supergraph
                & Color(t,y) = m.
    Edges(n,m) <- Edges(m,n).                          // edges are undirected
  }

  ! t n : Color(t,n) = PreColor(t,n).   // extant labels must not be changed
  ! t c : ?1 n : Color(t,n) = c. // use each label exactly once in each tree

Minimize:
  #{ x[Color] y[Color] : Edges(x,y) }          // the size of the supergraph
```

Fig. 5. The code used by the IDP system to model the optimisation version of CS-PLT.

Specifying an instance of CS-PLT in that format is easy, and Figure 6 shows an example of a valid input, which consists of the following parts:

1) the `Tree` line specifies the unique indices from $\{1, 2, \ldots, t\}$ summarising our input (n, k)-trees T_1, T_2, ..., T_t;
2) the `Node` and `Color` lines specify the set $\{1, 2, \ldots, n\}$ of indices and labels used to refer to vertices;
3) the `PreColor` set specifies the labellings $\mathcal{L}_1$, $\mathcal{L}_2$, ..., $\mathcal{L}_t$, where `i, v -> b` means that vertex `v` in tree T_i has label `b`, and
4) the `TEdge` section specifies the set of edges in each tree, where `i, u, v` means that $\{u, v\} \in E(T_i)$.

```
Tree = { 1; 2; 3; 4; 5 }                              // ID's used for the trees
Node = { 1 .. 8 }                                     // ID's used for the vertices
Color = { 1 .. 8 }                                    // the range used for labels
PreColor = {                                  // the labelled nodes in each tree
        1, 1->4; 1, 5->2; 1, 7->1; 1, 8->3;
        2, 1->1; 2, 5->2; 2, 6->3; 2, 8->4;
        3, 1->4; 3, 2->2; 3, 4->1; 3, 8->3;
        4, 1->4; 4, 3->2; 4, 4->1; 4, 8->3;
        5, 1->3; 5, 3->1; 5, 7->2; 5, 8->4;
}
TEdge = {                                          // the set of edges in each tree
        1, 1, 3; 1, 6, 7; 1, 2, 8; 1, 1, 4; 1, 1, 6; 1, 2, 4; 1, 3, 5;
        2, 1, 2; 2, 4, 6; 2, 4, 8; 2, 5, 7; 2, 2, 3; 2, 3, 7; 2, 2, 4;
        3, 4, 7; 3, 6, 7; 3, 5, 7; 3, 3, 8; 3, 1, 5; 3, 3, 6; 3, 2, 5;
        4, 1, 2; 4, 4, 7; 4, 5, 6; 4, 5, 7; 4, 3, 6; 4, 2, 5; 4, 7, 8;
        5, 2, 7; 5, 2, 6; 5, 4, 8; 5, 4, 5; 5, 1, 5; 5, 3, 6; 5, 2, 5;
}
```

Fig. 6. An example of an instance of our problem formatted for use by the IDP system; in this case, the instance consists of five $(8, 4)$-trees.

The output of the IDP system, which takes as input our model and a data file representing an instance, consists of a list of edges expressed as pairs of vertex indices, together with complete labellings for each input tree. The union of those edges constitutes a (minimum) common supergraph of the input trees.

D. An improved model

The model described in Section IV-C can be used to solve CS-PLT instances, but lacks efficiency. We identify two reasons for this lack of speed: differently coloured solutions can yield isomorphic supergraphs, and the definition of edges produces an unnecessarily difficult SAT instance. We solve these issues by adding *symmetry breaking* predicates, and by defining the supergraph edges per tree instead of overall.

a) Symmetry breaking: Labellings are merely a way of identifying vertices in different trees; the actual labels do not matter, and permuting the labels assigned to the initially unlabelled vertices in any tree will not affect the size of the solution if we permute the corresponding labels in the other trees accordingly. Therefore, we can safely choose an arbitrary labelling for the unlabelled

vertices of any one tree in our instance, thereby reducing the search space by a factor of $(n-k)!$.

b) Supergraph edges per tree: The way edges are defined in the model of Figure 5 results in an instance that is difficult to solve, which makes the model inefficient. The reasons why a particular model is inefficient are unfortunately not always obvious; models that yield SAT instances with fewer clauses are usually regarded as more efficient, but sometimes larger models and redundant clauses have a positive effect on the runtime of a SAT solver. We identified by trial-and-error three inefficiencies in the definition of edges in the model of Figure 5, which we list and address below.

1) A first cause of inefficiency is the way in which the edges of the supergraph are specified as being undirected. In Figure 5, this is specified using the colours of nodes, and in an inductive way. Since these colours are free variables, and the nodes in a tree are fixed by the model input, it is more efficient to specify this property using these nodes instead of their colours. We do so by adding an additional declaration for undirected edges:

   `UEdge(i,u,v)`, which is TRUE if and only if `TEdge(i,u,v)` or `TEdge(i,v,u)` is TRUE.

   Constraints are then specified using the `UEdge` variables instead of the `TEdge` variables.
2) A second cause of inefficiency that we discovered is related to the way in which MINISATID makes use of the clauses. For reasons that remain to be investigated (likely due to propagation mechanisms), MINISATID is able to find satisfying assignments much more quickly when the edges of the common supergraph are specified per tree:

   `TreeEdge(i,n,m)`, which is TRUE if `UEdge(i,u,v)` is TRUE and `Color(i,v)=n` and `Color(i,u)=m`.

   An element of `Edges` is then TRUE if and only if there exists a corresponding `TreeEdge`.
3) A third and final cause of inefficiency is already visible in the `TreeEdge` definition. Instead of an equivalence constraint (if and only if), we require an implication (if). This means that a `TreeEdge(i,n,m)` can be TRUE even though the nodes with colours `n` and `m` are not joined by an edge in tree `i`. However, since the aim is to minimise the number of `Edges` and therefore the number of `TreeEdges`, this constraint is implicit in the model. Requiring `TreeEdges` (or `Edges`) to be false when there is no corresponding edge in a tree is redundant information. In our experience, removing this information results in an improved performance of MINISATID.

Figure 7 shows the improved model that we used in the experiments.

```
Given:
  type int Tree
  type int Node
  type int Color
  partial PreColor(Tree, Node) : Color    // some nodes are already labelled
  TEdge(Tree, Node, Node)

Find:
  Edges(Color, Color)
  Color(Tree,Node) : Color          // label the remaining nodes in each tree

Satisfying:
  ! t x y n m : (x < y & n < m & TEdge(t,x,y) & Color(t,x) = n
                & Color(t,y) = m) => Edges(n,m).
  ! t x y n m : (x < y & n < m & TEdge(t,x,y) & Color(t,x) = m
                & Color(t,y) = n) => Edges(n,m).
  ! n m       : n >= m => ~Edges(n,m).

  ! t n : Color(t,n) = PreColor(t,n).   // extant labels must not be changed
  ! t c : ?1 n : Color(t,n) = c. // use each label exactly once in each tree

Minimize:
  #{ x[Color] y[Color] : Edges(x,y) }          // the size of the supergraph
```

Fig. 7. An improvement over the model shown in Figure 5.

E. The greedy approach

In addition to the IDP model, we implemented the following greedy approach:

1) find a minimum common supergraph for every pair of trees using IDP;
2) merge the two trees that yield the smallest common supergraph G, and replace them with G;
3) for every remaining tree T, use IDP to compute a minimum common supergraph of T and G;
4) merge G with the tree that yields the smallest common supergraph G', and replace G and T with G';
5) go back to step 3 if any tree remains.

The greedy method focuses on merging single trees with the current common supergraph one at a time, which greatly reduces the search space. The idea of carrying out the merging process in a way that minimises the number of edges added at each step seems sensible, but it is not necessarily optimal: Figure 8 shows a small example where this approach performs suboptimally. An interesting open question is whether the ratio between the solution found using an optimal pairwise merging strategy and the optimal solution is bounded. In our experiments, the greedy method performed very well, significantly outperforming the exact approach on larger problem instances where the solver timed out before reaching an optimal solution.

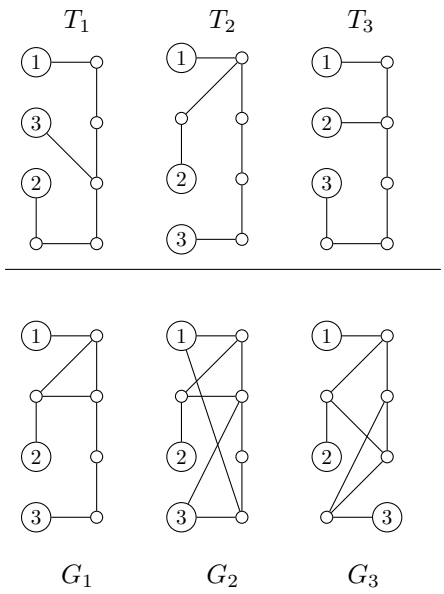

Fig. 8. An instance on which the pairwise approach performs suboptimally. The first step creates a minimum common supergraph G_1 of T_2 and T_3 that only requires a single additional edge, then creates a minimum common supergraph G_2 of G_1 and T_2 with 10 edges. However, G_3 is a common supergraph of T_1, T_2 and T_3 with only 9 edges.

F. Experimental results

For our experiments[1], we generated random CS-PLT instances of varying difficulty. We generated four different instances for every setting of the following parameters: 5, 10, or 20 trees; 10, 20, or 50 nodes per tree; and 5, 10, or 25 labelled nodes per tree. Unlabelled trees are generated by randomly adding edges between a growing connected component and an isolated vertex; since the number of leaves in the resulting tree may exceed the number of labels, we then modify it by repeatedly connecting random pairs of leaves (after disconnecting one leaf from each pair) until we have enough labels, which we then add, again randomly, starting with leaves and ending with internal nodes.

We tried to solve every generated instance using both the exact method and the greedy method. The exact method was given 2000 seconds of maximum runtime. Furthermore, since even pairwise merges can take a long time, the greedy method was given at most 10 seconds for every pairwise merge. Table I reports on the average sizes per parameter setting of the solutions found by both

[1]Experiments run on a desktop machine equipped with an Intel(R) Core TM i7 CPU 870 2.93GHz CPU (64bits) with 8GB of RAM.

methods.

			solution sizes	
#trees	#nodes	#labels	exact	greedy
5	10	5	17.50	18.00
10	10	5	19.50	21.50
20	10	5	23.00	25.25
5	20	5	34.75	**32.50**
5	20	10	53.00	**46.00**
10	20	5	38.75	**35.25**
10	20	10	64.25	**56.50**
20	20	5	42.25	42.25
20	20	10	75.50	**71.75**
5	50	5	130.00	131.25
5	50	10	128.00	132.75
5	50	25	207.75	**184.75**
10	50	5	183.75	**154.50**
10	50	10	177.75	**154.75**
10	50	25	270.00	**269.25**
20	50	5	241.50	**171.75**
20	50	10	232.00	**152.25**
20	50	25	346.25	**279.00**

TABLE I

AVERAGE SOLUTION SIZES OBTAINED BY THE EXACT AND THE GREEDY METHODS ON RANDOM INSTANCES WITH VARIOUS PARAMETERS AND PRESCRIBED TIMEOUTS. THE GREEDY APPROACH WAS ABLE IN SOME CASES (SHOWN IN **BOLD**) TO OUTPERFORM THE EXACT APPROACH.

IDP was able to solve all instances with 10 nodes per tree to optimality. In fact, all of these instances are solved optimally within approximately 10 seconds. No timeout occurs either in the pairwise greedy merges of these problems. As Table I shows, the greedy method performs worse on these instances, yielding solutions with two additional edges on average. None of the other instances are solved to optimality by IDP; the solver either times out (using all 2000 seconds), or runs out of memory. Interestingly, the greedy method performs very well on these larger instances, and we also investigated how the loss of quality evolves with the number of trees in the input. Figure 9 compares the sizes of the solutions obtained by the exact method and the greedy methods on random instances made of $(12, 6)$-trees, without any timeout. It can be seen that solutions obtained by the greedy method were at most 13% larger than those obtained by the exact method.

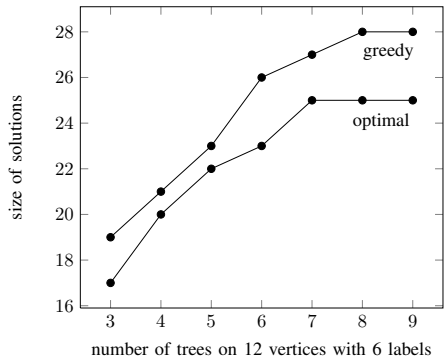

Fig. 9. Number of edges obtained by the exact and the greedy methods on random instances as the number of trees increases (no timeouts). The greedy approach produced solutions that were at most 13% larger than the optimal solution.

The quality of the solutions obtained by the greedy approach vastly exceeds that of the solutions obtained by the exact solver on the largest instances of Table I. Part of the reason for this is the memory required by IDP, due to the fact that the SAT solver keeps learning clauses while it runs. The solver eventually runs out of memory and returns the best solution found so far. Since this occurs frequently, even after running IDP for only 300 seconds, these solutions are worse than what IDP would have found in 2000 seconds. However, this only partially explains the differences: on some instances (e.g. those with 20 trees with 5 coloured nodes), IDP does reach the 2000 second time limit and still performs a lot worse than the greedy method. We therefore conclude that on large instances the pairwise approach is a very promising method for solving CS-PLT.

Figure 10 concludes our experiments and shows how

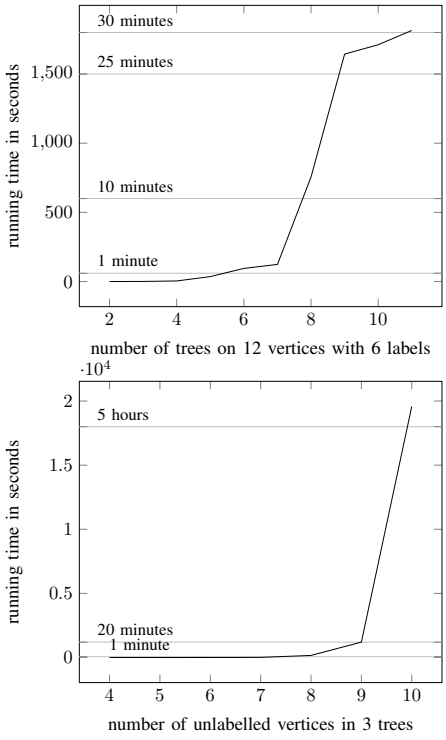

Fig. 10. Growth of the running time for finding an optimal solution (averages over 20 runs) with respect to the number t of trees or the number k of unlabelled vertices. Note that the search space has size $O((n-k)!^{t-1})$.

the running time of the exact solver grows with respect to the instance size, measured on the one hand by the number of trees in the instance and, on the other hand, by the proportion of unlabelled nodes in those trees.

V. Conclusions

In this work, we showed that the decision version of the problem of finding a minimum common supergraph of a given set of partially labelled trees is NP-complete. This negative result justifies and magnifies the importance of good approximate solutions to the original optimisation problem, as well as fast heuristics and exact algorithms for solving it in practice. In that regard, we investigated how promising the popular SAT solver-based approach could be in our case; we bypassed the difficulties that arise when trying to encode instances and problem descriptions as Boolean formulas by relying on the IDP system to handle the translation to a SAT instance and then to solve instances of our problem using a SAT solver. We proposed an optimised model that allowed us to obtain both an exact solution to our problem and a greedy approach that proved very useful in practice, yielding very high quality solutions much faster than the exact approach.

Several interesting theoretical questions arise. Most notably, the complexity of CS-PLT on two partially labelled trees remains open. Moreover, the computational complexity classification of CS-PLT could perhaps be further refined: in particular, does the problem admit a c-approximation algorithms for some constant c? Are there nice parameterisations of the problem that could prove useful in practice? The excellent performance of the greedy method justifies the importance of finding efficient algorithms for the pairwise case, since merging partial solutions in a greedy fashion usually gives solutions of high quality to the general problem. In addition, it would be interesting to further investigate the case where at least one of the input graphs is a graph instead of a tree, both from a complexity point of view and from an approximation point of view.

As far as practical aspects are concerned, fast and accurate solutions for real-world instances with actual data are still needed, especially in light of the problem's complexity. Future work will in particular investigate how the SAT solver-based approach proposed in this paper applies and scales in practice.

Finally, other considerations might need to be taken into account in order to assess the relevance of the results yielded by the UMP method in practice, which will require input from biologists. Are there other parameters that should be taken into account when searching for a

minimum common supergraph? Which criteria should be used to discriminate between nonisomorphic optimal solutions? We note that additional criteria could be easily incorporated directly into IDP, using the multitude of available logical operators and arithmetic operations.

ACKNOWLEDGMENTS

We wish to thank Broes De Cat and Johan Wittocx for explanations about the IDP system and for their help in improving the model shown in Figure 5.

REFERENCES

[1] H. J. BANDELT, P. FORSTER, B. C. SYKES, AND M. B. RICHARDS, *Mitochondrial portraits of human populations using median networks*, Genetics, 141 (1995), pp. 743–753.

[2] A. BIERE, M. HEULE, H. VAN MAAREN, AND T. WALSH, *Handbook of Satisfiability*, IOS Press, 2009.

[3] H. BUNKE, X. JIANG, AND A. KANDEL, *On the minimum common supergraph of two graphs*, Computing, 65 (2000), pp. 13–25.

[4] I. CASSENS, P. MARDULYN, AND M. C. MILINKOVITCH, *Evaluating intraspecific "network" construction methods using simulated sequence data: Do existing algorithms outperform the global maximum parsimony approach?*, Syst. Biol., 54 (2005), pp. 363–372.

[5] S. A. COOK, *The complexity of theorem-proving procedures*, in Proc. 3rd STOC, Shaker Heights, Ohio, USA, 1971, ACM, pp. 151–158.

[6] R. DIESTEL, *Graph theory*, vol. 173 of Graduate Texts in Mathematics, Springer-Verlag, Berlin, third ed., 2005.

[7] J. FELSENSTEIN, *Inferring Phylogenies*, Sinauer Associates, Sunderland, MA, 2004.

[8] P. GAMBETTE, *Who is who in phylogenetic networks: Articles, authors and programs*. Published electronically at `http://www.atgc-montpellier.fr/phylnet`.

[9] C. P. GOMES, H. KAUTZ, A. SABHARWAL, AND B. SELMAN, *Handbook of Knowledge Representation*, Foundations of Artificial Intelligence, Elsevier Science, 2007, ch. Satisfiability Solvers.

[10] D. H. HUSON, R. RUPP, AND C. SCORNAVACCA, *Phylogenetic Networks: Concepts, Algorithms and Applications*, Cambridge University Press, Nov. 2010.

[11] A. LABARRE, *Combinatorial aspects of genome rearrangements and haplotype networks*, PhD thesis, Université Libre de Bruxelles, Brussels, Belgium, Sept. 2008.

[12] M. MARIËN, J. WITTOCX, M. DENECKER, AND M. BRUYNOOGHE, *SAT(ID): Satisfiability of propositional logic extended with inductive definitions*, in Proc. 11th SAT, vol. 4996 of Lecture Notes in Computer Science, Guangzhou, China, May 2008, Springer, pp. 211–224.

[13] S. PRESTWICH, *Handbook of Satisfiability*, IOS Press, 2009, ch. CNF encodings.

[14] T. J. SCHAEFER, *The complexity of satisfiability problems*, in Proc. 10th STOC, San Diego, California, USA, May 1978, ACM, pp. 216–226.

[15] J. WITTOCX, M. MARIËN, AND M. DENECKER, *The IDP system: a model expansion system for an extension of classical logic*, in Proc. 2nd LaSh, Leuven, Belgium, Nov. 2008, pp. 153–165.

[16] J. WITTOCX, M. MARIËN, AND M. DENECKER, *Grounding FO and FO(ID) with bounds*, J. Artificial Intelligence Res., 38 (2010), pp. 223–269.