



HAL
open science

A General Framework for Enumerating Equivalence Classes of Solutions

Yishu Wang, Arnaud Mary, Marie-France Sagot, Blerina Sinimeri

► **To cite this version:**

Yishu Wang, Arnaud Mary, Marie-France Sagot, Blerina Sinimeri. A General Framework for Enumerating Equivalence Classes of Solutions. *Algorithmica*, 2023, 85 (10), pp.3003-3023. 10.1007/s00453-023-01131-1 . hal-04365403

HAL Id: hal-04365403

<https://inria.hal.science/hal-04365403>

Submitted on 28 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

A general framework for enumerating equivalence classes of solutions*

Yishu Wang[†], Arnaud Mary^{*}, Marie-France Sagot[‡], Blerina Sinimeri[§]

Abstract

When a problem has more than one solution, it is often important, depending on the underlying context, to enumerate (i.e., to list) them all. Even when the enumeration can be done in polynomial delay, that is, spending no more than polynomial time to go from one solution to the next, this can be costly as the number of solutions themselves may be huge, including sometimes exponential. Furthermore, depending on the application, many of these solutions can be considered equivalent. The problem of an efficient enumeration of the equivalence classes or of one representative per class (without generating all the solutions), although identified as a need in many areas, has been addressed only for very few specific cases. In this paper, we provide a general framework that solves this problem in polynomial delay for a wide variety of contexts, including optimization ones that can be addressed by dynamic programming algorithms, and for certain types of equivalence relations between solutions.

1 Introduction

Enumerating the solutions of an optimization problem solved by a dynamic programming algorithm (DP-algorithm) is a classical and well-known question. However, many enumeration problems have a huge number of solutions in practice, which might be an issue. From a computational point of view, since the number of solutions is a lower bound on the time complexity of any enumeration algorithm, it might make the algorithm impractical on real instances. Furthermore, even if the number of solutions is reasonable enough to be enumerated, the purpose of some applications is to give the output of the algorithm to a human specialist (this is necessary, for example, when some of the constraints of the problem are subjective and cannot be modeled).

Indeed, one of the advantages of an enumeration algorithm compared to an optimization one which in general outputs only one optimal solution, is to be able to understand the space of solutions. While this is important in many cases, no human can understand an output composed of billions of solutions.

*This is a preprint version of the following published paper: Wang, Y., Mary, A., Sagot, MF. et al. A General Framework for Enumerating Equivalence Classes of Solutions. *Algorithmica* (2023). <https://doi.org/10.1007/s00453-023-01131-1>

[†]Université de Lyon, Université Lyon 1, CNRS, Laboratoire de Biométrie et Biologie Evolutive UMR 5558, F-69622 Villeurbanne, France. Inria Grenoble Rhône-Alpes, Villeurbanne, France. yishu.wang@univ-lyon1.fr, arnaud.mary@univ-lyon1.fr

[‡]Inria Grenoble Rhône-Alpes, Villeurbanne, France. Université de Lyon, Université Lyon 1, CNRS, Laboratoire de Biométrie et Biologie Evolutive UMR 5558, F-69622 Villeurbanne, France. marie-france.sagot@inria.fr

[§]Luiss University, Rome, Italy. Inria Grenoble Rhône-Alpes, Villeurbanne, France. Université de Lyon, Université Lyon 1, CNRS, Laboratoire de Biométrie et Biologie Evolutive UMR 5558, F-69622 Villeurbanne, France. bsinimeri@luiss.it

The approach generally used to address this consists of enumerating all solutions, and then applying some type of clustering (grouping) algorithm to the set of optimal solutions. The final output presented to the user would then be some “representative description” of the clusters (groups) themselves. However, since the number of solutions is a lower bound for the total execution time of any enumeration algorithm, the first step of such a strategy becomes impossible when the number of solutions is too big. A natural question is then whether it would be possible to enumerate directly what we just called a “representative description” of the clusters of solutions. This could be for instance an element per cluster. Sometimes a cluster can also be seen as a set of characteristics that the solutions within the cluster share. In such a case, the representative description of a cluster could then be such a set of characteristics. A particularly convenient situation is however when the clusters correspond to equivalence classes of an equivalence relation over the set of solutions that we could establish *a priori*. The output could be in this case the quotient space of the equivalence relation.

Notice that the enumeration of equivalence classes of solutions is a combinatorial problem that could be solved exactly given a well-defined equivalence relation, and unlike data analysis methods such as incremental clustering, it does not require the definition of a similarity or dissimilarity measure between solutions which, depending on the mathematical nature of the solutions (numerical values, graphs, functions on graphs, etc.), can be difficult to define or costly to compute.

The problem this paper addresses is how to perform the task of enumerating equivalence classes of solutions with polynomial delay for a wide variety of problems (including optimization problems solved by dynamic programming algorithms), for certain types of equivalence relations between solutions.

The problem of enumerating equivalence classes, and particularly the generation of representative solutions is a challenge in the context of enumeration algorithms. It has been identified as a need in different areas, such as Genome Rearrangements [10], Artificial Intelligence [1] or Pattern Matching [7, 26]. It was listed as an important open problem in a recent Dagstuhl workshop on “Algorithmic Enumeration: Output-sensitive, Input-Sensitive, Parameterized, Approximative” (see, e.g., Sections 4.2 and 4.10 in [15]). To the best of our knowledge, this challenge has been addressed only for some few specific problems in the literature (e.g., [2, 10, 24, 25]).

To enumerate equivalence classes, we go through an intermediate problem, namely the enumeration of colored subtrees in acyclic decomposable AND/OR graphs (ad-AND/OR graph). The paper is organized as follows: Section 2 provides an algorithm to enumerate with polynomial delay colored subtrees in ad-AND/OR graphs; Section 3 details how this algorithm applies to the enumeration of equivalence classes in DP-problems. In that direction, we present some examples from well-known optimization problems in the literature. Finally, in Section 4 we conclude with some open problems.

2 Enumeration of colored subtrees in an acyclic decomposable AND/OR graph

2.1 AND/OR graphs and solution subtrees

An *AND/OR graph* (see, for example [23, 27]) is a well-known structure in the field of Logic and Artificial Intelligence (AI) that represents problem solving and problem decomposition. In this paper, we consider a particular flavor of AND/OR graphs known as *explicit AND/OR graphs for trees* [12].

This is a directed acyclic graph (DAG) G which explicitly represents an *AND/OR state space*

for solving a certain problem by decomposing it into subproblems. The set of nodes (or states) $S := V(G)$ contains OR and AND nodes (the OR nodes represent alternative ways for solving the problem while the AND nodes represent problem decomposition into subproblems, all of which need to be solved). There is a set of goal nodes $S_g \subseteq S$ and a set of start nodes $S_0 \subseteq S$ representing respectively the terminal states and the initial states. The children (out-neighbors) of an OR node are AND nodes, and the children of an AND node are OR nodes or goal nodes. We say that a node is an OR^+ node when it is either an OR node or a goal node. Furthermore, the AND/OR graphs that we consider must have the property of being *decomposable* (they can model a problem for which every decomposition yields disjoint subproblems that can be solved independently): for any AND node, the sets of nodes that are reachable from each one of its child nodes are pairwise disjoint. The example graph in Figure 1 is decomposable.




Figure 1: An acyclic decomposable AND/OR graph with four start nodes. Squares are OR^+ nodes (OR nodes or goal nodes); crossed circles are AND nodes. One solution subtree of size 8 is shown in bold.

Formally, in this paper, any graph that satisfies the properties in Definition 1 will be called an ad-AND/OR graph. Notice that this definition corresponds only to a particular case of the general AND/OR graphs in the AI literature; the latter may be neither acyclic nor decomposable.

Definition 1 (ad-AND/OR graph). *A directed graph G is an acyclic decomposable AND/OR graph, henceforth denoted by ad-AND/OR graph, if it satisfies the following:*

- G is a DAG.
- G is bipartite: its node set $V(G)$ can be partitioned into $(\mathcal{A}, \mathcal{O})$ so that all arcs of G are between these two sets. Nodes in \mathcal{A} are called AND nodes; nodes in \mathcal{O} are called OR^+ nodes.
- Every AND node has in-degree at least one and out-degree at least one. The set of nodes with out-degree zero is then a subset of \mathcal{O} and is called the set of goal nodes; the remaining OR^+ nodes are simply the OR nodes. The subset of OR nodes of in-degree zero is the set of start nodes.
- G is decomposable: for any AND node, the sets of nodes that are reachable from each one of its child nodes are pairwise disjoint.

Definition 2 (solution subtree). *A solution subtree T of an ad-AND/OR graph G is a subgraph of G which: (1) contains exactly one start node; (2) for any OR node in T it contains one of its child nodes in G , and for any AND node in T it contains all its children in G .*

It is immediate to see that a solution subtree is indeed a subtree of G : it is a rooted tree, the root of which is a start node. If we would drop the requirement of G being decomposable, the object defined in Definition 2 would not be guaranteed to be a tree. One solution subtree of the example graph in Figure 1 is shown in bold.

The set of all solution subtrees of G is denoted by $\mathbb{T}(G)$. Given an ad-AND/OR graph G , counting the number of its solution subtrees and enumerating all solution subtrees can be solved by folklore approaches based on depth-first search (DFS).

Before going further, we recall that the motivation of this paper is concerned with solutions of dynamic programming problems. The correspondence between the solutions of DP-style recurrence equations and the solution subtrees of general AND/OR graphs has been formally proven in [16]. In the case where the underlying graph is acyclic, the recurrence equations can be solved efficiently by DP-algorithms. While we will now concentrate on the solution subtrees of an ad-AND/OR graph and on the equivalence classes of solution subtrees, we will demonstrate in Section 3 how to apply our algorithms to analyze equivalence classes of solutions of a very general class of problems solvable by DP. It is important to point out that solution subtrees of general AND/OR graphs are equivalent to various other well-known formalisms, e.g., acceptance trees of a nondeterministic tree automata, languages of regular tree grammars, complete subcircuits of tropical circuits. The reader who is more familiar with those may also find this paper interesting even outside of a dynamic programming context.

2.2 Equivalence classes

Let G be an ad-AND/OR graph. Let C be an **ordered** set of colors. We will consider equivalence relations on the set of solution subtrees of G which are based on a local comparison of the colors of the OR^+ nodes. Intuitively, two OR^+ nodes having the same color represent two alternative ways of solving the problem that can be considered equivalent.

Definition 3 (e-coloring). *An ad-AND/OR graph G is e-colored if its OR^+ nodes are colored in such a way that for any AND node all its children have distinct colors.*

Notations If s is a OR^+ node of G , we denote by $c(s)$ its color. If s is an AND node, we denote by $\tilde{C}(s)$ the tuple of colors of the children of s sorted in increasing order of the colors. If $T_1 \in \mathbb{T}(G)$ is a solution subtree of G , we use the notation $\pi(T_1)$ for the result of contracting the AND nodes in T_1 : for each OR node s of T , contract the only child node of s in T (i.e., remove the child and connect s to each one of its “grandchildren”).

Definition 4 (equivalence class). *A node-colored rooted tree T is an equivalence class of solution subtrees of an e-colored ad-AND/OR graph G (or shortly an equivalence class of G) if there exists a solution subtree T_1 of G such that $\pi(T_1)$ is equal to T . Such a T_1 is said to be a solution subtree belonging to the class T .*

More notations We denote by $\mathbb{C}(G)$ the set of equivalence classes of G . The notation π can be seen as a function $\pi: \mathbb{T}(G) \rightarrow \mathbb{C}(G)$. We denote by $\pi^{-1}(T) := \{T_1 \in \mathbb{T}(G) \mid \pi(T_1) = T\}$ the subset of solution subtrees of G belonging to the class T . The notations $c(s)$ and $\tilde{C}(s)$ are naturally extended to the case where s is a node in an equivalence class T . The root node of a rooted tree T is denoted by $r(T)$. The set of the children of a node s is denoted by $Ch(s)$.

An example of an e-colored ad-AND/OR graph with five equivalence classes is given in Figure 2.




Figure 2: An e-colored ad-AND/OR graph and its five equivalence classes. The colors of the OR^+ nodes are written inside the squares.

2.3 Enumerating equivalence classes

Given an e-colored ad-AND/OR graph G , we propose a polynomial delay algorithm to enumerate all equivalence classes of G . Given a total ordering c_1, \dots, c_m of the colors of G , we define a total ordering \prec over $\mathbb{C}(G)$, the set of equivalence classes of G . If T and T' have their roots colored differently, we say that T is smaller than T' , denoted by $T \prec T'$, if the root color of T precedes the one of T' . If T and T' have the same root color, let (T_1, \dots, T_k) (resp. (T'_1, \dots, T'_ℓ)) be the child subtrees of $r(T)$ (resp. $r(T')$) sorted recursively with respect to \prec . We then say that T is smaller than T' if the tuple (T_1, \dots, T_k) is lexicographically smaller than (T'_1, \dots, T'_ℓ) , i.e., if $T_i \prec T'_i$ with i being the smallest index such that $T_i \neq T'_i$. We also assume that \emptyset is smaller than any tree, and therefore a single node tree colored with color c comes before any other tree whose root is colored with c in \prec .

2.3.1 Definitions and notations

Recall that given an AND-node x , $\tilde{C}(x)$ is the tuple of colors of the children of x sorted in increasing order. Given an OR-node o , we denote by $\mathcal{T}(o)$ the set of color tuples of its children, i.e., $\mathcal{T}(o) := \{\tilde{C}(x) : x \in \text{Ch}(o)\}$. In other words, a color tuple (c_1, \dots, c_j) belongs to $\mathcal{T}(o)$ if o has an AND-child node whose children are colored with (c_1, \dots, c_j) . If we consider an equivalence class T of $\mathbb{C}(G/\{o\})$ rooted at o , the tuples of $\mathcal{T}(o)$ are precisely the possible colorings of the children of $r(T)$. Indeed, if the AND-child node $x \in \text{Ch}(o)$ is chosen in a solution subtree, then $\tilde{C}(x)$ will be the colors of the children of o in that solution. Notice that several AND-children nodes of o may have the same color tuple.

We extend this definition to a set \mathcal{O} of OR^+ nodes with $\mathcal{T}(\mathcal{O}) = \bigcup_{o \in \mathcal{O}} \mathcal{T}(o)$. In the same way, $t = (c_1, \dots, c_j)$ is a color tuple of $\mathcal{T}(\mathcal{O})$ if and only if there exists an equivalence class T of $\mathbb{C}(G/\mathcal{O})$ such that the children of $r(T)$ are colored with (c_1, \dots, c_j) . Given a set \mathcal{O} of OR^+ nodes, we denote by $t_1, \dots, t_{|\mathcal{T}(\mathcal{O})|}$ the different color tuples of $\mathcal{T}(\mathcal{O})$ ordered lexicographically, and we denote by $\text{Ch}^\ell(\mathcal{O})$ the set of AND-nodes in $\text{Ch}(\mathcal{O})$ whose color tuple is t_ℓ , i.e., $\text{Ch}^\ell(\mathcal{O}) = \{x \in \text{Ch}(\mathcal{O}) : \tilde{C}(x) = t_\ell\}$. The sets $\text{Ch}^1(\mathcal{O}), \dots, \text{Ch}^{|\mathcal{T}(\mathcal{O})|}(\mathcal{O})$ form a partition of $\text{Ch}(\mathcal{O})$, each part corresponding to a color tuple $t_i \in \mathcal{T}(\mathcal{O})$.

Finally, given a color tuple $t_\ell := (c_1, \dots, c_j) \in \mathcal{T}(\mathcal{O})$, for each $i \leq j$, by Definition 3, each node of $\text{Ch}^\ell(\mathcal{O})$ has exactly one child with color c_i . We denote by C_i^ℓ the set of children of $\text{Ch}^\ell(\mathcal{O})$

Algorithm 1: Next solution

```
1 Input: A set  $\mathcal{O}$  of  $\text{OR}^+$  nodes having all the same color  $c$  and an equivalence class  $T$  of  $G/\mathcal{O}$ 
2 Output: The equivalence class  $T'$  of  $G/\mathcal{O}$  that follows  $T$  w.r.t the  $\prec$  ordering.
3 Function  $\text{Next}(T, \mathcal{O})$ :
4   if  $T = \emptyset$  and  $\mathcal{O}$  contains terminal nodes then
5     | Return A tree with a single root node colored with  $c$ 
6   end
7    $r \leftarrow 0$ 
8   if  $T = \emptyset$  or  $T$  is a single node tree then
9     |  $r \leftarrow r + 1$ 
10    | if  $r > |\mathcal{T}(\mathcal{O})|$  then
11      | Return  $\perp$ 
12    | end
13    | Let  $(c_1, \dots, c_j)$  be the color tuple of  $t_r \in \mathcal{T}(\mathcal{O})$  and let  $T_1, \dots, T_j \leftarrow \emptyset$ 
14    |  $\mathcal{O}_1 \leftarrow C_1^r$ 
15    |  $\ell \leftarrow 1$ 
16  | else
17    | Let  $r$  be such that  $t_r \in \mathcal{T}(\mathcal{O})$  is the root color tuple  $\tilde{C}(r(T))$ 
18    | Let  $(T_1, \dots, T_j)$  be the child subtrees of  $r(T)$ , the roots of which are colored
19    |   respectively with  $(c_1, \dots, c_j) := t_r$ 
20    | For all  $i \leq j$ , let  $\mathcal{O}_i \subseteq C_i^r$  be the set of nodes in  $C_i^r$  compatible with  $(T_1, \dots, T_{i-1})$ 
21    | Let  $\ell$  be the largest index  $i \leq j$  such that  $\text{Next}(T_i, \mathcal{O}_i) \neq \perp$  if such index exists.
22    |   Otherwise,  $T \leftarrow \emptyset$  and go to line 8
23  | end
24  |  $T_\ell \leftarrow \text{Next}(T_\ell, \mathcal{O}_\ell)$ 
25  | for  $\ell < i \leq j$  do
26  |   | Let  $\mathcal{O}_i \subseteq C_i^r$  be the set of nodes in  $C_i^r$  compatible with  $(T_1, \dots, T_{i-1})$ 
27  |   |  $T_i \leftarrow \text{Next}(\emptyset, \mathcal{O}_i)$ 
28  | end
29  | Return A tree with root color  $c$  and root child subtrees  $(T_1, \dots, T_j)$ 
```

colored with c_i , i.e., $C_i^\ell = \{o \in \text{Ch}(x) : x \in \text{Ch}^\ell(\mathcal{O}), c(o) = c_i\}$ (it is a set of “grandchildren” of \mathcal{O}).

In the left panel of Figure 3, an example graph is shown where each node is labeled by an integer. The colors are, in increasing order, w , x , y , and z . For $\mathcal{O} = \{1, 2\}$, the set $\mathcal{T}(\mathcal{O})$ contains the three tuples $t_1 = (w, y)$, $t_2 = (x, y)$, and $t_3 = (x, y, z)$. We have $\text{Ch}^1(\mathcal{O}) = \{6\}$, $\text{Ch}^2(\mathcal{O}) = \{4, 5\}$, and $\text{Ch}^3(\mathcal{O}) = \{3\}$. The sets C_i^ℓ are $C_1^1 = \{12\}$, $C_2^1 = \{11\}$, $C_1^2 = \{9, 10\}$, $C_2^2 = \{8, 11\}$, $C_1^3 = \{9\}$, $C_2^3 = \{8\}$, and $C_3^3 = \{7\}$.

2.3.2 Algorithm description

Notice that by definition of \prec , given a set of OR^+ nodes \mathcal{O} , all having the same color c , and a color tuple $t_\ell \in \mathcal{T}(\mathcal{O})$, all equivalence classes T of G/\mathcal{O} such that $\tilde{C}(r(T)) = t_\ell$ are consecutive with respect to \prec .

The algorithm outputs the equivalence classes in ascending order with respect to \prec . Given an equivalence class T of G/\mathcal{O} for a set of OR^+ nodes \mathcal{O} of color c , it will output the equivalence

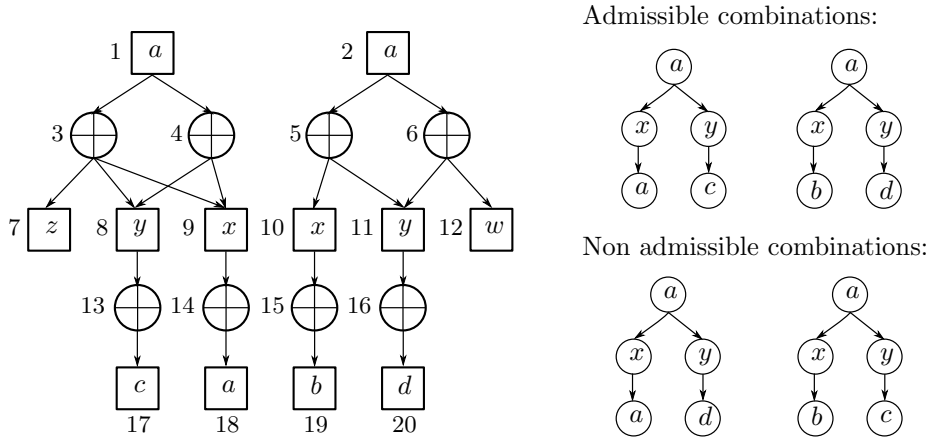


Figure 3: Left panel: An e-colored ad-AND/OR graph. Right panel: For $\mathcal{O} = \{1, 2\}$, there are four combinations between $\mathbb{C}(G/C_1^2)$ and $\mathbb{C}(G/C_2^2)$; only two of them are admissible.

class T' of G/\mathcal{O} that succeeds T w.r.t. \prec if it exists or output the symbol \perp if T is the last solution.

Assume that the children of $r(T)$ are colored with the *root color tuple* $(c_1, \dots, c_j) =: t_r \in \mathcal{T}(\mathcal{O})$ and let (T_1, \dots, T_j) be the child subtrees of T , the roots of which are colored with the tuple t_r . Notice that for all $i \leq j$, T_i is an equivalence class of G/C_i^r . The algorithm will output the next equivalence class T' such that $\tilde{C}(r(T')) = t_r$ if there remains one (same color tuple at the root), or it will output the first solution such that $\tilde{C}(r(T')) = t_{r+1} \in \mathcal{T}(\mathcal{O})$ otherwise (the next color tuple at the root).

To find the next solution corresponding to the root color tuple t_r , the algorithm will replace recursively T_j by its successor T'_j w.r.t. \prec if there exists one. We obtain the solution T' whose subtrees are $(T_1, \dots, T_{j-1}, T'_j)$ which is by definition the successor of T in \prec whenever T'_j is the successor of T_j . If T_j has no successor (that is, if it is the last one), we replace if possible T_{j-1} by its successor T'_{j-1} and we replace T_j by the smallest admissible solution (i.e., the successor of \emptyset). In general, we select at each step the greatest index ℓ such that T_ℓ has a successor w.r.t. \prec , we replace it by its successor T'_ℓ and we take the smallest admissible solution for every $\ell < i \leq j$.

Without further care, the above described procedure would output solutions whose child subtrees (T_1, \dots, T_j) of the root correspond to the elements of the Cartesian product of $\mathbb{C}(G/C_i^r)$, $i \leq j$. However, while it is true that if T is a solution, its child subtree T_i is an equivalence class of G/C_i^r for all $i \leq j$, the converse is not true. Indeed, not all elements of $\mathbb{C}(G/C_1^r) \times \dots \times \mathbb{C}(G/C_j^r)$ lead to an admissible solution (an example is given in the right panel of Figure 3). In order to find an admissible solution, we should guarantee that the choice of a given T_i is *compatible* with the previous choices (T_1, \dots, T_{i-1}) . This is done by selecting the subset of OR^+ nodes $\mathcal{O}_i \subseteq C_i^r$ that are compatible with (T_1, \dots, T_{i-1}) (see Definition 5 below). An admissible choice of T_i will then be any equivalence class of G/\mathcal{O}_i . The two key properties are that the set \mathcal{O}_i can be easily computed, and that it is never empty, i.e., there is always a choice for T_i that is compatible with the previous choices of (T_1, \dots, T_{i-1}) (there is at least one choice that corresponds to the current solution). Notice that if the latter were not true, the algorithm would not have a polynomial delay complexity since we may spend exponential time without reaching a final solution. With this property, we are guaranteed that we can always extend a partial tuple (T_1, \dots, T_i) until we reach a complete tuple (T_1, \dots, T_j) that will form a solution.

Compatible nodes Given a set of OR^+ nodes \mathcal{O} all colored with the same color c and a tree T of $\mathbb{C}(G/\mathcal{O})$, we denote by $r(\pi^{-1}(T)) := \{r(S) : S \in \pi^{-1}(T)\}$ the subset of OR^+ nodes of \mathcal{O} , each one of which is the root of a solution subtree of class T . The following definition formalizes the notion of compatible nodes mentioned previously.

Definition 5. Let \mathcal{O} be a set of OR^+ nodes of color c , $t_r =: (c_1, \dots, c_j) \in \mathcal{T}(\mathcal{O})$, and let T_1, \dots, T_k , with $k < j$, be respectively equivalence classes of $\mathbb{C}(G/C_i^r)$ for all $i \leq k$. We say that a node $o \in C_{k+1}^r$ is compatible with (T_1, \dots, T_k) if there exists an AND-node $x \in Ch^r(\mathcal{O})$ such that o is a child of x and such that $r(\pi^{-1}(T_i))$ contains a child of x for all $i \leq k$.

2.3.3 Analysis

Lemma 6. Let \mathcal{O} be a set of OR^+ nodes of color c , $t_r =: (c_1, \dots, c_j) \in \mathcal{T}(\mathcal{O})$. Let $T \in \mathbb{C}(G/\mathcal{O})$, let T_1, \dots, T_k , $k < j$ be its first k child subtrees with $T_i \in \mathbb{C}(G/C_i^r)$ for all $i \leq k$ and let $\mathcal{O}_{k+1} \subseteq C_{k+1}^r$ be the set of OR^+ nodes compatible with (T_1, \dots, T_k) . Given $T_{k+1} \in \mathbb{C}(G/C_{k+1}^r)$, there exists a tree $T' \in \mathbb{C}(G/\mathcal{O})$ whose first $k+1$ child subtrees are $(T_1, \dots, T_k, T_{k+1})$ if and only if $T_{k+1} \in \mathbb{C}(G/\mathcal{O}_{k+1})$.

Proof. (\Rightarrow) Assume that there exists a tree $T' \in \mathbb{C}(G/\mathcal{O})$ whose first $k+1$ child subtrees are $(T_1, \dots, T_k, T_{k+1})$, and let C be a solution subtree of G/\mathcal{O} such that $\pi(C) = T'$. Let x be the AND-child node of the root of C . Notice that since we assumed that $T_i \in C_i^r$ for all $i \leq k$, $\tilde{C}(x) = t_r$, and so $x \in Ch^r(\mathcal{O})$. Let (o_1, \dots, o_j) be the children of x . Since for all $i \leq k$ we have $o_i \in r(\pi^{-1}(T_i))$, the node o_{k+1} is compatible with (T_1, \dots, T_k) . Thus $T_{k+1} \in \mathbb{C}(G/\mathcal{O}_{k+1})$ since $T_{k+1} \in \mathbb{C}(G/\{o_{k+1}\})$, and $o_{k+1} \in \mathcal{O}_{k+1}$.

(\Leftarrow) Assume now that $T_{k+1} \in \mathbb{C}(G/\mathcal{O}_{k+1})$. There exists $o_{k+1} \in \mathcal{O}_{k+1}$ and a solution subtree C_{k+1} of $G/\{o_{k+1}\}$ rooted at o_{k+1} with $\pi(C_{k+1}) = T_{k+1}$. Since $o_{k+1} \in \mathcal{O}_{k+1}$, there exists an AND-child node x of a node o in \mathcal{O} such that $\tilde{C}(x) = t_r$, and $o_i \in r(\pi^{-1}(T_i))$ for all $i \leq k$ where o_i is the unique child of x of color c_i . Therefore, for all $i \leq k$, there exists a solution subtree C_i of $G/\{o_i\}$ such that $\pi(C_i) = T_i$. Now consider any solution subtree C of G/\mathcal{O} rooted at o , with o having x as AND-child node and with x having its first $k+1$ child solution subtrees equal to (C_1, \dots, C_{k+1}) . Then the first $k+1$ child subtrees of $\pi(C)$ will be $(T_1, \dots, T_k, T_{k+1})$. \square

Proposition 7. Let T be an equivalence class of G/\mathcal{O} for a set \mathcal{O} of OR^+ nodes of G , all colored with the same color c . Then, the function `Next` of Algorithm 1 is such that:

1. `Next`(\emptyset, \mathcal{O}) returns the smallest equivalence class of G/\mathcal{O} w.r.t. \prec .
2. `Next`(T, \mathcal{O}) returns the equivalence class of G/\mathcal{O} that follows T w.r.t. \prec .
3. if T is the last equivalence class of G/\mathcal{O} , `Next`(T, \mathcal{O}) returns \perp .

Proof. Let us define the height of G/\mathcal{O} , $h(G/\mathcal{O})$ to be the maximum height of an equivalence class of G/\mathcal{O} , i.e., the number of OR^+ nodes in a longest path from \mathcal{O} to a goal node minus 1. The proof will be done by induction on $h(G/\mathcal{O})$.

Assume first that $h(G/\mathcal{O}) = 0$, i.e., \mathcal{O} contains only goal nodes. Then G/\mathcal{O} has only one equivalence class T , which is the single node tree of color c . The call of `Next`(\emptyset, \mathcal{O}) will output it in Line 5 of the algorithm, and the call of `Next`(T, \mathcal{O}) will return \perp in Line 11 since $|\mathcal{T}(\mathcal{O})| = 0$.

Assume now that $h(G/\mathcal{O}) > 0$.

Proof of 7.1 If \mathcal{O} contains goal nodes, then the smallest equivalence class of G/\mathcal{O} with respect to \prec is the single node tree colored with c and $\text{Next}(\emptyset, \mathcal{O})$ outputs it in Line 5. Otherwise, let T be the smallest equivalence class of G/\mathcal{O} and let (T'_1, \dots, T'_j) be the subtrees of T rooted at the children of $r(T)$. By definition of \prec , the roots of (T'_1, \dots, T'_j) are colored with the minimum color tuple $t_1 := (c_1, \dots, c_j)$ of $\mathcal{T}(\mathcal{O})$ and by Lemma 6, for all $i \leq j$ T'_i is the smallest equivalence class of G/\mathcal{O}_i where $\mathcal{O}_i \subseteq C_i^1$ is the set of nodes of C_i^1 compatible with (T'_1, \dots, T'_{i-1}) . Thus, $\text{Next}(\emptyset, \mathcal{O})$ will return T in Line 27 since r will receive 1 in Line 9, T_1 will receive $\text{Next}(\emptyset, C_1^1)$ in Line 22 which is equal to T'_1 by the induction hypothesis, and for all $i \leq j$, T_i will receive $\text{Next}(\emptyset, \mathcal{O}_i)$ in Line 25 which is equal to T'_i by the induction hypothesis.

Proof of 7.2 Let (T'_1, \dots, T'_j) be the child subtrees of $r(T)$ and let $t_r := (c_1, \dots, c_j)$ be the color tuple of $\mathcal{T}(\mathcal{O})$ with which their roots are colored. Let T' be the equivalence class of G/\mathcal{O} that follows T with respect to \prec . Notice that the children of $r(T')$ are either colored with t_r or with t_{r+1} if T is the largest equivalence class whose root children are colored with t_r .

Assume first that the children of $r(T')$ are colored with the color tuple t_r . Let (T''_1, \dots, T''_j) be the child subtrees of $r(T')$ and let $\mathcal{O}_i \subseteq C_i^r$ be the set of nodes of C_i^r compatible with $(T''_1, \dots, T''_{i-1})$ for all $i \leq j$. Let ℓ be the smallest index such that $T'_\ell \neq T''_\ell$. We claim that ℓ is also the largest index such that T'_ℓ has a successor in $\mathbb{C}(G/\mathcal{O}_\ell)$ with respect to \prec , and thus that it corresponds to the ℓ chosen by the algorithm in Line 20. Indeed, assume that there exist $i < k \leq j$ and $F \in \mathbb{C}(G/\mathcal{O}_k)$ such that $T'_k \prec F$. By Lemma 6, there exists an equivalence class of G/\mathcal{O} whose first k child subtrees would be $(T'_1, \dots, T'_{k-1}, F)$. However, in this case such an equivalence class would be greater than T and smaller than T' with respect to \prec , and it would be in contradiction with the fact that T' immediately follows T in \prec . Now T_ℓ will receive $\text{Next}(T'_\ell, \mathcal{O}_\ell)$ in Line 22 which is by the induction hypothesis the tree of $\mathbb{C}(G/\mathcal{O}_\ell)$ that follows T'_ℓ . Since we assumed that ℓ is the smallest index such that $T'_\ell \neq T''_\ell$, by definition of \prec and by Lemma 6, T''_ℓ is the equivalence class of G/\mathcal{O}_ℓ that follows T'_ℓ in \prec , and so T_ℓ will receive T''_ℓ in Line 22. Since for all $i \leq \ell$ $T'_i = T''_i$, and since $(T_1, \dots, T_{\ell-1})$ are not modified by the algorithm, at the end of it, $(T_1, \dots, T_{\ell-1}, T_\ell)$ will be equal to $(T'_1, \dots, T'_{\ell-1}, T''_\ell) = (T''_1, \dots, T''_{\ell-1}, T''_\ell)$. It now remains to show that T_i will be equal to T''_i for all $\ell < i \leq j$. Again, by Lemma 6, for all $\ell < i \leq j$, T''_i is the smallest tree of $\mathbb{C}(G/\mathcal{O}'_i)$ where \mathcal{O}'_i is the set of nodes of C_i^r compatible with $(T''_1, \dots, T''_{i-1})$ since otherwise, another tree of $\mathbb{C}(G/\mathcal{O})$ greater than T and smaller than T' could be built. Thus, in Line 25, T_i will receive $\text{Next}(\emptyset, \mathcal{O}'_i)$ which is equal to T''_i by the induction hypothesis, and T' will be returned in Line 27.

Assume now that the children of $r(T')$ are colored with the color tuple t_{r+1} . In this case, T is the greatest tree of $\mathbb{C}(G/\mathcal{O})$ with respect to \prec whose root children are colored with t_r . Let $\mathcal{O}_i \subseteq C_i^r$ be the set of nodes of C_i^r compatible with (T'_1, \dots, T'_{i-1}) for all $i \leq j$. We claim that T'_i is the greatest tree of $\mathbb{C}(G/\mathcal{O}_i)$ for all $i \leq j$ with respect to \prec . Indeed, assume otherwise that there exists $T''_i \in \mathbb{C}(G/\mathcal{O}_i)$ with $T'_i \prec T''_i$. Then by Lemma 6 $\mathbb{C}(G/\mathcal{O})$ would contain a tree whose root children would be colored with t_r which would be larger than T with respect to \prec . By the induction hypotheses, $\text{Next}(T'_i, \mathcal{O}_i)$ will therefore return \perp for all $i \leq j$. So T will receive \emptyset in Line 20 and the algorithm will jump to Line 8, and the next color tuple t_{r+1} will be selected in Line 9. Using now similar arguments to the ones used in the proof of 7.1, the smallest tree with respect to \prec whose root children are colored with t_{r+1} will be returned.

Proof of 7.3 Assume now that T is the last equivalence class of G/\mathcal{O} . Notice that in this case, the root children of T are colored with t_r where $r = |\mathcal{T}(\mathcal{O})|$. As previously, T'_i is the greatest tree of $\mathbb{C}(G/\mathcal{O}_i)$ for all $i \leq j$ with respect to \prec since otherwise a greater tree of $\mathbb{C}(G/\mathcal{O})$ would exist. By the induction hypothesis, $\text{Next}(T'_i, \mathcal{O}_i)$ would return \perp for all $i \leq j$. Therefore

T will receive \emptyset in Line 20 and the algorithm will return to Line 8. Since $r = |\mathcal{T}(\mathcal{O})|$, r will receive $r + 1$ in Line 9 and \perp will be returned in Line 11. \square

Theorem 8. *Given an e-colored ad-AND/OR graph G , the set $\mathbb{C}(G)$ can be enumerated with delay $O(n \cdot s)$ where n is the number of nodes of G and s is the maximum size of a solution.*

Proof. To enumerate $\mathbb{C}(G)$, we first split the start nodes of G into sets S_0, \dots, S_k according to their colors. For each set S_i , starting with $T = \emptyset$, we repeatedly assign $\text{Next}(T, S_i)$ to T and output it until $T = \perp$. By Proposition 7, this guarantees that we output every solution of $\mathbb{C}(G/S_i)$ exactly once. Since any solution of $\mathbb{C}(G)$ belongs to $\mathbb{C}(G/S_i)$ for a given $i \leq k$, every solution of $\mathbb{C}(G)$ will be outputted exactly once.

For the complexity, notice that at most one recursive call is performed by the node of the next solution. More precisely, if $\text{Next}(T, \mathcal{O}) = T'$, there will be exactly one recursive call per node in T' that is not in T , and thus at most s recursive calls will be performed.

In each recursive call, both the set $\mathcal{T}(\mathcal{O})$ and the partition $\{C_i^r\}_{i \leq j}$ of grandchildren of \mathcal{O} can be computed in $O(n)$ time. It remains to show that the sets of compatible nodes \mathcal{O}_i , $i \leq j$, can be computed in $O(n)$ time in total which will conclude the proof. To do this, we should be able to compute the sets $r(\pi^{-1}(T_i))$ for all $i \leq j$. If $\text{Next}(T, \mathcal{O}) = T'$, the easiest way is to return the set $r(\pi^{-1}(T'))$ together with T' when the call $\text{Next}(T, \mathcal{O})$ returns. This could be done by observing that if $T \in \mathbb{C}(G/\mathcal{O})$ where \mathcal{O} is a set of goal nodes all having the same color c , then $r(\pi^{-1}(T)) = \mathcal{O}$, and if T has child subtrees T_1, \dots, T_j then $r(\pi^{-1}(T))$ is the set of nodes of \mathcal{O} that has at least an AND-child x such that the children of x contain exactly one node in $r(\pi^{-1}(T_i))$ for each $i \leq j$, which can be found in $O(n)$ time. Thus only $O(n)$ time is necessary at each recursive call to return $r(\pi^{-1}(T'))$ in addition to T' . \square

2.4 Restricting the graph to an equivalence class

After the enumeration of the equivalence classes, it might be interesting to go back to the solutions in each class, in particular, one might want to use the number of solutions as a measure for the “importance” or “significance” of an equivalence class. We present an algorithm that, given an e-colored ad-AND/OR graph G and an equivalence class T , constructs the *subgraph G^T of G restricted to T* , that is, a subgraph of G of which the solution subtrees are exactly the ones of G belonging to the equivalence class T : $\mathbb{T}(G^T) = \pi_G^{-1}(T) := \{T_1 \in \mathbb{T}(G) \mid \pi(T_1) = T\}$. Once the graph G^T is obtained, the following questions can be answered (by applying the same method as for the unrestricted ad-AND/OR graph G): counting the number of, and enumerating the solution subtrees belonging to the class T .

Algorithm 2 relies on two recursive functions `VisitOR` and `VisitAND`, both taking as input a node in G and a node in T . The **Require** statements are used to specify the preconditions that the two parameters of the two `Visit` functions must verify; it can be checked by inspection that these conditions are always satisfied whenever the functions are called. The algorithm performs an operation called `Mark` on the nodes in G . All nodes are initially unmarked; the `Mark` operation changes the state of a node into marked.

Recall that $\pi(T_1)$ transforms a solution subtree $T_1 \in \mathbb{T}(G)$ into an equivalence class $T \in \mathbb{C}(G)$ by contracting the AND nodes in T_1 . For a fixed T_1 , we extend this notation and write $\pi(s) = \pi(s_1) := v$ for every OR node s in T_1 with its unique AND-child node s_1 in T_1 that are identified with the node v in T under the transformation.

Lemma 9. *Let T_1 be a solution subtree of G belonging to the class T . For every node s of T_1 , there is a call to either `VisitOR` or to `VisitAND` of Algorithm 2 with parameters s and $\pi(s)$, depending on whether s is an OR^+ node or an AND node.*

Algorithm 2: Restricting the graph to an equivalence class

Data: an e-colored ad-AND/OR graph G , an equivalence class T

Result: the graph G^T

```
1 Function Main( $G, T$ ):
2   for each start node  $s_0$  of  $G$  such that  $c(s_0) = c(r(T))$  do
3     | VisitOR( $s_0, r(T)$ )
4   end
5   return  $G^T$  obtained from  $G$  by removing all unmarked nodes

6 Function VisitOR( $s, v$ ):
7   Require:  $c(s) = c(v)$ 
8   if  $s$  is a goal node and  $v$  is a leaf then
9     | Mark( $s$ )
10    | return
11  end
12  for each child node  $s_i$  of  $s$  in  $G$  such that  $\tilde{C}(s_i) = \tilde{C}(v)$  do
13    | VisitAND( $s_i, v$ )
14  end
15  if at least one child of  $s$  is marked then
16    | Mark( $s$ )
17  end

18 Function VisitAND( $s, v$ ):
19  Require:  $\tilde{C}(s) = \tilde{C}(v)$ 
20  for each child node  $s_i$  of  $s$  do
21    |  $v_i \leftarrow$  the unique child of  $v$  such that  $c(v_i) = c(s_i)$ 
22    | VisitOR( $s_i, v_i$ )
23  end
24  if all children of  $s$  are marked then
25    | Mark( $s$ )
26  end
```

Proof. By top-down induction. The start node of T_1 is visited in Line 3 since it has the correct color. In the induction step we separate two cases. For an OR⁺ node s of T_1 that is not a start node, suppose that the parent s_1 of s in T_1 is visited in a call **VisitAND**($s_1, \pi(s_1)$). Then s is visited (Line 22), and the second parameter is $\pi(s)$. In the other case, for an AND node s of T_1 , suppose that the parent s_1 of s in T_1 is visited in a call **VisitOR**($s_1, \pi(s_1)$). Since we have $\pi(s) = \pi(s_1)$ and $\tilde{C}(\pi(s_1)) = \tilde{C}(s)$, the condition at Line 12 is satisfied and s is visited in a call **VisitAND**($s, \pi(s)$). \square

The correctness of Algorithm 2 is shown in Proposition 10. We omit the analysis of complexity as the algorithm clearly requires a running time that is linear in the size of the graph.

Proposition 10. *The set of solution subtrees of G^T , the graph returned by Algorithm 2, is equal to $\pi_G^{-1}(T)$, i.e., the set of solution subtrees of G belonging to the equivalence class T .*

Proof. (First direction) We show that any solution subtree of G^T is also a solution subtree of G , and that it belongs to the equivalence class T . For every marked OR node, at least one child is marked (Line 15); for every marked AND node, all its children are marked (Line 24). A solution

subtree of G^T is thus a solution subtree of G . Let T_1 be a solution subtree of G^T , consider the recursion tree of the Visit function calls during which the nodes in T_1 are marked. By the preconditions of the Visit functions ($c(s) = c(v)$ and $\tilde{C}(s) = \tilde{C}(v)$), the tree $\pi(T_1)$ is equal to the tree formed by the colored nodes that are used as the second parameter v in this recursion tree. The latter is simply equal to T (we start with the root of T , then visit each child of the current node), so we have $\pi(T_1) = T$. Therefore, every solution subtree of G^T belongs to the class T .

(*Second direction*) Let T_1 be a solution subtree of G such that $\pi(T_1) = T$, we show that every node in T_1 is marked by bottom-up induction. By Lemma 9, any goal node s in T_1 is visited in a call `VisitOR`($s, \pi(s)$) so s is marked (Line 9) because $\pi(s)$ is necessarily a leaf. For the induction step we separate two cases. Let s be an OR node in T_1 , and suppose that all nodes in T_1 at a smaller height are marked. By the same lemma, s is visited. Then s is marked at Line 16, since exactly one child of s is in T_1 and is thus marked. In the other case, let s be an AND node in T_1 and suppose that all nodes in T_1 at a smaller height are marked. By the lemma, s is visited. Then s is marked at Line 25 because all children of s are in T_1 and are thus marked. This completes the proof. \square

3 Application to dynamic programming

3.1 A formalism for tree-sequential dynamic programming

Since its introduction by Karp and Held [21], *monotone sequential decision processes* (mSDP) have been the classical model for problems solvable by dynamic programming (DP). This formalism is based on finite-state automata. The solutions of DP-problems are thus equivalent to languages of regular expressions, or to paths in directed graphs. It is known that Bellman's *principle of optimality* [5] also applies to problems for which the solutions are not sequential but *tree-like* [9]. Various generalizations have been proposed to characterize broader classes of problems solvable by DP or DP-like techniques [11, 19]. In this paper, we consider a framework which is the immediate generalization of the mSDP model, i.e., generalizing finite automata (regular expressions, paths in DAGs) to finite tree automata (regular tree grammars, solution trees of general AND/OR graphs). Further generalizations exist (from trees to graphs of treewidth > 1); the collection of these methods is known as *Non-serial dynamic programming* [6].

In this model, a tree-sequential problem can be specified by a finite (bottom-up) tree automaton $A = (Q, \Sigma, \delta, q_0, Q_F)$, where Q is a finite set of states, Σ is a ranked alphabet, δ is a set of transition rules of the form (q_1, \dots, q_n, a, q) where $q_1, \dots, q_n, q \in Q$ and $a \in \Sigma$, $q_0 \in Q$ is the initial state, $Q_F \subseteq Q$ is a set of final states. The problem specification also includes a cost function. The set $L(A)$ of trees accepted by the tree automaton A defines the set of feasible solutions. The minimization problem seeks to minimize the cost function over the set $L(A)$ of feasible solutions.

We will consider the simple case of a positive additive cost function that always equals zero in the initial state. An additive cost function can be defined via an incremental cost function $I: Q^* \times \Sigma \rightarrow \mathbb{R}$, where Q^* consists of tuples of states in Q of the form (q_1, \dots, q_n) . $I(q_1, \dots, q_n, a)$ can be viewed as the cost of attaching n child subtrees to a new root of symbol a . While it might seem restrictive to require an additive structure on the cost function, this simple case does cover many important problems admitting a DP-algorithm, for instance, Travelling Salesman [4, 18], Knapsack [22], or Levenshtein distance [34].

In this case, the answer of the minimization problem can be shown to be equal to $\min_{q \in Q_F} D(q)$,

where $D : Q \rightarrow \mathbb{R}_{\geq 0}$ is defined by the following recurrence equations:

$$D(q_0) = 0, \\ \text{for } q \neq q_0, \quad D(q) = \min_{(q_1, \dots, q_n, a, q) \in \delta} \sum_{1 \leq i \leq n} D(q_i) + I(q_1, \dots, q_n, a). \quad (1)$$

A *dynamic programming algorithm* for the minimization problem corresponds to an algorithm that computes D ; the function D is commonly called a *dynamic programming table* (a DP-tabled, also called a DP-array, or a DP-matrix). Needless to say, such an algorithm does not exist in general for given arbitrary tree automata and cost functions [20].

Using an algebraic approach, Gnesi and Montanari [16] have shown that solving the functional Equation (1) corresponds to finding the solution subtrees of a general AND/OR graph. An important special case in which DP-algorithms exist is when the underlying AND/OR graph is acyclic.

When a fixed tree is given as an input to the problem, the underlying AND/OR graph is acyclic and decomposable (that is, it is an ad-AND/OR graph). Such problems are hence naturally solvable by DP-algorithms. These algorithms are known in folklore under the name *Dynamic programming on a tree*. Many graph-theoretical problems (e.g., maximum matching, longest path) can be solved optimally on trees by DP-algorithms. Numerous real-world applications also rely on DP-algorithms on trees; examples can be found, for instance, in Data Science [30], Computer Vision [14, 33], and Computational Biology [3, 13].

Explicit construction of the ad-AND/OR graph for DP on a fixed tree Due to its usefulness for the examples that we will develop next, in the case of DP on a fixed tree, an explicit construction of the ad-AND/OR graph from Equation (1) is described below. The construction is done in two steps. In the first step, we build a graph in which every node retains an additional attribute, its *value*, and every OR^+ node is labeled by a state $q \in Q$. In the second step, we *prune* the graph by removing nodes that do not yield optimal values.

1. For each $(q_0, a, q) \in \delta$, create a goal node of value 0 labeled by q . Then, for each $q \neq q_0$ in post-order,
 - i. For each $(q_1, \dots, q_n, a, q) \in \delta$, create an AND node, connect it to the n OR^+ nodes labeled by q_1, \dots, q_n . Its value is equal to the sum of the values of its children, plus $I(q_1, \dots, q_n, a)$.
 - ii. Create a single OR node, connect it to every AND node created in the previous step. Its label is q , and its value is the minimum of the values of its children.
2. For each $q \in Q_F$, remove the OR node labeled by q unless its value is equal to $\min_{q \in Q_F} G(q)$. For each OR node s , remove the arc to its AND-child node s_i if the value of s_i is not equal to the value of s . Finally, remove recursively all AND nodes without incoming arcs.

3.2 Examples

3.2.1 Optimal tree coloring problem

Description A prototypical problem that fits into the framework of DP on a tree is OPTIMAL TREE COLORING, that is, finding an optimal node-coloring of the input tree. Many problems of practical interest reduce to OPTIMAL TREE COLORING; three concrete examples are given later in this section.

If T is the input (rooted, ordered) tree and C is the set of colors, such a problem seeks a coloring $\phi: V(T) \rightarrow C$ that minimizes the cost function. There can be many constraints on the coloring function: some nodes of T may be forced to have a certain color, the possible colors of a node may depend on the colors of its descendants. In our tree-sequential dynamic programming formalism, a tree automaton and a cost function are given as part of the input. The tree automaton defines the set $L(A)$ of feasible coloring functions satisfying all those constraints. A state q can be interpreted as a colored subtree of T with a particular root color; the unique initial state is an empty coloring and transitions into a colored leaf of T ; a final state corresponds to a fully colored T with a particular root color. A commonly used form of cost functions considers the (possibly weighted) sum over the edges of the tree of the cost of putting two colors on each end of an edge, that is, an incremental cost function I of the form $I(q_1, \dots, q_n, a, p) = \sum_{1 \leq i \leq n} p(a_i, a)$ where a_i is the color of the root of the subtree in state q_i and $p: C^2 \rightarrow \mathbb{R}_{\geq 0}$ is a function that gives the cost of putting two colors at each end of an edge.

Equivalence relations on the set of solutions A possible strategy to define equivalence classes on the solution space of the OPTIMAL TREE COLORING problem is to consider some colors to be locally equivalent on a node. In practical applications (see the next section), the space of colors can be quite large. Even though the precise colors of each node are necessary for correctly computing the cost function, when the solutions are analyzed by a human expert, it can be desirable to omit the colors and just look at whether the color of a node belongs to some group of colors. Therefore, this kind of equivalence relations is natural in many situations. Our Definition 4 of equivalence classes of an e-colored ad-AND/OR graph deals exactly with equivalence relations of this type.

Let e be a function that maps a color c to its “color group” $e(c)$. Two solutions of the OPTIMAL TREE COLORING problem $\phi_1, \phi_2: V(T) \rightarrow C$ are said to be *equivalent* if $\forall u \in V(T)$, $e(\phi_1(u)) = e(\phi_2(u))$. Let G be the ad-AND/OR graph associated with this instance. For each OR^+ node s of G labeled with the state q , where q is interpreted as a colored subtree of T with a particular root color c , color the node s with $e(c)$. Then G is e-colored and $\mathbb{C}(G)$ corresponds to the set of equivalence classes of the solutions of the instance. Notice that the constraint we had on the e-coloring of an ad-AND/OR graph is naturally satisfied by any meaningful function e because in a DP setting we only consider ordered trees: the i -th and j -th children of a node of T cannot be in the same color group unless $i = j$.

Concrete examples of tree coloring problems

Example 1 The first example is related to the alignment of gene sequences on a phylogenetic tree [31]. The input is a tree T , a set Σ of letters (DNA alphabet or protein alphabet), a function that labels each leaf node of T with a letter, and a distance function $d: \Sigma^2 \rightarrow \mathbb{R}_{\geq 0}$ between two letters. The goal is to extend the leaf labeling to a full labeling $\phi: V(T) \rightarrow \Sigma$ such that the sum of the distances over the edges of T is minimized. Defining equivalence relations of the solutions based on a grouping of the letters is uncontrived in this problem: for the DNA or protein alphabet, the letters can be subdivided into structurally similar groups.

Example 2 The FREQUENCY ASSIGNMENT problems are a family of problems that naturally arise in telecommunication networks, and that have been extensively studied in graph theory as a generalization of graph coloring known as the T-coloring problem [17, 29, 32]. In the variant called the list T-coloring, the input is a graph G representing the interference between radio stations, a set C of colors, a function S that gives for each vertex $v \in V(G)$ a set $S(v) \subseteq C$

of colors (possible frequencies for a station), and a set $T \subseteq C^2$ of forbidden pairs of colors (interfering frequencies). The goal is to find a coloring $\phi: V(G) \rightarrow C$ such that $\forall v \in V(G)$, $\phi(v) \in S(v)$, and $\forall (u, v) \in E(G)$, $(\phi(u), \phi(v)) \notin T$. While this problem is hard in general, it can be solved by DP when the underlying graph is a tree. In this case, we can enumerate colorings of the input trees without any forbidden pair of colors on the edges. Defining equivalence relations by grouping some of the colors together (similar frequencies) can be a practical way of reducing the size of the output.

Example 3 The TREE RECONCILIATION problem is the main method for analyzing the co-evolution of two sets of species, the hosts and their parasites [28]. The input are two phylogenetic trees H, P (of the hosts and of the parasites, respectively), together with a mapping $\phi_0: \text{Leaves}(P) \rightarrow \text{Leaves}(H)$ that reflects the present-day parasite infections. What needs to be enumerated are then all past associations, that is, all mappings of the non-leaf nodes of the parasite tree to the nodes of the host tree that optimize a function which overall represents the sum of the number of different possible “events” weighted by the inverse of their estimated probability. The number of optimal solutions is often huge and, by applying our Algorithm 1, the enumeration of biologically inspired equivalence classes have allowed a significant reduction (in some cases from 10^{42} to only 96 classes) of the size of the output while still preserving the important biological information (see [35]).

3.2.2 Dynamic programming on tree decomposition of a graph

Many graph problems can be solved in polynomial time with a dynamic programming algorithm when the input graph has bounded treewidth (see for example [8]). The underlying idea is that, given a tree decomposition of a graph, the dynamic programming algorithm traverses the nodes (bags) of the decomposition and consecutively solves the respective sub-problems. For vertex subset optimization problems, given a bag X , a dynamic programming algorithm generally computes for each $Z \subseteq X$ the optimal solution of the sub-problem whose intersection with X is Z . In this context, we could define two solutions to be equivalent if they intersect each bag of the decomposition in an “equivalent” way. The equivalence relation on the solutions is then defined by an equivalence relation over the subsets of each bag, and two solutions S_1 and S_2 are equivalent if for all bags X of the decomposition, $S_1 \cap X$ is equivalent to $S_2 \cap X$.

One of the simplest examples is to consider that all the nonempty subsets of vertices of a bag are equivalent. Thus, what we are interested in is whether a solution “hits” a bag (i.e., whether it has a nonempty intersection with the vertices in the bag). Consequently, two solutions would be considered equivalent if they hit the same bags.

We can also consider two subsets of a bag to be equivalent if they have the same size. In this case, two solutions would be equivalent if each bag contains the same number of vertices in the two solutions.

We believe that considering solutions in the way they are distributed along the tree decomposition of a graph could give a good overview of the diversity of the solution space.

4 Conclusion and perspectives

In this paper, we provide a general framework for the enumeration of equivalence classes of solutions in polynomial delay for a wide variety of contexts. This work opens a door to different research directions.

It would be interesting to ask whether we can efficiently enumerate groups of solutions that result from classical clustering procedures, or one representative per group. Moreover, in this

paper we heavily rely on the decomposability property of the structure of the solution space. It remains open whether the problem of enumerating equivalence classes is hard without this restriction.

References

- [1] Steen A. Andersson, David Madigan, and Michael D. Perlman. A characterization of markov equivalence classes for acyclic digraphs. *Annals of Statistics*, 25(2):505–541, 04 1997. doi:10.7916/D8280JSB.
- [2] Albert Angel and Nick Koudas. Efficient diversity-aware search. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data*, SIGMOD '11, page 781–792, New York, NY, USA, 2011. Association for Computing Machinery. doi:10.1145/1989323.1989405.
- [3] Mukul S. Bansal, Eric J. Alm, and Manolis Kellis. Efficient algorithms for the reconciliation problem with gene duplication, horizontal transfer and loss. *Bioinformatics*, 28(12):i283–i291, 2012. doi:10.1093/bioinformatics/bts225.
- [4] Richard Bellman. Dynamic programming treatment of the travelling salesman problem. *Journal of the ACM*, 9(1):61–63, January 1962. doi:10.1145/321105.321111.
- [5] Richard Bellman. *Dynamic Programming*. Dover Books on Computer Science. Dover Publications, 2013.
- [6] Umberto Bertele and Francesco Brioschi. *Nonserial Dynamic Programming*. Academic Press, Inc., USA, 1972.
- [7] Anselm Blumer, Janet A. Blumer, David H. Haussler, Ross M. McConnell, and Andrzej Ehrenfeucht. Complete inverted files for efficient text retrieval and analysis. *Journal of the ACM*, 34(3):578–595, July 1987. doi:10.1145/28869.28873.
- [8] Hans L. Bodlaender. Dynamic programming on graphs with bounded treewidth. In Timo Lepistö and Arto Salomaa, editors, *Automata, Languages and Programming*, pages 105–118, Berlin, Heidelberg, 1988. Springer Berlin Heidelberg. doi:10.1007/3-540-19488-6_110.
- [9] Pierre E. Bonzon. Necessary and sufficient conditions for dynamic programming of combinatorial type. *Journal of the ACM*, 17:675–682, 1970. doi:10.1145/321607.321616.
- [10] Marília D.V. Braga, Marie-France Sagot, Celine Scornavacca, and Eric Tannier. Exploring the solution space of sorting by reversals, with experiments and an application to evolution. *IEEE/ACM transactions on computational biology and bioinformatics*, 5 3:348–56, 2008. doi:10.1109/TCBB.2008.16.
- [11] Joshua Buresh-Oppenheimer, Sashka Davis, and Russell Impagliazzo. A stronger model of dynamic programming algorithms. *Algorithmica*, 60:938–968, 08 2011. doi:10.1007/s00453-009-9385-1.
- [12] Rina Dechter and Robert Mateescu. And/or search spaces for graphical models. *Artificial Intelligence*, 171(2):73 – 106, 2007. doi:10.1016/j.artint.2006.11.003.
- [13] Beatrice Donati, Christian Baudet, Blerina Sinimeri, Pierluigi Crescenzi, and Marie-France Sagot. EUCALYPT: efficient tree reconciliation enumerator. *Algorithms for Molecular Biology*, 10(1):3, 2015. doi:10.1186/s13015-014-0031-3.

- [14] Pedro F. Felzenszwalb and Daniel P. Huttenlocher. Pictorial structures for object recognition. *International Journal of Computer Vision*, 61(1):55–79, January 2005. doi:10.1023/B:VISI.0000042934.15159.49.
- [15] Henning Fernau, Petr A. Golovach, and Marie-France Sagot. Algorithmic Enumeration: Output-sensitive, Input-Sensitive, Parameterized, Approximative (Dagstuhl Seminar 18421). *Dagstuhl Reports*, 8(10):63–86, 2019. doi:10.4230/DagRep.8.10.63.
- [16] Stefania Gnesi, Ugo Montanari, and Alberto Martelli. Dynamic programming as graph searching: An algebraic approach. *Journal of the ACM*, 28(4):737–751, October 1981. doi:10.1145/322276.322285.
- [17] William K. Hale. Frequency assignment: Theory and applications. *Proceedings of the IEEE*, 68(12):1497–1514, 1980. doi:10.1109/PROC.1980.11899.
- [18] Michael Held and Richard M. Karp. A dynamic programming approach to sequencing problems. *Journal of the Society for Industrial and Applied Mathematics*, 10(1):196–210, 1962. doi:10.1137/0110015.
- [19] Paul Helman. A common schema for dynamic programming and branch and bound algorithms. *Journal of the ACM*, 36(1):97–128, January 1989. doi:10.1145/58562.59304.
- [20] Toshihide Ibaraki. Classes of discrete optimization problems and their decision problems. *Journal of Computer and System Sciences*, 8(1):84–116, 1974. doi:10.1016/S0022-0000(74)80024-3.
- [21] Richard M. Karp and Michael Held. Finite-state processes and dynamic programming. *SIAM Journal on Applied Mathematics*, 15(3):693–718, 1967. doi:10.1137/0115060.
- [22] Hans Kellerer, Ulrich Pferschy, and David Pisinger. *Basic Algorithmic Concepts*, pages 15–42. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004. doi:10.1007/978-3-540-24777-7_2.
- [23] Alberto Martelli and Ugo Montanari. Optimizing decision trees through heuristically guided search. *Communications of the ACM*, 21(12):1025–1039, December 1978. doi:10.1145/359657.359664.
- [24] Cristian Molinaro, Amy Sliva, and Vs S. Subrahmanian. Super-solutions: Succinctly representing solutions in abductive annotated probabilistic temporal logic. *ACM Transactions on Computational Logic*, 15(3), July 2014. doi:10.1145/2627354.
- [25] Katherine Morrison. An enumeration of the equivalence classes of self-dual matrix codes. *Advances in Mathematics of Communications*, 9:415, 2015. doi:10.3934/amc.2015.9.415.
- [26] Kazuyuki Narisawa, Shunsuke Inenaga, Hideo Bannai, and Masayuki Takeda. Efficient computation of substring equivalence classes with suffix arrays. In Bin Ma and Kaizhong Zhang, editors, *Combinatorial Pattern Matching*, pages 340–351, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg. doi:10.1007/s00453-016-0178-z.
- [27] Nils J. Nilsson. *Principles of Artificial Intelligence*. Springer-Verlag Berlin Heidelberg, Tioga, Palo Alto, CA, 1982.
- [28] Roderic D. M. Page. *Tangled trees: phylogeny, cospeciation, and coevolution*. The University of Chicago Press, 2003.

- [29] Fred S. Roberts. T-colorings of graphs: recent results and open problems. *Discrete Mathematics*, 93(2):229–245, 1991. doi:10.1016/0012-365X(91)90258-4.
- [30] Lior Rokach and Oded Z. Maimon. *Data Mining with Decision Trees: Theory and Applications*. Series in machine perception and artificial intelligence. World Scientific, 2008. doi:10.1142/9097.
- [31] David Sankoff. Minimal mutation trees of sequences. *SIAM Journal on Applied Mathematics*, 28(1):35–42, 1975. doi:10.1137/0128004.
- [32] Barry A. Tesman. List t-colorings of graphs. *Discrete Applied Mathematics*, 45(3):277–289, 1993. doi:10.1016/0166-218X(93)90015-G.
- [33] Olga Veksler. Stereo correspondence by dynamic programming on a tree. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 2, pages 384–390, 2005. doi:10.1109/CVPR.2005.334.
- [34] Robert A. Wagner and Michael J. Fischer. The string-to-string correction problem. *Journal of the ACM*, 21(1):168–173, January 1974. doi:10.1145/321796.321811.
- [35] Yishu Wang, Arnaud Mary, Marie-France Sagot, and Blerina Sinimeri. Copybara: equivalence class enumeration of cophylogeny event-based reconciliations. *Bioinformatics*, 36(14):4197–4199, 2020. doi:10.1093/bioinformatics/btaa498.