

**INSTITUT FÜR INFORMATIONSSYSTEME**  
ABTEILUNG DATENBANKEN UND ARTIFICIAL INTELLIGENCE

**D-FLAT<sup>2</sup>: Subset Minimization in  
Dynamic Programming on Tree  
Decompositions Made Easy**

**DBAI-TR-2015-93**

**Bernhard Bliem  
Markus Hecher**

**Günther Charwat  
Stefan Woltran**

Institut für Informationssysteme  
Abteilung Datenbanken und  
Artificial Intelligence  
Technische Universität Wien  
Favoritenstr. 9  
A-1040 Vienna, Austria  
Tel: +43-1-58801-18403  
Fax: +43-1-58801-18493  
sekret@dbai.tuwien.ac.at  
www.dbai.tuwien.ac.at

DBAI TECHNICAL REPORT  
2015



TECHNISCHE  
UNIVERSITÄT  
WIEN  
Vienna University of Technology

## **D-FLAT<sup>2</sup>: Subset Minimization in Dynamic Programming on Tree Decompositions Made Easy**

**Bernhard Bliem**<sup>1</sup>      **Günther Charwat**<sup>1</sup>      **Markus Hecher**<sup>1</sup>  
**Stefan Woltran**<sup>1</sup>

**Abstract.** Many problems from the area of AI have been shown tractable for bounded treewidth. In order to put such results into practice, quite involved dynamic programming (DP) algorithms on tree decompositions have to be designed and implemented. These algorithms typically show recurring patterns that call for tasks like subset minimization. In this paper we present a novel approach to obtain such DP algorithms from simpler principles, where the DP formalization of subset minimization is performed automatically. We first give a theoretical account of our novel method, and then present D-FLAT<sup>2</sup>, a system that allows one to specify the core DP algorithm via answer set programming (ASP). We illustrate the approach at work by providing several DP algorithms that are more space-efficient than existing solutions, while featuring improved readability, reuse and therefore maintainability of ASP code. Experiments show that our approach also yields a significant improvement in runtime performance.

---

<sup>1</sup>TU Wien. E-mail: {bliem,gcharwat,hecher,woltran}@dbai.tuwien.ac.at

**Acknowledgements:** This work has been supported by the Austrian Science Fund (FWF): Y698, P25607, P25518.

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>Dynamic Programming on Tree Decompositions</b>	<b>4</b>
<b>3</b>	<b>A Formal Account of DP on TDs</b>	<b>7</b>
3.1	Computations . . . . .	7
3.2	Table Compression . . . . .	10
3.3	Normalizing Computations . . . . .	11
<b>4</b>	<b>Practical Realization</b>	<b>15</b>
4.1	D-FLAT: A Quick Tutorial . . . . .	16
4.2	Towards D-FLAT <sup>2</sup> . . . . .	21
<b>5</b>	<b>Application to Common AI Problems</b>	<b>23</b>
5.1	SAT . . . . .	23
5.2	Circumscription . . . . .	23
5.3	Disjunctive ASP . . . . .	24
5.4	Abstract Argumentation . . . . .	25
<b>6</b>	<b>Evaluation</b>	<b>27</b>
6.1	Abstract Argumentation . . . . .	27
6.2	Subset-minimal SAT . . . . .	27
<b>7</b>	<b>Conclusion</b>	<b>29</b>
<b>A</b>	<b>Appendix: Proofs</b>	<b>33</b>

# 1 Introduction

Many prominent NP-hard problems in the area of AI have been shown tractable for bounded treewidth. Thanks to Courcelle’s theorem [9], it is sufficient to encode a problem as an MSO sentence in order to obtain such a result. To put this into practice, tailored systems for MSO logic are required, however. While there has been remarkable progress in this direction [26] there is still evidence that designing DP algorithms for the considered problems from scratch results in more efficient software solutions (cf. [11, 30]).

The actual design of these algorithms can be quite tedious, especially for problems located at the second level of the polynomial hierarchy like the AI problems circumscription, abduction, answer set programming or abstract argumentation (see [16, 21, 23, 24]). In many cases, the increased complexity of such problems is caused by co-NP hard subset minimization or maximization subproblems (e.g., minimality of models in circumscription). It is exactly the handling of these subproblems that makes the design of the DP algorithms difficult.

What we aim for in this article is thus the automatic generation of intricate DP algorithms from simpler principles. To the best of our knowledge, there is only a little amount of work in this direction. The D-FLAT system [2] – a declarative framework for rapid prototyping of DP algorithms on tree decompositions – offers a few built-ins for cost minimization and the handling of join nodes in standard DP algorithms; the LISP-based Autograph approach (see, e.g., [10]) on the other hand makes it possible to obtain a specification of the problem at hand via combinations of (pre-defined) fly-automata.

What we have in mind is different and motivated by recent developments in the world of answer set programming (ASP) [8]: For exploiting the full expressive power of ASP, a saturation programming technique (see, e.g., [27]) is often required for the encoding of co-NP subproblems. Several approaches for relieving the user from this task have been proposed [7, 18, 19] that employ metaprogramming techniques. For instance, in order to compute minimal models of a propositional formula, one can simply express the SAT problem in ASP together with a special minimize statement (recognized by systems like *metasp*). In this way, one obtains a program computing subset minimal models. Unfortunately, easy-to-use facilities like such minimize statements had no analog in the area of DP so far.

In this article, we propose a solution to this issue: We provide a new method for automatically obtaining DP algorithms for problems requiring minimization, given only an algorithm for a problem variant without minimization. For example, given a DP algorithm for SAT [33], our approach enables us to generate a new algorithm for finding only subset-minimal models. Making minimization implicit in this way makes the programmer’s life considerably easier. Moreover, we present D-FLAT<sup>2</sup>, which is an implementation of the new method and extends D-FLAT<sup>1</sup>. To underline the practical relevance of this work, we show elegant solutions for common AI problems using D-FLAT<sup>2</sup>.

The contributions of this article are the following:

- We introduce a formal model of DP computations, abstracting from concrete algorithms.

---

<sup>1</sup>The systems D-FLAT and D-FLAT<sup>2</sup> are publicly available at <http://dbai.tuwien.ac.at/proj/dflat/system/> and <https://github.com/hmarkus/dflat-2/>, respectively.

Our results are therefore generally applicable, not just to a particular problem.

- We then show how our model captures typical DP computations for subset minimization problems.
- We discuss how computations can be compressed to ensure fixed-parameter tractability.
- Our main contribution is a formal definition of a transformation that turns non-minimizing computations into ones that perform minimization. We identify under which conditions this procedure is sound and give a formal proof.
- Finally, we discuss implementation issues. In contrast to DP algorithms that implement subset minimization directly, we present an implementation of our method, called D-FLAT<sup>2</sup>, which not only yields more succinct algorithm specifications, but also avoids redundant computations. We illustrate our method at work by providing DP encodings for several AI problems, and show promising preliminary performance results.

This article is structured as follows. In Section 2, DP on tree decompositions is introduced. Section 3 contains our formal model of DP computations. In Section 4 we illustrate how the concepts are put into practice in the D-FLAT and D-FLAT<sup>2</sup> systems. Next, in Section 5 we show how common AI problems such as Circumscription and disjunctive ASP can be solved using our approach. Finally, Section 6 reports on experimental results.

Compared to the workshop paper [4], this article contains a full formal proof of the correctness of the method implemented in D-FLAT<sup>2</sup>.

## 2 Dynamic Programming on Tree Decompositions

In this section we outline DP on tree decompositions. The ideas underlying this concept stem from the field of parameterized complexity. Many computationally hard problems become tractable in case a certain problem parameter is bound to a fixed constant. This property is referred to as *fixed-parameter tractability* [14], and the complexity class FPT consists of problems that are solvable in  $f(k) \cdot n^{\mathcal{O}(1)}$ , where  $f$  is a function that only depends on the parameter  $k$ , and  $n$  is the input size.

For problems whose input can be represented as a graph, an important parameter is *treewidth*, which measures “tree-likeness” of a graph. It is defined by means of tree decompositions (TDs) [32].

**Definition 1.** A tree decomposition of a graph  $G = (V, E)$  is a pair  $\mathcal{T} = (T, \chi)$  where  $T = (N, F)$  is a (rooted) tree and  $\chi : N \rightarrow 2^V$  assigns to each node a set of vertices (called the node’s bag), such that the following conditions are met:

- (1) For every  $v \in V$ , there exists a node  $n \in N$  such that  $v \in \chi(n)$ .
- (2) For every edge  $e \in E$ , there exists a node  $n \in N$  such that  $e \subseteq \chi(n)$ .
- (3) For every  $v \in V$ , the subtree of  $T$  induced by  $\{n \in N \mid v \in \chi(n)\}$  is connected.

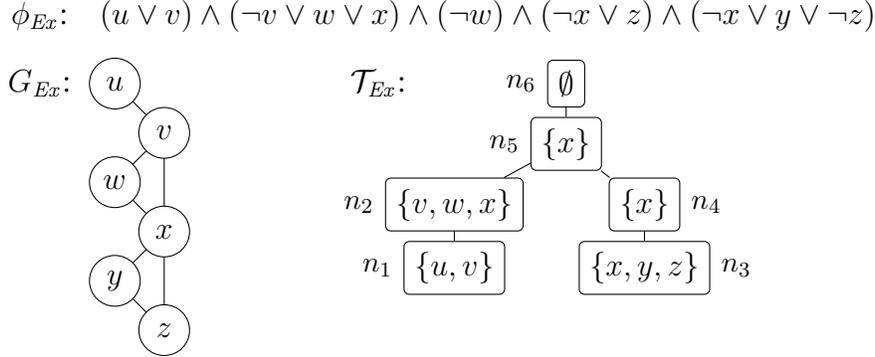


Figure 1: Primal graph  $G_{Ex}$  and a TD  $\mathcal{T}_{Ex}$  of  $\phi_{Ex}$ .

The width of  $\mathcal{T}$  is  $\max_{n \in N} |\chi(n)| - 1$ . The treewidth of a graph is the minimum width over all its tree decompositions.

Although constructing a minimum-width TD is intractable in general [3], it is in FPT [5] and there are polynomial-time heuristics giving “good” TDs [6, 12, 13].

**Example 1.** Let us consider the enumeration variant of the SAT problem. Given a propositional formula  $\phi$  in CNF, we first have to find an appropriate graph representation. Here, we construct the primal graph  $G$  of  $\phi$ , that is, vertices in  $G$  represent atoms of  $\phi$ , and atoms occurring together in a clause form a clique in  $G$ . An example formula  $\phi_{Ex}$ , its graph representation  $G_{Ex}$  and a possible TD  $\mathcal{T}_{Ex}$  are given in Figure 1. The width of  $\mathcal{T}_{Ex}$  is 2.

Algorithms for DP on TDs generally traverse the TD in post-order. At each node, partial solutions for the subgraph induced by the vertices encountered so far are computed and stored in a data structure associated with the node. The size of the data structure is typically bounded by the TD’s width and the number of TD nodes is linear in the input size. So if the width is bounded by a constant, the search space for subproblems is constant as well, and the number of subproblems only grows linearly for larger instances. We now illustrate the DP algorithm for SAT [33] on our running example; formal details are given in the next section.

**Example 2.** Figure 2 illustrates the DP computation for SAT. The tables are computed as follows. For a TD node  $n$ , each table row  $r$  consists of data  $D(r)$ , which stores partial truth assignments over atoms in  $\chi(n)$ . Here,  $D(r)$  only contains atoms that get assigned “true”, atoms in  $\chi(n) \setminus D(r)$  get assigned “false”. In  $r$ , all clauses covered by  $\chi(n)$  must be satisfied by the partial truth assignment. The set  $P(r)$  contains so-called extension pointer tuples (EPTs) that denote the rows in the children where  $r$  was constructed from. First consider node  $n_1$ : here,  $\chi(n_1) = \{u, v\}$  covers clause  $(u \vee v)$ , yielding three partial assignments for  $\phi_{Ex}$ . In  $n_2$ , the child rows are extended, the partial assignments are updated (by removing atoms not contained in  $\chi(n_2)$  and guessing truth assignments for atoms in  $\chi(n_2) \setminus \chi(n_1)$ ). Here, clauses  $(\neg v \vee w \vee x)$  and  $(\neg w)$  have to be satisfied. Observe that row  $2^1$  is constructed from two different child rows. In  $n_3$  we proceed as described before. In  $n_4$ , data related to removed vertices  $y$  and  $z$  are projected away. In  $n_5$ , additionally

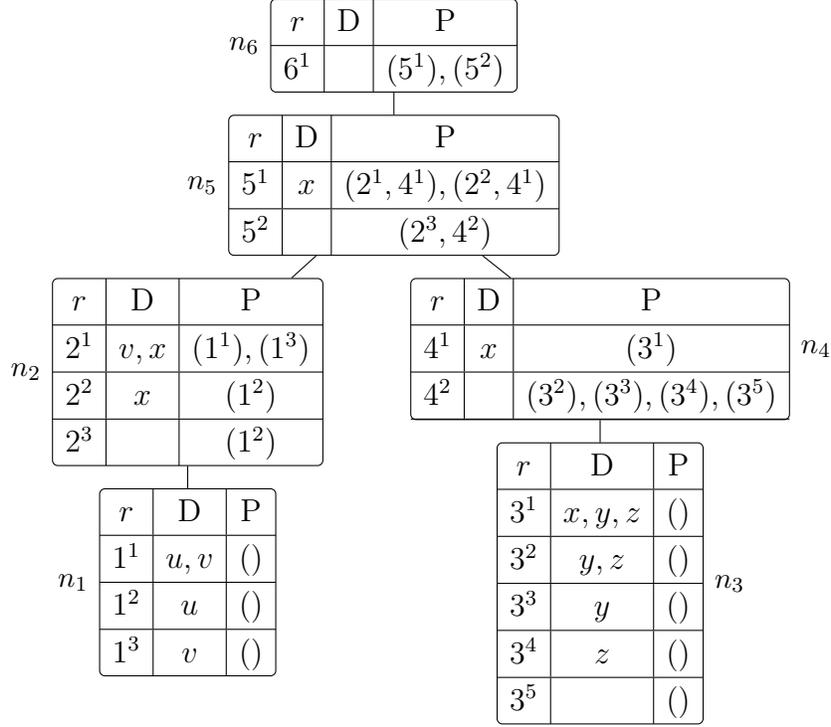


Figure 2: DP computation for SAT.

only partial assignments that agree on the truth assignment for common atoms are to be joined. We continue this procedure recursively until we reach the TD's root.

To decide whether a formula is satisfiable, it suffices to check if the table in the root node is non-empty. The overall procedure is in FPT time because the number of nodes in the TD is bounded by the input size (i.e., the number of atoms), and each node  $n$  is associated with a table of size at most  $\mathcal{O}(2^{|\chi(n)|})$  (i.e., the number of possible truth assignments). For our example,  $\phi_{Ex}$  is satisfiable due to existence of row  $5^1$ . Solutions (models of  $\phi_{Ex}$ ) can be enumerated with linear delay by starting at the root and following the EPTs while combining the partial assignments associated with the rows. For instance, model  $\{u, v, x, y, z\}$  is constructed by starting at  $6^1$  and following EPTs  $(5^1)$ ,  $(2^1, 4^1)$ ,  $(1^1)$  and  $(3^1)$ , thereby combining  $D(6^1) \cup D(5^1) \cup D(2^1) \cup D(1^1) \cup D(4^1) \cup D(3^1)$ .

The problems we will focus on in this paper are more involved. One such problem is the enumeration of the *subset-minimal* models of CNF formulas. Note that a naive implementation of this problem where the enumeration of models is used to filter out the minimal ones would violate the desired linear delay for enumerating the solutions. Indeed, genuine DP algorithms for such problem which involve subset minimization are needed. In the forthcoming section, we will present such an DP algorithm for  $\subseteq$ -MINIMAL SAT. Together with the above algorithm for SAT these algorithms will be used as a running example to illustrate our method.

### 3 A Formal Account of DP on TDs

In this section we formally introduce DP on TDs for subset optimization problems. First we define our data structure, called *computation*, which is a tree of tables resulting from the bottom-up traversal of the TD. We then define the *extensions* of rows in the tables, that are used to obtain (partial) solution candidates for the problem at hand. Additionally, we define their relation to so-called counterexamples. Counterexamples are witnesses for a solution candidate not being subset-minimal. Finally, we state how solutions are obtained, and define properties that have to be fulfilled by the tables in a computation in order to yield correct results. Section 3.2 then deals with compressing the tables, such that the size of the tables is bounded by the width of the TD. Finally, Section 3.3 introduces requirements for a computation without minimization to be eligible for transformation into a “minimizing” one (we call such a computation *augmentable*). We then show how this transformation can be achieved automatically, and prove that the solutions of the augmentable computation correspond to subset-minimal solutions of the original computation.

#### 3.1 Computations

First, let us define the data structure used in our approach.

**Definition 2.** A computation is a rooted ordered tree whose nodes are called tables. Each table  $R$  is a set of rows and each row  $r \in R$  possesses

- some problem-specific data  $D(r)$ ,
- a non-empty set of extension pointer tuples (EPTs)  $P(r)$  such that each tuple is of arity  $k$ , where  $k$  is the number of children of  $R$ , and for each  $(p_1, \dots, p_k) \in P(r)$  it holds that each  $p_i$  is a row of the  $i$ -th child of  $R$ ,
- a subtable  $S(r)$ , which is a set of subrows, where each subrow  $s \in S(r)$  possesses
  - some problem-specific data  $D(s)$ ,
  - a non-empty set of EPTs  $P(s)$  such that for each  $(p_1, \dots, p_k) \in P(s)$  there is some  $(q_1, \dots, q_k) \in P(r)$  with  $p_i \in S(q_i)$  for  $1 \leq i \leq k$ ,
  - an inclusion status flag  $\text{inc}(s) \in \{\text{eq}, \subset\}$ .

For rows or subrows  $a, b$  we write  $a \approx b$ ,  $a \leq b$  and  $a < b$  to denote  $D(a) = D(b)$ ,  $D(a) \subseteq D(b)$  and  $D(a) \subset D(b)$ , respectively. For sets of rows or subrows  $R, S$  we write  $D(R)$  to denote  $\bigcup_{r \in R} D(r)$ , and we write  $R \approx S$ ,  $R \leq S$  and  $R < S$  to denote  $D(R) = D(S)$ ,  $D(R) \subseteq D(S)$  and  $D(R) \subset D(S)$ , respectively.

The reason that each row possesses a subtable is that we consider subset optimization problems, and we assume that algorithms for such problems use subtables to store potential counterexamples to a solution candidate being subset-minimal. The intuition of a subrow  $s$  for a row  $r$  is that  $s$  represents solution candidates that are subsets of the candidates represented by  $r$ . If one of these subset relations is proper, we indicate this by  $\text{inc}(s) = \subset$ .

$n_3$						
$R$			$S(r)$			
$r$	D	P	$s$	D	P	inc
$3^1$	$x, y, z$	$()$	$3_1^1$	$x, y, z$	$()$	eq
			$3_2^1$	$y, z$	$()$	$\subset$
			$3_3^1$	$y$	$()$	$\subset$
			$3_4^1$	$z$	$()$	$\subset$
			$3_5^1$		$()$	$\subset$
$3^2$	$y, z$	$()$	$3_1^2$	$y, z$	$()$	eq
			$3_2^2$	$y$	$()$	$\subset$
			$3_3^2$	$z$	$()$	$\subset$
			$3_4^2$		$()$	$\subset$
$3^3$	$y$	$()$	$3_1^3$	$y$	$()$	eq
			$3_2^3$		$()$	$\subset$
$3^4$	$z$	$()$	$3_1^4$	$z$	$()$	eq
			$3_2^4$		$()$	$\subset$
$3^5$		$()$	$3_1^5$		$()$	eq

$n_4$						
$R$			$S(r)$			
$r$	D	P	$s$	D	P	inc
$4^1$	$x$	$(3^1)$	$4_1^1$	$x$	$(3_1^1)$	eq
			$4_2^1$		$(3_2^1)$	$\subset$
			$4_3^1$		$(3_3^1)$	$\subset$
			$4_4^1$		$(3_4^1)$	$\subset$
			$4_5^1$		$(3_5^1)$	$\subset$
$4^2$		$(3^2)$	$4_1^2$		$(3_1^2)$	eq
			$4_2^2$		$(3_2^2)$	$\subset$
			$4_3^2$		$(3_3^2)$	$\subset$
			$4_4^2$		$(3_4^2)$	$\subset$
$4^3$		$(3^3)$	$4_1^3$		$(3_1^3)$	eq
			$4_2^3$		$(3_2^3)$	$\subset$
$4^4$		$(3^4)$	$4_1^4$		$(3_1^4)$	eq
			$4_2^4$		$(3_2^4)$	$\subset$
$4^5$		$(3^5)$	$4_1^5$		$(3_1^5)$	eq

Figure 3: (Partial) DP computation for  $\subseteq$ -MINIMAL SAT without compression.

**Example 3.** We extend our running example by now considering  $\subseteq$ -MINIMAL SAT (i.e., enumerating models that are subset-minimal w.r.t. the set of atoms that get assigned “true”). Figure 3 illustrates the computation for parts of our example. At  $n_3$ ,  $R$  is computed as before. For any  $r \in R$ , each subrow  $s \in S(r)$  represents a partial “true” assignment that is a subset of the one in  $r$  (i.e.,  $D(s) \subseteq D(r)$ ), and  $\text{inc}(s)$  is set appropriately. Now consider  $n_4$ , where  $y$  and  $z$  are removed. As in the SAT problem, we simply project away data related to the removed vertices. Observe that for instance  $4^1$  now contains redundant subrows  $4_2^1, 4_3^1, 4_4^1$  and  $4_5^1$  that contain the same data and inclusion flag (and only different EPTs). We overcome this problem in Section 3.2, where tables are compressed in order to remove redundancies. A complete example including compression is given in Section 3.2.

The EPTs of a table row  $r$  are used for recursively combining the problem-specific data  $D(r)$  with data from “compatible” rows that are in descendant tables. The fact that each set of EPTs is required to be non-empty entails that for each (sub)row  $r$  at a leaf table it holds that  $P(r) = \{()\}$ . We disallow rows with an empty set of EPTs because in the end we are only interested in rows that can be extended to complete solutions, consisting of one row per table. For this we introduce the notion of an *extension* of a table row.

**Definition 3.** Let  $\mathcal{C}$  be a computation and  $R$  be a table in  $\mathcal{C}$  with  $k$  children. We inductively define the extensions of a row  $r \in R$  as  $E(r) = \{\{r\} \cup A \mid A \in \bigcup_{(p_1, \dots, p_k) \in P(r)} \{X_1 \cup \dots \cup X_k \mid X_i \in$

$E(p_i)$  for all  $1 \leq i \leq k$ }}.

Note that any extension  $X \in E(r)$  contains  $r$  and exactly one row from each table that is a descendant of  $R$ . If  $r$  is a row of a leaf table,  $E(r) = \{\{r\}\}$  because  $P(r) = \{\{\}\}$ .

While the extensions from the root table of a computation represent complete solution candidates, the purpose of subtables is to represent possible counterexamples that would cause a solution candidate to be invalidated. More precisely, for each extension  $X$  that can be obtained by extending a root table row  $r$ , we check if we can find an extension  $Y$  of an element  $s \in S(r)$  with  $\text{inc}(s) = \subset$  such that every element of  $Y$  is listed as a subrow of a row in  $X$  (i.e., we check if for every  $y \in Y$  there is some  $x \in X$  with  $y \in S(x)$ ). If this is so, then  $Y$  witnesses that  $X$  represents no solution because  $Y$  then represents a solution candidate that is a proper subset. For this reason, we need to introduce the notion of extensions (like  $Y$ ) relative to another extension (like  $X$ ).

**Definition 4.** Let  $\mathcal{C}$  be a computation,  $R$  be a table in  $\mathcal{C}$  with  $k$  children,  $r \in R$  be a row and  $s \in S(r)$  be a subrow of  $r$ . We first define, for any  $X \in E(r)$ , a restriction of  $P(s)$  to EPTs where each element is a subrow of a row in  $X$ , as  $P_X(s) = \{(p_1, \dots, p_k) \in P(s) \mid r_i \in X, p_i \in S(r_i) \text{ for all } 1 \leq i \leq k\}$ . Now we define the set of extensions of  $s$  relative to some extension  $X \in E(r)$  as  $E_X(s) = \{\{s\} \cup A \mid A \in \bigcup_{(p_1, \dots, p_k) \in P_X(s)} \{Y_1 \cup \dots \cup Y_k \mid Y_i \in E_X(p_i) \text{ for all } 1 \leq i \leq k\}\}$ .

We can now formalize that the solutions of a computation are the extensions of those rows that do not have a subrow indicating a counterexample.

**Definition 5.** Let  $R$  be the root table in a computation  $\mathcal{C}$ . We define the set of solutions of  $\mathcal{C}$  as  $\text{sol}(\mathcal{C}) = \{D(X) \mid r \in R, X \in E(r), \nexists s \in S(r) : \text{inc}(s) = \subset\}$

**Example 4.** If  $n_4$  in Figure 3 were the root of the TD, only  $A^5$  would yield a solution (i.e.  $\{\}$ , the interpretation where  $x, y$ , and  $z$  are all set to false), since all other rows contain a subrow where  $\text{inc}(s) = \subset$ . This is indeed the only subset-minimal model of the formula consisting of the clauses encountered until  $n_4$ , i.e.,  $(\neg x \vee z) \wedge (\neg x \vee y \vee \neg z)$ .

Next we formalize requirements on subrows and their inclusion status to ensure that subrows correspond to subsets of their parent row, that each potential counterexample is represented by a subrow and that  $\text{inc}(\cdot)$  is used as intended.

**Definition 6.** A table  $R$  is normal if the following properties hold:

1. For each  $r \in R, s \in S(r), X \in E(r)$  and  $Y \in E_X(s)$ , it holds that  $Y \leq X$ , and  $Y < X$  holds if and only if  $\text{inc}(s) = \subset$ .
2. For each  $r \in R, s \in S(r)$  and  $Y \in E(s)$  there is some  $r' \in R$  and  $X' \in E(r')$  such that  $s \approx r'$  and  $Y \approx X'$ .
3. For each  $q, r \in R, Z \in E(q)$  and  $X \in E(r)$ , if  $Z \leq X$  holds, then there is some  $s \in S(r)$  and  $Y \in E_X(s)$  with  $s \approx q$  and  $Y \approx Z$ .

A computation is normal if all its tables are normal.

This definition ensures that it suffices to examine the root table of a normal computation in order to decide a subset minimization problem correctly, provided that the rows represent all solution candidates.

We will later show how a non-minimizing computation (i.e., one with empty subtables) satisfying certain properties can be transformed into a normal computation. In this transformation, we must avoid redundancies lest we destroy fixed-parameter tractability. For this, we first introduce how tables can be compressed without losing solution candidates.

### 3.2 Table Compression

To compress tables by merging equivalent (sub)rows, which is required for keeping the size of the tables bounded by the width of the TD, we first define an equivalence relation on rows, as well as one on subrows.

**Definition 7.** *Let  $R$  be a table and  $r \in R$ . We define an equivalence relation  $\equiv_r$  over subrows of  $r$  such that  $s_1 \equiv_r s_2$  if  $s_1 \approx s_2$  and  $\text{inc}(s_1) = \text{inc}(s_2)$ .*

We use this notion of equivalence between subrows to compress subtables by merging equivalent subrows.

**Definition 8.** *Let  $R$  be a table and  $r \in R$ . We define a subtable  $S^*(r)$  called the compressed subtable of  $r$  that contains exactly one subrow for each  $\equiv_r$ -equivalence class. For any  $s \in S(r)$ , let  $[s]$  denote the  $\equiv_r$ -equivalence class of  $s$  and let  $s'$  denote the subrow in  $S^*(r)$  corresponding to  $[s]$ . We define  $s'$  by  $s' \approx s$ ,  $\text{inc}(s') = \text{inc}(s)$  and  $P(s') = \bigcup_{t \in [s]} P(t)$ .*

Once subtables have been compressed, we can compress the table by merging equivalent rows. For this, we first need a notion of equivalence between rows.

**Definition 9.** *We define an equivalence relation  $\equiv_R$  over rows of a table  $R$  such that  $r_1 \equiv_R r_2$  if  $r_1 \approx r_2$  and there is a bijection  $f : S^*(r_1) \rightarrow S^*(r_2)$  such that for any  $s \in S^*(r_1)$  it holds that  $s \approx f(s)$  and  $\text{inc}(s) = \text{inc}(f(s))$ .*

When rows are equivalent, their compressed subtables only differ in the EPTs. We now define how such compressed subtables can be merged.

**Definition 10.** *Let  $R$  be a table,  $r \in R$ , and let  $[r]$  denote the  $\equiv_R$ -equivalence class of  $r$ . For any  $r' \in [r]$ , let  $f_{r'} : S^*(r) \rightarrow S^*(r')$  be the bijection such that for any  $s \in S^*(r)$  it holds that  $s \approx f_{r'}(s)$  and  $\text{inc}(s) = \text{inc}(f_{r'}(s))$ . (The existence of  $f_{r'}$  is guaranteed by Definition 9.) We define a subtable  $\text{mst}([r])$  (for “merged subtable”) that contains exactly one subrow for each element of  $S^*(r)$ . For any  $s \in S^*(r)$ , let  $s'$  denote the subrow in  $\text{mst}([r])$  corresponding to  $s$ . We define  $s'$  by  $s' \approx s$ ,  $\text{inc}(s') = \text{inc}(s)$  and  $P(s') = \bigcup_{r' \in [r]} P(f_{r'}(s))$ .*

We use these equivalence relations to compress tables in such a way that all equivalent (sub)rows (according to the respective equivalence relation) are merged.

**Definition 11.** Let  $R$  be a table. We now define a table  $\text{compr}(R)$  that contains exactly one row for each  $\equiv_R$ -equivalence class. For any  $r \in R$ , let  $[r]$  denote the  $\equiv_R$ -equivalence class of  $r$  and let  $r'$  be the row in  $\text{compr}(R)$  corresponding to  $[r]$ . We define  $r'$  by  $r' \approx r$ ,  $P(r') = \bigcup_{q \in [r]} P(q)$  and  $S(r') = \text{mst}([r])$ . For any computation  $\mathcal{C}$ , we write  $\text{compr}(\mathcal{C})$  to denote the computation isomorphic to  $\mathcal{C}$  where each table  $R$  in  $\mathcal{C}$  corresponds to  $\text{compr}(R)$ .

**Example 5.** Figure 4 illustrates the complete DP computation for our running example, including compression. Root node  $n_6$  contains two rows,  $6^1$  and  $6^2$ .  $6^1$  represents the subset-minimal models of our running example, since it is associated with a single subrow  $6_1^2$ , where  $\text{inc}(6_1^2) = \text{eq}$ . Opposed to that, models represented by  $6^1$  are not subset-minimal, due to  $\text{inc}(6_2^1) = \subset$ . We obtain the solutions by following the EPTs and combining  $P$  of the respective rows: We have  $P(6^2) \cup P(5^4) \cup P(2^3) \cup P(1^2) \cup P(4^3) \cup P(3^5) = \{u\}$  and  $P(6^2) \cup P(5^5) \cup P(2^4) \cup P(1^3) \cup P(4^1) \cup P(3^1) = \{x, v, y, z\}$ . Thus the subset-minimal models (cf. solutions) are  $\{u\}$  (i.e.,  $u$  set to true, and  $v, x, y, z$  set to false) as well as  $\{v, x, y, z\}$  (with  $u$  set to false).

Regarding compression, consider the table stored at node  $n_4$  in comparison to the non-compressed table for  $n_4$  in Figure 3. For instance, due to Definition 10  $4^1$  now has subrow  $4_2^1$  that contains a set of EPTs, i.e.,  $P(4_2^1) = \{(3_2^1), (3_3^1), (3_4^1), (3_5^1)\}$ . Additionally, due to Definition 11 we can merge complete rows. For instance, observe that the single row  $4^2$  is obtained by compressing rows  $4^2$ ,  $4^3$  and  $4^4$  of Figure 3, since the data of these rows as well as the data and the inclusion flags of all subrows coincide, i.e., they are contained in the same  $\equiv_R$ -equivalence class.

We now show that compressing a table according to these definitions retains normality.

**Lemma 1.** If a table  $R$  is normal, then so is  $\text{compr}(R)$ .

*Proof sketch.* Let  $R$  be a normal table and  $R' = \text{compr}(R)$ . We prove conditions 1–3 of normality of  $R'$  separately. Let  $r' \in R'$ ,  $s' \in S(r')$ ,  $X' \in E(r')$  and  $Y'_{X'} \in E_{X'}(s')$ . Then we can find  $r \in R$ ,  $s \in S(r)$ ,  $X \in E(r)$  and  $Y \in E_X(s)$  such that  $r \approx r'$ ,  $s \approx s'$ ,  $X \approx X'$  and  $Y \approx Y'_{X'}$ . As  $R$  is normal,  $Y \leq X$  holds, and  $Y < X$  if and only if  $\text{inc}(s) = \subset$ . This entails  $Y'_{X'} \leq X'$ , and  $Y'_{X'} < X'$  if and only if  $\text{inc}(s') = \subset$  because  $\text{inc}(s') = \text{inc}(s)$ . This proves condition 1. Let  $Y' \in E(s')$ . Then we can find  $Y \in E(s)$  such that  $Y \approx Y'$ . As  $R$  is normal, there is a row  $t \in R$  with  $t \approx s$  and an extension  $T \in E(t)$  such that  $T \approx Y$ . Then there is some  $t' \in R'$  with  $t' \approx t$  and  $(T \setminus \{t\}) \cup \{t'\} \in E(t')$ , which proves condition 2. Let  $q' \in R$  and  $Z' \in E(q')$  such that  $Z' \leq X'$ . Then we can find  $q \in R$  and  $Z \in E(q)$  such that  $Z \approx Z'$ , so  $Z \leq X$ . As  $R$  is normal, there are  $u \in S(r)$  and  $U \in E_X(u)$  such that  $u \approx q$  and  $U \approx Z$ . Then there is some  $u' \in S(r')$  with  $u' \approx u$  and  $(U \setminus \{u\}) \cup \{u'\} \in E_{X'}(u')$ , which proves condition 3.  $\square$

### 3.3 Normalizing Computations

Before we introduce our transformation from “non-minimizing” computations to “minimizing” ones, we define certain conditions that are prerequisites for the transformation. For this, we first define the set of all data of rows that have occurred in a table or any of its descendants.

**Definition 12.** Let  $R$  be a table in a computation such that  $R_1, \dots, R_k$  are the child tables of  $R$ . We inductively define  $D^*(R) = \bigcup_{r \in R} D(r) \cup \bigcup_{1 \leq i \leq k} D^*(R_i)$ .

R			S(r)				
r	D	P	s	D	P	inc	
$n_6$	$6^1$		$(5^1), (5^2), (5^3)$	$6_1^1$		$(5_1^1), (5_1^2), (5_1^3)$	eq
				$6_2^1$		$(5_2^1), (5_3^1), (5_2^2), (5_2^3)$	$\subset$
	$6^2$		$(5^4), (5^5)$	$6_1^2$		$(5_1^4), (5_1^5)$	eq

R			S(r)				
r	D	P	s	D	P	inc	
$n_5$	$5^1$	$x$	$(2^1, 4^1)$	$5_1^1$	$x$	$(2_1^1, 4_1^1)$	eq
				$5_2^1$	$x$	$(2_2^1, 4_1^1), (2_4^1, 4_1^1)$	$\subset$
				$5_3^1$		$(2_3^1, 4_2^1)$	$\subset$
	$5^2$	$x$	$(2^2, 4^1)$	$4_1^2$	$x$	$(2_1^2, 4_1^1)$	eq
				$5_2^2$		$(2_2^2, 4_2^1)$	$\subset$
	$5^3$		$(2^3, 4^2)$	$5_1^3$		$(2_1^3, 4_2^1)$	eq
			$5_2^3$		$(2_1^3, 4_2^2)$	$\subset$	
$5^4$		$(2^3, 4^3)$	$5_1^4$		$(2_1^3, 4_1^3)$	eq	
$5^5$	$x$	$(2^4, 4^1)$	$5_1^5$	$x$	$(2_1^4, 4_1^1)$	eq	

R			S(r)				
r	D	P	s	D	P	inc	
$n_2$	$2^1$	$v, x$	$(1^1)$	$2_1^1$	$v, x$	$(1_1^1)$	eq
				$2_2^1$	$x$	$(1_2^1)$	$\subset$
				$2_3^1$		$(1_3^1)$	$\subset$
				$2_4^1$	$v, x$	$(1_3^1)$	$\subset$
	$2^2$	$x$	$(1^2)$	$2_1^2$	$x$	$(1_1^2)$	eq
			$2_2^2$		$(1_1^2)$	$\subset$	
$2^3$		$(1^2)$	$2_1^3$		$(1_1^2)$	eq	
$2^4$	$v, x$	$(1^3)$	$2_1^4$	$v, x$	$(1_1^3)$	eq	

R			S(r)				
r	D	P	s	D	P	inc	
$n_4$	$4^1$	$x$	$(3^1)$	$4_1^1$	$x$	$(3_1^1)$	eq
				$4_2^1$		$(3_2^1), (3_3^1), (3_4^1), (3_5^1)$	$\subset$
	$4^2$		$(3^2), (3^3), (3^4)$	$4_1^2$		$(3_1^2), (3_1^3), (3_4^1)$	eq
				$4_2^2$		$(3_2^2), (3_3^2), (3_4^2), (3_3^3), (3_2^4)$	$\subset$
$4^3$		$(3^5)$	$4_1^3$		$(3_1^5)$	eq	

R			S(r)				
r	D	P	s	D	P	inc	
$n_1$	$1^1$	$u, v$	$()$	$1_1^1$	$u, v$	$()$	eq
				$1_2^1$	$u$	$()$	$\subset$
				$1_3^1$	$v$	$()$	$\subset$
	$1^2$	$u$	$()$	$1_1^2$	$u$	$()$	eq
	$1^3$	$v$	$()$	$1_1^3$	$v$	$()$	eq

R			S(r)				
r	D	P	s	D	P	inc	
$n_3$	$3^1$	$x, y, z$	$()$	$3_1^1$	$x, y, z$	$()$	eq
				$3_2^1$	$y, z$	$()$	$\subset$
				$3_3^1$	$y$	$()$	$\subset$
				$3_4^1$	$z$	$()$	$\subset$
				$3_5^1$		$()$	$\subset$
$3^2$	$y, z$	$()$	$()$	$3_1^2$	$y, z$	$()$	eq
				$3_2^2$	$y$	$()$	$\subset$
				$3_3^2$	$z$	$()$	$\subset$
				$3_4^2$		$()$	$\subset$
$3^3$	$y$	$()$	$()$	$3_1^3$	$y$	$()$	eq
				$3_2^3$		$()$	$\subset$
$3^4$	$z$	$()$	$()$	$3_1^4$	$z$	$()$	eq
				$3_2^4$		$()$	$\subset$
$3^5$		$()$	$3_1^5$		$()$	eq	

Figure 4: DP computation for  $\subseteq$ -MINIMAL SAT.

Now we define conditions that the tables in a computation must satisfy for being eligible for our transformation.

**Definition 13.** Let  $R$  be a table in a computation such that  $R_1, \dots, R_k$  are the child tables of  $R$ , and let  $r, r' \in R$ . We say that  $d \in D(r)$  has been illegally introduced at  $r$  if there are  $(r_1, \dots, r_k) \in P(r)$  such that for some  $1 \leq i \leq k$  it holds that  $d \notin D(r_i)$  while  $d \in D^*(R_i)$ . Moreover, we say that  $d \in D(r') \setminus D(r)$  has been illegally removed at  $r$  if there is some  $X \in E(r)$  such that  $d \in X$ .

We now define the notion of an *augmentable* table, i.e., a table that can be used in our transformation.

**Definition 14.** We call a table  $R$  *augmentable* if the following conditions hold:

1. For all rows  $r \in R$  it holds that  $S(r) = \emptyset$ .
2. For all  $r, r' \in R$  with  $r \neq r'$  it holds that  $D(r) \neq D(r')$ .
3. For all  $r \in R$ ,  $(r_1, \dots, r_k) \in P(r)$ ,  $1 \leq h < j \leq k$ ,  $H \in E(r_h)$  and  $J \in E(r_j)$  it holds that  $D(H) \cap D(J) \subseteq D(r)$ .
4. No element of  $D(R)$  has been illegally introduced.
5. No element of  $D(R)$  has been illegally removed.

We call a computation *augmentable* if all its tables are *augmentable*.

These requirements are satisfied by reasonable TD-based DP algorithms (cf. [30]) as these usually do not put arbitrary data into the rows. Rather, the data in a row is typically restricted to information about bag elements of the respective TD node. For instance, condition 3 mirrors condition 3 of Definition 1, and condition 2 is usually satisfied by reasonable FPT algorithms because they avoid redundancies in order to stay fixed-parameter tractable.

Now we describe how *augmentable* computations can automatically be transformed into normal computations that take minimization into account. For any table  $R$  in an *augmentable* computation, this allows us to compute a new table  $\text{aug}(R)$  if for each child table  $R_i$  the table  $\text{aug}(R_i)$  has already been computed and compressed to  $\text{compr}(\text{aug}(R_i))$ .

**Definition 15.** We inductively define a function  $\text{aug}(\cdot)$  that maps each table  $R$  from an *augmentable* computation to a table. Let the child tables of  $R$  be called  $R_1, \dots, R_k$ . For any  $1 \leq i \leq k$  and  $r \in R_i$ , we write  $\text{res}(r)$  to denote  $\{q \in \text{compr}(\text{aug}(R_i)) \mid q \approx r\}$ . We define  $\text{aug}(R)$  as the smallest table that satisfies the following conditions:

1. For any  $r \in R$ ,  $(r_1, \dots, r_k) \in P(r)$  and  $(c_1, \dots, c_k) \in \text{res}(r_1) \times \dots \times \text{res}(r_k)$ , there is a row  $q \in \text{aug}(R)$  with  $q \approx r$  and  $P(q) = \{(c_1, \dots, c_k)\}$ .
2. For any  $q, q' \in \text{aug}(R)$  such that  $q' \leq q$ ,  $P(q) = \{(q_1, \dots, q_k)\}$  and  $P(q') = \{(q'_1, \dots, q'_k)\}$  the following holds: If for all  $1 \leq i \leq k$  there is some  $s_i \in S(q_i)$  with  $s_i \approx q'_i$ , then there is a subrow  $s \in S(q)$  with  $s \approx q'$  and  $P(s) = \{(s_1, \dots, s_k)\}$ . Moreover,  $\text{inc}(s) = \subset$  if  $q' < q$  or  $\text{inc}(s_i) = \subset$  for some  $s_i$ , otherwise  $\text{inc}(s) = \text{eq}$ .

For any *augmentable* computation  $\mathcal{C}$ , we write  $\text{aug}(\mathcal{C})$  to denote the computation isomorphic to  $\mathcal{C}$  where each table  $R$  in  $\mathcal{C}$  corresponds to  $\text{aug}(R)$ .

Augmentable tables never have two different rows  $r, r'$  with  $D(r) = D(r')$ . Moreover, we defined  $\text{aug}(R)$  in such a way that for all  $r \in R$  there is some  $q \in \text{aug}(R)$  with  $r \approx q$ . In the compression  $\text{compr}(\text{aug}(R))$ , we only merge rows and subrows having the same data. So with each row and subrow in  $\text{aug}(R)$  or  $\text{compr}(\text{aug}(R))$  we can associate a unique originating row in  $R$ . In fact, the extensions from an augmentable table  $R$  are in a one-to-one correspondence to the extensions of rows from  $\text{aug}(R)$ . This is formalized by the following lemma, which we prove in the appendix.

**Lemma 2.** *Let  $R$  be a table from an augmentable computation and  $Q = \text{aug}(R)$ . Then for any  $r \in R$  and  $Z \in E(r)$  there are  $q \in Q$  and  $X \in E(q)$  such that  $r \approx q$  and  $Z \approx X$ . Also, for any  $q \in Q$  and  $X \in E(q)$  there are  $r \in R$  and  $Z \in E(r)$  such that  $q \approx r$  and  $X \approx Z$ .*

The following lemma is central for showing that  $\text{aug}(\cdot)$  works as intended. For a full proof, see the appendix.

**Lemma 3.** *Let  $R$  be a table from an augmentable computation. Then the table  $\text{aug}(R)$  is normal.*

*Proof sketch.* Let  $R$  be a table in some augmentable computation such that  $R_1, \dots, R_k$  denote the child tables of  $R$  and let  $Q = \text{aug}(R)$ . We use induction. If  $Q$  is a leaf table, then rows and extensions coincide and the construction of  $Q$  obviously ensures that  $Q$  is normal. If  $Q$  has child tables  $Q_i = \text{compr}(\text{aug}(R_i))$  and all  $\text{aug}(R_i)$  are normal, all  $Q_i$  are normal by Lemma 1. Let  $q \in Q$ ,  $s \in S(q)$ ,  $P(q) = \{(q_1, \dots, q_k)\}$ ,  $P(s) = \{(s_1, \dots, s_k)\}$ ,  $X_i \in E(q_i)$ ,  $Y_i \in E_{X_i}(s_i)$ ,  $X = \{q\} \cup X_1 \cup \dots \cup X_k$  and  $Y = \{s\} \cup Y_1 \cup \dots \cup Y_k$ . As for normality condition 1, the construction of  $Q$  ensures  $s \leq q$  and normality of  $Q_i$  ensures  $Y_i \leq X_i$ , so  $Y \leq X$ . To show that  $\text{inc}(s)$  has the correct value, first suppose  $Y < X$  and  $\text{inc}(s) = \text{eq}$ . The latter would entail  $Y_i = X_i$ , so  $s < q$ , but then  $\text{inc}(s) = \subset$ , which is a contradiction. So suppose  $X \approx Y$  and  $\text{inc}(s) = \subset$ . If  $q < s$ , there would be an illegal removal at the origin of  $s$  in  $R$ , contradicting that  $R$  is augmentable. So for some  $j$  there is a  $d \in D(X_j) \setminus D(Y_j)$ . Due to  $X \approx Y$ ,  $d \in D(s)$  or  $d \in D(Y_h)$  for some  $h \neq j$ . In the first case, there is an illegal introduction at the origin of  $s$  in  $R$ . In the other case,  $d \in D(Y_h)$  entails  $d \in D(X_h)$ . As row extensions in  $Q$  are in a one-to-one correspondence with those in  $R$ , and by augmentability of  $R$ ,  $d \in D(X_j) \cap D(X_h)$  entails  $d \in D(q)$ . But then  $d \in D(s)$ , which we already led to a contradiction.

For condition 2, let  $Z_i \in E(s_i)$  and  $Z = \{s\} \cup Z_1 \cup \dots \cup Z_k$ . As  $s \in S(q)$ , there are  $p \in Q$  and  $(p_1, \dots, p_i) \in P(p)$  with  $p \approx s$  and  $p_i \approx s_i$ . This entails existence of  $r \in R$  and  $(r_1, \dots, r_k) \in E(r)$  with  $r \approx p \approx s$  and  $r_i \approx p_i$ . By hypothesis, there are  $q'_i \in Q_i$  and  $X'_i \in E(q'_i)$  with  $q'_i \approx s_i$  and  $X'_i \approx Z_i$ . Each  $q'_i$  originates from the unique  $r_i \in R$  with  $r_i \approx q'_i$ . So  $q'_i \in \text{res}(r_i)$  holds and there are  $q' \in Q$  and  $X' \in E(q')$  with  $q' \approx r \approx s$  and  $X' \approx Z$ .

For condition 3, let  $q' \in Q$ ,  $P(q') = \{(q'_1, \dots, q'_k)\}$ ,  $X' \in E(q')$  and  $X'_i \in E(q'_i)$  for all  $1 \leq i \leq k$ . Suppose  $X' \leq X$  and, for the sake of contradiction, for some  $j$  there is a  $d \in D(X'_j) \setminus D(X_j)$ . Then  $d \in D(q)$  or  $d \in D(X_h)$  for some  $h \neq j$ . In the first case, there is an illegal introduction at the origin of  $q$  in  $R$ . In the other case,  $d \in D(X_j) \cap D(X_h)$  entails  $d \in D(q)$ , which we already led to a contradiction. So  $X'_i \leq X_i$  for each  $i$ . By hypothesis then there are  $t_i \in S(q_i)$  and  $T_i \in E_{X_i}(t_i)$  with  $t_i \approx q'_i$  and  $T_i \approx X'_i$ . So there is a  $t \in S(q)$  with  $t \approx q'$  and  $P(t) = \{(t_1, \dots, t_k)\}$ . Then  $T = \{t\} \cup T_1 \cup \dots \cup T_k$  is in  $E_X(t)$  and  $T \approx X'$ .  $\square$

We can now state our main theorem, which says that exactly the subset-minimal solutions of an augmentable computation are solutions of the augmented computation.

**Theorem 1.** *Let  $\mathcal{C}$  be an augmentable computation. Then  $\text{sol}(\text{aug}(\mathcal{C})) = \{S \in \text{sol}(\mathcal{C}) \mid \nexists S' \in \text{sol}(\mathcal{C}) : S' \subset S\}$ .*

*Proof.* Let  $R$  be the root table of an augmentable computation  $\mathcal{C}$  and  $R'$  be the root table of  $\mathcal{C}' = \text{aug}(\mathcal{C})$ . By Lemma 3,  $\mathcal{C}'$  is normal. For the first direction, let  $S \in \text{sol}(\mathcal{C}')$ . Then there are  $r' \in R'$  and  $X' \in E(r')$  such that  $D(X') = S$  and for all  $s' \in S(r')$  it holds that  $\text{inc}(s') = \text{eq}$ . By Lemma 2, then there are  $r \in R$  and  $X \in E(r)$  such that  $X \approx X'$ . As  $R$  is augmentable,  $S(r)$  is empty, so  $S \in \text{sol}(\mathcal{C})$  by Definition 5. We must now show that there is no solution in  $\mathcal{C}$  smaller than  $S$ . For the sake of contradiction, suppose there is some  $T \in \text{sol}(\mathcal{C})$  with  $T \subset S$ . Then there are  $q \in R$  and  $Z \in E(q)$  such that  $D(Z) = T$ , hence  $Z < X'$ . By Lemma 2, then there are  $q' \in R'$  and  $Z' \in E(q')$  such that  $Z' \approx Z$ . As  $R'$  is normal, due to  $Z' < X'$ , there is some  $s' \in S(r')$  such that  $\text{inc}(s') = \subset$ . This contradicts  $\text{inc}(s') = \text{eq}$ , which we have seen earlier.

For the other direction, let  $S \in \text{sol}(\mathcal{C})$  be such that there is no  $S' \in \text{sol}(\mathcal{C})$  with  $S' \subset S$ . Then there are  $r \in R$  and  $X \in E(r)$  such that  $D(X) = S$ . By Lemma 2, then there are  $r' \in R'$  and  $X' \in E(r')$  such that  $X' \approx X$ , and there are no  $q' \in R'$  and  $Z' \in E(q')$  with  $Z' < X$ . Hence, as  $R'$  is normal, there cannot be a  $s' \in S(r')$  with  $\text{inc}(s') = \subset$ . This proves that  $S \in \text{sol}(\mathcal{C}')$ .  $\square$

Finally, we sketch that  $\text{aug}(\cdot)$  does not destroy fixed-parameter tractability.

**Theorem 2.** *Let  $\mathcal{A}$  be an algorithm that takes as input an instance of size  $n$  and treewidth  $w$  along with a TD  $\mathcal{T}$  of width  $w$ . Suppose  $\mathcal{A}$  produces an augmentable computation  $\mathcal{C}$  isomorphic to  $\mathcal{T}$  in time  $f(w) \cdot n^{\mathcal{O}(1)}$ , where  $f$  is a function depending only on  $w$ . Then  $\text{aug}(\mathcal{C})$  can be computed in time  $g(w) \cdot n^{\mathcal{O}(1)}$ , where  $g$  again depends only on  $w$ .*

*Proof sketch.* We assume w.l.o.g. that each node in  $\mathcal{T}$  has at most 2 children, as any TD can be transformed to this form in linear time without increasing the width [25]. As  $\mathcal{A}$  runs in FPT time, i.e., in  $f(w) \cdot n^{\mathcal{O}(1)}$  for a function  $f$ , no table in  $\mathcal{C}$  can be bigger than  $f(w) \cdot n^c$  for a constant  $c$ . To inductively compute  $Q = \text{aug}(R)$  for some table  $R$  in  $\mathcal{C}$  with child tables  $R_1, \dots, R_k$  ( $k \leq 2$ ), suppose we have already constructed each  $Q_i = \text{aug}(R_i)$  in FPT time. Then  $|Q_i| = f_i(w) \cdot n^{c_i}$  for some  $f_i$  and  $c_i$ . We can compute  $Q'_i = \text{compr}(Q_i)$  in time polynomial in  $|Q_i|$ . Then  $|Q'_i| = f'_i(w) \cdot n^{c'_i}$  for some  $f'_i$  and  $c'_i$ . Definition 15 suggests a straightforward way to compute  $Q$  in time polynomial in  $|R|$  and  $\sum_{1 \leq i \leq k} |Q'_i|$ . So we can compute  $Q$  in FPT time. As  $\mathcal{T}$  has size  $\mathcal{O}(n)$ , we can compute  $\text{aug}(\mathcal{C})$  in FPT time.  $\square$

## 4 Practical Realization

In this section we illustrate how the concepts introduced so far can be put into practice. We first recall the D-FLAT system, a system which enables specifications of DP algorithms on TDs via ASP. Then, we introduce a new variant of D-FLAT, called D-FLAT<sup>2</sup>. While D-FLAT requires an explicit encoding of subset minimization within the DP algorithm, subset minimization is handled

automatically in D-FLAT<sup>2</sup>, provided that the encoded problem produces an augmentable computation. Both systems use the library Htdecomp [13] for the generation of the tree decomposition and use state-of-the-art ASP technology from the Potassco family [20] for running the encodings of the DP algorithms.

## 4.1 D-FLAT: A Quick Tutorial

D-FLAT [1] is a framework for simplifying DP on TDs such that the user only has to provide an ASP program formulating the DP algorithm on TDs for solving the given graph-based problem. It automatically constructs a TD of the input graph in polynomial time (using heuristics). Then, the generated TD is traversed in post-order. At each DP node the tables containing the solution candidates are computed. There, an ASP solver is called with the following input:

1. the user-specified ASP encoding,
2. the input instance,
3. information about the current and child TD nodes (see Table 1), and
4. the partial solutions computed in the child TD nodes (see below).

Each model returned by the solver represents a partial solution of the current node.<sup>2</sup> In order to represent partial solutions, D-FLAT provides a general interface that supports (amongst others) rows *and* subrows of a table as introduced in Section 3. Additionally, for problems in NP, a simplified interface that *only* supports rows is available. We will now briefly describe the concepts behind these interfaces, as well as their relation to the formal concepts presented in Section 3.

**Simplified interface (for problems in NP).** The simplified interface supports tables that contain rows *without* subrows. When traversing the TD, tables for the already-visited child nodes are given to the user-specified encoding via predicates as listed in Table 2. Then, the partial solutions for the current TD node are computed via the user-specified encoding, and returned via the output predicates listed in Table 3. A row  $r$  consists of a set of *items* and *auxiliary items* that store problem-specific data  $D(r)$ <sup>3</sup>. The output predicate `extend/1` specifies the child row(s) that give rise to the partial solution encoded by the respective answer set.

Listing 1 contains the user-specified ASP encoding  $\Pi_{\text{SAT}}$  for the SAT problem. The encoding makes use of D-FLAT’s input interface (see Tables 1 and 2) in the bodies of the rules, and the output interface (see Table 3) in the heads of the rules. Additionally, it expects that the input instance  $\phi$  is encoded as follows. Atoms and clauses occurring in  $\phi$  are given as ASP facts `atom( $\cdot$ )` and `clause( $\cdot$ )` respectively. Facts `pos( $c, a$ )` (`neg( $c, a$ )`) denote that some atom  $a$  occurs positively (negatively) in clause  $c$ .

<sup>2</sup>We use colors to highlight `input` (red), `output` (orange) and `input instance` (blue) predicates.

<sup>3</sup>D-FLAT provides an efficient implementation for joining tables that emerge from different child nodes, which requires to specify items that are subject to set equivalence and auxiliary items that are subject to set union. Details are given in [1].

Table 1: Input predicates describing the tree decomposition.

Input predicate	Meaning
<code>final</code>	The current tree decomposition node is the root.
<code>childNode(<math>n</math>)</code>	$n$ is a child of the current decomposition node.
<code>bag(<math>n, v</math>)</code>	Vertex $v$ is contained in the bag of the decomposition node $n$ .
<code>current(<math>v</math>)</code>	Vertex $v$ is an element of the current bag.
<code>introduced(<math>v</math>)</code>	Vertex $v$ is a current vertex but was in no child node’s bag.
<code>removed(<math>v</math>)</code>	Vertex $v$ was in a child node’s bag but is not in the current one.

Table 2: Input predicates describing tables of decomposition child nodes.

Input predicate	Meaning
<code>childRow(<math>r, n</math>)</code>	$r$ is a table row belonging to decomposition node $n$ .
<code>childItem(<math>r, i</math>)</code>	$D(r)$ of child row $r$ contains item $i$ .
<code>childAuxItem(<math>r, i</math>)</code>	$D(r)$ of child row $r$ contains auxiliary item $i$ .

Table 3: Output predicates for constructing the table of the current decomposition node.

Output predicate	Meaning
<code>item(<math>i</math>)</code>	The data associated with the current row shall contain item $i$ .
<code>auxItem(<math>i</math>)</code>	The data associated with the current row shall contain auxiliary item $i$ .
<code>extend(<math>r</math>)</code>	The EPT of the current row shall contain a reference to child table row $r$ .

**Example 6.** We will now go through Listing 1 and explain how the partial solutions of our running example (depicted in Figure 2) are computed by D-FLAT. Let us first consider leaf node  $n_1$  of  $\mathcal{T}_{Ex}$ . Since  $n_1$  has no children, only Lines 6 and 12-19 are of interest to us. For all atoms in  $\chi(n_1) = \{u, v\}$  we guess their truth assignment (Line 6). In case an atom gets assigned true, it is stored as `item` in the data of the computed row. Lines 12-14 specify the “current” clauses (i.e. clauses, such that all clause atoms occur in the bag of the current TD node). In our example, this yields clause  $(u \vee v)$ . Lines 16-17 denote whether such a clause is satisfied by the current truth assignment. Clause  $(u \vee v)$  is satisfied by  $\{u, v\}$ ,  $\{u\}$  and  $\{v\}$ . However,  $\{\}$  is no model of the clause, and removed by the constraint in Line 19. In  $n_2$ ,  $\{u\}$  is removed and  $\{w, x\}$  is introduced to the bag. Whenever an atom was assigned false in a child row, Line 2 makes this explicit. We now `extend` each partial solution of the child node (Line 4). In Figure 2, this extension is denoted in column P. Their data (truth assignment), restricted to the current bag, is kept via Line 10. For introduced vertices, their truth assignment is guessed in Line 6. Similar to before, the “current” clauses are computed, and rows yielding unsatisfied current clauses are removed. In TD nodes having several children, we extend exactly one row per child table at a time (Line 4). Extended partial solutions have to agree on the truth assignment of atoms in the current bag

---

```

1 % Define false atoms
2 f(R, X) ← childRow(R, N), bag(N, X), not childItem(R, X).
3 % Guess partial solutions to be extended
4 1 { extend(R) : childRow(R, N) } 1 ← childNode(N).
5 % Guess truth value of introduced atoms
6 { item(A) : introduced(A) }.
7 % Only join rows coinciding on truth values of atoms
8 ← extend(X; Y), childItem(X, A), f(Y, A).
9 % True atoms are kept
10 item(X) ← extend(R), childItem(R, X), current(X).
11 % Current clauses
12 notCurr(C) ← pos(C, A), not current(A).
13 notCurr(C) ← neg(C, A), not current(A).
14 curr(C) ← clause(C), not notCurr(C).
15 % Define satisfied clauses
16 sat(C) ← curr(C), pos(C, A), item(A).
17 sat(C) ← curr(C), neg(C, A), not item(A).
18 % Current clauses need to be satisfied
19 ← curr(C), not sat(C).

```

---

Listing 1:  $\Pi_{\text{SAT}}$ : D-FLAT encoding for solving SAT.

(Line 8). Consider join node  $n_5$ . For instance, since  $\chi(n_5) \cap D(2^1) = \chi(n_5) \cap D(4^1)$ , we can construct row  $5^1$  with  $D(5^1) = \{x\}$  and EPT  $(2^1, 4^1)$ .

**General interface (for problems beyond NP).** D-FLAT provides an interface that supports a more general data structure, called *item tree*. Here, we first briefly introduce item trees. Then, we show how item trees can be used to represent rows *with* subtables. For details we refer to [2]. Each node in the TD can be associated with an item tree. The predicates specifying item trees computed in the child nodes are given in Table 4, output predicates are given in Table 5. Similar to the simplified interface, each node in an item tree contains data, specified by `item/2` and `auxItem/2` predicates. Each item tree node additionally has an EPT that represents its origin, that contains extension pointers denoted by `extend/2` (Note that these are now binary predicates). Each root-to-leaf path has a particular `length/1`, and the *level* of an item tree node is its depth on the path.

At the TD’s root, each item tree node must be labeled with either `accept` or `reject` if it is a leaf, otherwise with `or/1` or `and/1`. D-FLAT uses this to filter out solution candidates for which “counterexamples” exist: Only so-called *accepting* nodes are kept at the TD root, where a node is accepting if a) its label is `accept`, or b) its label is `or` and at least one child is accepting, or c) its label is `and` and all children are accepting. One can view the simplified interface as a special case of the interface for item trees, where the `length` of each root-to-leaf path is one, and the root node is of type `or`. Furthermore, contrary to tables, where each model returned by the ASP solver represents a row, for item trees a model represents a single root-to-leaf path in the item tree.

For solving subset minimization problems, we use item trees to represent tables  $R$  with subtables as follows. The `length` of each root-to-leaf path is two. The root node (level 0) of the item

Table 4: Input predicates describing item trees of child nodes in the decomposition.

Input predicate	Meaning
<code>atNode(s, n)</code>	$s$ is an item tree node belonging to decomposition node $n$ .
<code>rootOf(s, n)</code>	$s$ is the root of the item tree at decomposition node $n$ .
<code>sub(r, s)</code>	$r$ is an item tree node with child $s$ .
<code>childItem(s, i)</code>	The data of item tree node $s$ contains item $i$ .
<code>childAuxItem(s, i)</code>	The data of item tree node $s$ contains auxiliary item $i$ .

Table 5: Output predicates for constructing the item tree of the current decomposition node.

Output predicate	Meaning
<code>item(l, i)</code>	Item tree node at level $l$ in the current root-to-leaf path shall contain item $i$ .
<code>auxItem(l, i)</code>	Item tree node at level $l$ in the current root-to-leaf path shall contain auxiliary item $i$ .
<code>extend(l, s)</code>	The EPT associated with the item tree node at level $l$ in the current root-to-leaf path shall contain a reference to item tree node $s$ .
<code>length(l)</code>	The current root-to-leaf path has length $l$ .
<code>or(l)/and(l)</code>	The node at level $l$ in the current root-to-leaf path has type “or”/“and”.
<code>accept/reject</code>	The leaf in the current root-to-leaf path has type “accept”/“reject”.

tree is of type `or`. The nodes at level 1 are of type `and`. Here, the idea is to construct item trees, such that some (`or`) solution candidate (stored at level 1) is accepted if there is no counterexample (`and`) at level 2 that is smaller. For a row  $r \in R$ , we specify each item  $i \in D(r)$  by `item(1, i)` (or `auxItem(1, i)`, respectively).  $P(r)$  is represented by predicates `extend(1, pi)` where  $p_i$  is a row in the  $i$ -th child of  $R$ . For a subrow  $s \in S(r)$ ,  $i \in D(s)$  is given via `item(2, i)` predicates (or `auxItem(2, i)`), and EPTs  $P(s)$  are given by `extend(2, pi)` for some  $p_i$ . The inclusion status flag  $\text{inc}(s) = \subset$  is represented by `auxItem(2, smaller)`, for  $\text{inc}(s) = \text{eq}$  no flag is stored.

**Example 7.** *In the following we explain this concept on basis of  $\subseteq$ -MINIMAL SAT. In D-FLAT, it is again only required to specify a single ASP encoding, depicted in Listing 2. The encoding defines that data for both solution candidates and counterexamples are computed as in the SAT problem. Additionally, this encoding ensures that partial interpretations represented by counterexamples are subsets of partial interpretations represented by solution candidates. Line 1 defines the structure of the item trees used for the subset minimization problem. Similar to  $\Pi_{\text{SAT}}$ , Line 3 makes explicit if an atom is false in the data of an item tree node. Exactly one root-to-leaf path of each child item tree is extended (Lines 5-6). Then, for each level the truth assignment of introduced atoms is guessed. True atoms are stored via `item/2` in the data of the row or subrow (Line 8). Root-to-leaf paths are only joined in case the respective item sets coincide on their partial interpretations for current atoms (Line 10). Truth assignments of extended child (sub)rows for current atoms are*

---

```

1 length(2). or(0). and(1).
2% Define false atoms
3 f(S,X) ← atNode(S,N), sub(_,S), childNode(N), bag(N,X), not childItem(S,X).
4% Guess root-to-leaf paths in item trees to be extended
5 extend(0,R) ← root(R).
6 1 { extend(L+1,S) : sub(R,S) } 1 ← extend(L,R), L<2.
7% Guess truth value of introduced atoms
8 { item(2,A;1,A) : introduced(A) }.
9% Only join root-to-leaf paths coinciding on atom truth values
10 ← extend(L,X;L,Y), atom(A), childItem(X,A), f(Y,A), L=1..2.
11% True atoms are kept
12 item(L,X) ← extend(L,R), childItem(R,X), current(X), L=1..2.
13% Current clauses
14 notCurr(C) ← pos(C,A), not current(A).
15 notCurr(C) ← neg(C,A), not current(A).
16 curr(C) ← clause(C), not notCurr(C).
17% Define satisfied clauses
18 sat(L,C) ← curr(C), pos(C,A), item(L,A), L=1..2.
19 sat(L,C) ← curr(C), neg(C,A), not item(L,A), L=1..2.
20% Current clauses need to be satisfied
21 ← curr(C), not sat(L,C), L=1..2.
22% Interpretation at level 2 must be subset of that at level 1
23 ← item(2,A), not item(1,A).
24% Update subset information; reject larger models at root
25 auxItem(2,smaller) ← extend(2,S), childAuxItem(S,smaller).
26 auxItem(2,smaller) ← atom(A), item(1,A), not item(2,A).
27 reject ← final, auxItem(2,smaller).
28 accept ← final, not reject.

```

---

Listing 2: D-FLAT encoding for solving  $\subseteq$ -MINIMAL SAT.

kept (Line 12). As in Listing 1, “current” clauses are computed (Lines 14-16), which have to be satisfied both by the partial interpretation stored for the solution candidate and the counterexample (Line 21). Line 23 guarantees that the truth assignment in a counterexample is a subset of (or equal to) that of a solution candidate. Flag `smaller` denotes that the counterexample represents a proper subset of the solution candidate (Lines 25-26). In the final (i.e., root) node of the TD, solution candidates are rejected that still have a smaller counterexample, and accepted otherwise (Lines 27-28).

Recall Figure 4 which contains the complete computation for our running example. Our encoding guarantees that the algorithm behaves exactly as described in the previous section. Observe that each model of the encoding matches a root-to-leaf path in the item tree, i.e., it contains exactly one row and subrow entry.

## 4.2 Towards D-FLAT<sup>2</sup>

We implemented our approach in a system called D-FLAT<sup>2</sup> by extending D-FLAT. D-FLAT<sup>2</sup> avoids redundant computations of solution candidates and potential counterexamples. It supports minimization and maximization on user-specified items (e.g., for  $\subseteq$ -MINIMAL SAT, on atoms). D-FLAT<sup>2</sup> directly makes use of  $\text{aug}(\cdot)$  as given in Definition 15; it is thus a problem-independent software framework for subset optimization whose input is (1) an algorithm  $\mathcal{A}$  for solving a problem without subset minimization, (2) an instance of this problem and (3) a tree decomposition (which actually can be generated from the input with the help of heuristics). By running  $\mathcal{A}$  and transforming the resulting computation  $\mathcal{C}$  into  $\text{aug}(\mathcal{C})$ , Theorem 1 guarantees that the solutions are exactly the subset-minimal ones of  $\mathcal{C}$ .

**Interface.** D-FLAT<sup>2</sup> supports the simplified interface of D-FLAT. Internally, instead of using item trees for representing rows and subrows that explicitly represent potential counterexamples, D-FLAT<sup>2</sup> refines them by using so-called *reduced item trees*. Recall item trees and the simplified user interface for problems in NP, since reduced item trees work in a similar way. In a reduced item tree every root-to-leaf path is of length one. For a table  $R$ , rows  $r \in R$  are stored as usual: each item  $i \in D(r)$  is given by `item(i)` (or `auxItem(i)`, respectively).  $P(r)$  is represented by predicates `extend(pi)` where  $p_i$  is a row in the  $i$ -th child of  $R$ . Instead of storing subrows explicitly, D-FLAT<sup>2</sup> only retains a set of so-called *counterexample pointers* associated with  $r$ . For a subrow  $s \in S(r)$ , a counterexample pointer is a pair  $(c, \text{inc}(s))$ , where  $c$  is a pointer to some  $r' \in R$ . By using counterexample pointers, it is no longer required to store counterexamples explicitly. Instead, solution candidates now also serve as countercandidates whenever they are referred to in a set of counterexample pointers. With this, it suffices to update solution candidates during the bottom-up traversal of the TD.

**Program flow.** Figure 5 illustrates the program flow of D-FLAT<sup>2</sup>. The system parses the input instance and computes some tree decomposition heuristically using the `Htdecomp` library [13]. Then, in a first pass over the TD, solution candidates are computed as in D-FLAT. In contrast to D-FLAT, which computes counterexample candidates directly during the first (and only) pass, D-FLAT<sup>2</sup> traverses the TD a second time. There, the counterexample pointers are computed (along the lines of Definition 15). After the second pass, D-FLAT<sup>2</sup> is able to decide the problem or to materialize the solutions.

The two-pass approach allows us to exclude possible counterexamples that turn out not to correspond to a solution candidate after all. First we compute all rows without their subtables, then we delete rows that do not lead to solutions and finally we apply  $\text{aug}(\cdot)$  to the resulting tables. This way, we avoid storing pointers to counterexample candidates that appear in no extension at the root table.

**Extensions.** D-FLAT<sup>2</sup> generalizes our approach by also supporting problems where only a certain part of the data (instead of all data) is subject to minimization<sup>4</sup>. Towards this, D-FLAT<sup>2</sup> supports a new output predicate `optItem/1`, where `optItem(i)` means that item  $i$  is subject to op-

---

<sup>4</sup>Note that in the previous section we provided definitions and proofs only for the special case where minimization is applied to all data, as the generalization leads to more cumbersome notation but does not change the nature of the approach.

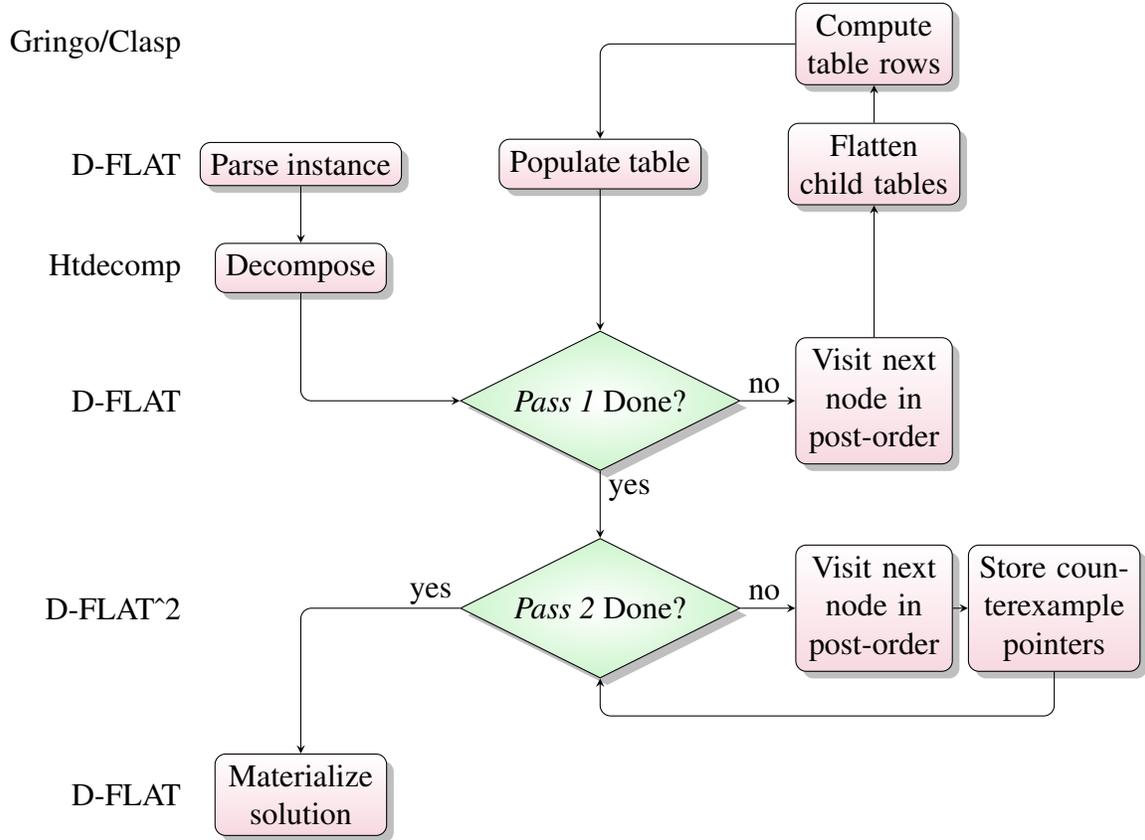


Figure 5: Flowchart for D-FLAT<sup>2</sup>.

timization. Using this generalization, we are able to obtain DP algorithms for further AI problems like circumscription [29] (see the next section).

Beside minimization, D-FLAT<sup>2</sup> also supports subset maximization over sets of items. Basically this works analogous to Section 3 by maintaining inclusion flags  $\text{eq}$  and  $\supset$  (instead of  $\subset$ ); subtables are then used to store potential counterexamples to a solution candidate being subset-maximal.

For some problems we require counterexample candidates that are not solution candidates at the same time. An example will be given in Section 5, where we consider disjunctive ASP. There, possible counterexamples are models of a program reduct and not the program itself. In such cases, the rows representing such counterexample candidates can be marked with a special flag as “pseudo solution candidates” by means of the atom `auxItem(pseudo)`, which is taken into account by D-FLAT<sup>2</sup>. Then, the only thing required from the user of D-FLAT<sup>2</sup> is an ASP encoding where each model corresponds to a single solution or counterexample candidate marked with `pseudo`.

## 5 Application to Common AI Problems

In the following we outline some practical applications where D-FLAT<sup>2</sup> is well suited.

### 5.1 SAT

Program  $\Pi_{\text{optAllItems}} = \{\text{optItem}(X) \leftarrow \text{item}(X) .\}$  uses the `optItem`/1 predicate in a trivial way. With this single additional rule at hand, we obtain the D-FLAT<sup>2</sup> encoding

$$\Pi_{\subseteq\text{-MINIMAL SAT}} = \Pi_{\text{optAllItems}} \cup \Pi_{\text{SAT}}$$

which allows us to solve  $\subseteq\text{-MINIMAL SAT}$  by simply using the existing encoding for SAT and adding information on what to optimize. In this case, the basic specification  $\Pi_{\text{optAllItems}}$  suffices, as in  $\subseteq\text{-MINIMAL SAT}$  we consider all items for minimization. D-FLAT<sup>2</sup> encoding  $\Pi_{\subseteq\text{-MINIMAL SAT}}$  gives several advantages over the traditional D-FLAT encoding (see Listing 2): It allows us to again use the simplified interface (designed for problems in NP) and the overall length of the encoding is greatly reduced.

### 5.2 Circumscription

In the propositional case of Circumscription [29], we are given a theory  $T$  and sets of atoms  $P$  and  $Z$ , and we are interested in models  $M$  of  $T$  such that there is no model  $M'$  with  $M' \cap P \subset M \cap P$  and  $M \cap Z = M' \cap Z$ . The formula whose classical models correspond to exactly those solutions is denoted by  $\text{CIRC}(T; P; Z)$ . We can model Circumscription in our approach by a slight modification of our  $\subseteq\text{-MINIMAL SAT}$  algorithm: We only put an atom  $x$  in an optimization item set if  $x \in P$ ; and for any  $x \in Z$  we add optimization items  $t(x)$  or  $f(x)$  if the item set contains  $x$  or not, respectively (thus making solution candidates with different interpretations of  $Z$  incomparable). By applying this technique we have that for any sibling rows  $r$  and  $r'$  in a table of the computation, the optimization item set of  $r$  is a subset of the one of  $r'$  iff  $M_r \cap P \subseteq M_{r'} \cap P$  and  $M_r \cap Z = M_{r'} \cap Z$ , where  $M_r$  is a partial interpretation for row  $r$  following the EPTs. The program  $\Pi_{\text{optForCirc}}$  (Listing 3) takes these considerations into account; moreover,

$$\Pi_{\text{CIRC}} = \Pi_{\text{optForCirc}} \cup \Pi_{\text{SAT}}$$

then gives a D-FLAT<sup>2</sup> implementation of the DP algorithm for Circumscription. As input it expects the theory  $T$  to be given as a CNF formula like for  $\Pi_{\text{SAT}}$  and the sets  $P$  and  $Z$  to be given using the predicates `p`/1 and `z`/1, respectively.

---

```

1 optItem(X)      ←      item(X), p(X) .
2 optItem(t(X))  ←      item(X), z(X) .
3 optItem(f(X))  ← not item(X), z(X), current(X) .

```

---

Listing 3:  $\Pi_{\text{optForCirc}}$ : used for solving Circumscription via  $\Pi_{\text{CIRC}} = \Pi_{\text{optForCirc}} \cup \Pi_{\text{SAT}}$ .

### 5.3 Disjunctive ASP

While a traditional TD-based DP algorithm for solving disjunctive ASP can be found in [24], here we solve the problem with D-FLAT<sup>2</sup>. We first do so via reduction to Circumscription. In the following, for any interpretation, rule or set of atoms  $X$ , we write  $X'$  to denote the result of replacing each atom  $a$  in  $X$  with a new atom  $a'$ . Given a disjunctive logic program  $\Pi$  consisting of rules of the form  $a_1 \vee a_2 \vee \dots \vee a_k \leftarrow b_1, \dots, b_m, \text{not } b_{m+1}, \dots, \text{not } b_n$ , we use the notation  $h(r) := \bigvee_{1 \leq i \leq k} a_i$ ,  $b^+(r) := \bigwedge_{1 \leq i \leq m} b_i$  and  $b^-(r) := \bigwedge_{m+1 \leq i \leq n} \neg b_i$ .

As shown in [31], an interpretation  $I$  is an answer-set of  $\bar{\Pi}$  iff  $I \cup I'$  is a model of  $\bigwedge_{p \in HB(\Pi)} (p \equiv p') \wedge \text{CIRC}(\bigcup_{r \in \Pi} \{(b^+(r) \wedge b^-(r)) \rightarrow h(r)\}; HB(\Pi); HB(\Pi)')$ , where  $HB(\Pi)$  denotes the Herbrand base of  $\Pi$ . In order to compute models of this formula, we can first calculate models of the Circumscription part and then remove those models (using the “pseudo” item for marking pseudo solution candidates), where the truth value of some atom  $a$  is different from the one of  $a'$ . This amounts to  $\Pi_{\text{pseudoForASP}}$  (Listing 4), which can be used for solving disjunctive ASP by means of the combined program  $\Pi_{\text{ASP}} = \Pi_{\text{pseudoForASP}} \cup \Pi_{\text{CIRC}}$ . The predicate `cor` in this encoding is assumed to be symmetric and occurring in input facts in order to associate each atom  $a$  with its corresponding primed variant  $a'$ . (We require that  $a$  and  $a'$  always occur together in some bag, which can be achieved by adding an edge  $(a, a')$  to the input graph.)

---

```

1 auxItem(pseudo) ← cor(A, B), current(A; B), item(A), not item(B).
2 auxItem(pseudo) ← extend(R), childAuxItem(R, pseudo).

```

---

Listing 4:  $\Pi_{\text{pseudoForASP}}$ : used for solving disjunctive ASP via  $\Pi_{\text{ASP}} = \Pi_{\text{pseudoForASP}} \cup \Pi_{\text{CIRC}}$ .

In Listing 5, we present an alternative approach to solving disjunctive ASP, which does not resort to Circumscription. In this encoding,  $\Pi'_{\text{ASP}}$ , we generate solution candidates for all interpretations that are candidates for being a classical model of the input program, which is specified by means of the predicates `head`, `pos` and `neg`. A classical model  $M$  of a program  $P$  might be no answer set because some  $M' \subset M$  is a model of the reduct  $P^M$ . To check this, we generate additional rows that only serve as counterexample candidates (like  $M'$ ) to the rows representing classical model candidates (like  $M$ ). For any atom  $a$  from the current bag, if an item set contains  $a$ , then the corresponding interpretation sets  $a$  to true (otherwise to false). The rows representing potential counterexamples can additionally contain items of the form  $r(a)$ . This signifies that the atom  $a$  is false in the respective counterexample candidate but true in the classical model candidates that reference this counterexample candidate (by means of their counterexample pointers). In Lines 22 and 24, we make sure that any row containing an item  $r(a)$  is marked with the “pseudo” item and will therefore not be considered as a solution but rather serves as a counterexample candidate only. Lines 27 and 28 are required to ensure that any counterexample candidate  $C$  of any solution candidate  $M$  only contains  $r(a)$  for atoms  $a$  that are also contained in  $M$ .

---

```

1 1 { extend(R) : childRow(R,N) } 1 ← childNode(N) .
2 % Guess truth value/rule flag of introduced atoms
3 0 { item(A;r(A)) : atom(A), introduced(A) } 1.
4 % Make explicit when an atom is false or a rule is unsat
5 false(R,X) ← childRow(R,N), bag(N,X), not childItem(R,X) .
6 falser(R,X) ← childRow(R,N), bag(N,X), not childItem(R,r(X)) .
7 unsat(R,X) ← childRow(R,N), bag(N,X), not childAuxItem(R,X) .
8 % Only join child item sets that coincide on common atoms
9 ← extend(X;Y), atom(A), childItem(X,A), false(Y,A) .
10 ← extend(X;Y), atom(A), childItem(X,r(A)), falser(Y,A) .
11 % Only extend child item sets satisfying all removed rules
12 ← extend(S), rule(X), removed(X), unsat(S,X) .
13 % True atoms and satisfied rules remain so unless removed
14 item(X) ← extend(S), childItem(S,X), current(X) .
15 item(r(X)) ← extend(S), childItem(S,r(X)), current(X) .
16 % Through the guess, rules may become satisfied
17 auxItem(R) ← current(R;A), head(R,A), item(A) .
18 auxItem(R) ← current(R;A), pos(R,A), not item(A) .
19 auxItem(R) ← current(R;A), neg(R,A), item(A) .
20 % Rule is not in reduct if a negative body atom is set to true
21 auxItem(R) ← current(R;A), neg(R,A), item(r(A)) .
22 auxItem(pseudo) ← item(r(X)), current(X) .
23 % Inherit pseudo flag from child nodes
24 auxItem(pseudo) ← extend(R), childAuxItem(R,pseudo) .
25 optItem(S) ← atom(S), item(S) .
26 % Prevents r(S) at level 2 (reduct) if S is not true at level 1
27 optItem(r(S)) ← atom(S), item(S), not auxItem(pseudo) .
28 optItem(r(S)) ← atom(S), item(r(S)) .

```

---

Listing 5:  $\Pi'_{\text{ASP}}$ : D-FLAT<sup>2</sup> encoding for solving disjunctive ASP directly.

## 5.4 Abstract Argumentation

Problems from abstract argumentation [15] are further examples where our approach is reasonable. Given an object  $(A, R)$ , where  $A$  is a set of arguments and  $R \subseteq A \times A$ , we call a set  $S \subseteq A$  *admissible* if (1)  $(a, b) \notin R$  for all  $a, b \in S$  and (2) for each  $s \in S$  and  $r \in A$ ,  $(r, s) \in R$  implies that there is some  $q \in S$  with  $(q, r) \in R$ .  $S$  is *preferred* if it is a subset-maximal admissible set. For any  $C \subseteq A$ , we call  $C^+ = C \cup \{a \mid \exists b \in C \text{ s.t. } (b, a) \in R\}$  the *range* of  $C$ . A set  $S$  is *semi-stable* if it is admissible and for every admissible  $S' \subset S$ ,  $S^+ \not\subseteq S'^+$  holds.

Listing 6 shows an encoding  $\Pi_{\text{admissible}}$  for computing admissible sets. At first glance, it seems to be overcomplicated compared to the algorithm in [16]: For computing admissible sets it actually suffices to guess which introduced arguments are in  $S$ , whereas we guess in  $\Pi_{\text{admissible}}$  which arguments are in the set, attackers or neither. In order for a guessed set to be a solution, every attacker has to be defeated by the time it is removed from the bag. Once it can be determined

that an attacker is defeated, its status changes from “attc” to “def”. This additional complexity allows us to reuse  $\Pi_{\text{admissible}}$  for computing semi-stable sets later on.

---

```

1 1 { extend(R) : childRow(R,N) } 1 ← childNode(N).
2 % Guess whether an element is in, out or attacking (attc)
3 0 { item(in(A)); attc(A) } 1 ← introduced(A).
4 % Join only arguments with compatible status
5 nIn(S,A) ← childRow(S,N), bag(N,A), not childItem(S,in(A)).
6 nDef(S,A) ← childRow(S,N), bag(N,A), not childAuxItem(S,def(A)), not
   childAuxItem(S,attc(A)).
7 ← extend(S1), extend(S2), childItem(S1,in(A)), nIn(S2,A).
8 ← extend(S1), extend(S2), childAuxItem(S1,def(A)), nDef(S2,A).
9 ← extend(S1), extend(S2), childAuxItem(S1,attc(A)), nDef(S2,A).
10 % Inherit arguments that are in, defeated or attackers
11 item(in(A)) ← extend(S), childItem(S,in(A)), current(A).
12 chdef(A) ← extend(S), childAuxItem(S,def(A)), current(A).
13 attc(A) ← extend(S), childAuxItem(S,attc(A)), current(A).
14 % Set defeated arguments
15 auxItem(def(A)) ← current(A;B), att(B,A), item(in(B)).
16 auxItem(def(A)) ← chdef(A).
17 % Still remaining (undefeated) attackers
18 auxItem(attc(A)) ← attc(A), not auxItem(def(A)).
19 % Out-arguments are not allowed to be defeated/attackers
20 out(A) ← not attc(A), not chdef(A), current(A).
21 ← auxItem(def(A)), out(A).
22 ← out(A), current(A), item(in(B)), att(A,B).
23 % Assure that the set is conflict-free
24 ← item(in(A)), item(in(B)), att(A,B).
25 % Remove candidates that leave attackers undefeated
26 ← extend(S), childAuxItem(S,attc(A)), removed(A).

```

---

Listing 6:  $\Pi_{\text{admissible}}$ : D-FLAT encoding for admissible sets.

Using  $\Pi_{\text{admissible}}$ , we can compute preferred sets by simple subset maximization via

$$\Pi_{\text{preferred}} = \Pi_{\text{optAllItems}} \cup \Pi_{\text{admissible}}.$$

Furthermore, it is now also easy to compute semi-stable sets by means of  $\Pi_{\text{optForSemiStable}}$  (Listing 7), using

$$\Pi_{\text{semiStable}} = \Pi_{\text{optForSemiStable}} \cup \Pi_{\text{admissible}}.$$

Although the code  $\Pi_{\text{admissible}}$  could have been simplified for computing preferred sets, here it is indeed required because we need to find those admissible sets that have maximal range.

---

```

1 optItem(A) ← item(in(A)).
2 optItem(A) ← auxItem(attc(A)).
3 optItem(A) ← auxItem(def(A)).

```

---

Listing 7:  $\Pi_{\text{optForSemiStable}}$ : used for computing semi-stable sets.

## 6 Evaluation

In this section we compare our new system D-FLAT<sup>2</sup> (version 1.0.2) with D-FLAT (version 1.0.1). Both frameworks internally use ASP grounder Gringo 4.4.0 and solver Clasp 3.1.1. Additionally, we compare the performance to well-established systems for the respective problem domains.

### 6.1 Abstract Argumentation

This part focuses on experiments for problems from the area of abstract argumentation. In addition to D-FLAT and D-FLAT<sup>2</sup>, we benchmarked the ASPARTIX system [17] that solves argumentation-related problems directly via ASP. The results for ASPARTIX were produced with Gringo 3.0.5, as it is not fully compatible with newer versions of Gringo. For evaluation we used so-called “grid-based” instances, where vertices are arranged on a  $n \times m$  matrix, and edges connect horizontally, vertically and diagonally neighboring vertices. Each instance was run five times with different TDs, and every run was limited to one hour and three GB of memory.

**System comparison.** We considered the problem of enumerating all preferred extensions and compared the systems on grid-based instances with 40 to 65 nodes and treewidth 4. Figure 6 illustrates average runtimes and allocated memory together with the 95 % confidence interval. D-FLAT<sup>2</sup> showed the best performance, while D-FLAT is slightly slower and requires more memory. For ASPARTIX we observed timeouts for instances having more than 55 nodes.

**Problem comparison.** As D-FLAT<sup>2</sup> is based on D-FLAT, we compared these systems on several problems using grid-based instances with treewidth 4. Moreover, we analyzed the cost of computing preferred and semi-stable sets compared to only obtaining admissible sets. As instances have much more admissible sets than preferred sets, which would bias a performance comparison when doing explicit enumeration, we considered the counting variants of these problems. Results are summarized in Figure 7.

When counting admissible sets, D-FLAT<sup>2</sup> requires slightly more time and memory than D-FLAT due to the overhead imposed by using reduced item trees instead of item trees. For preferred sets, the inefficiency of computing redundant counter candidates in D-FLAT becomes evident. On the contrary, in D-FLAT<sup>2</sup> the difference in runtime for counting preferred instead of admissible sets is barely measurable (i.e. within the 95% confidence interval). Here, we observed that subset maximization comes for free for instances of small treewidth. Finally, for semi-stable sets, D-FLAT was not able solve instances with 500 vertices within the given memory limits. One reason is that for this problem many potential counter candidates have to be computed that turn out to be not even admissible. Thus, our two-phased approach of first computing (not necessarily maximal) solutions and then performing maximization obviously pays off in this case.

### 6.2 Subset-minimal SAT

**Comparison to QBF.** We also performed experiments for counting minimal models of randomly generated 3-CNF instances (conjunctive normal form with 3 literals in every clause). Results were compared to depQBF 4 [28], a state-of-the-art QBF solver and bloqper 35 [22] in combination with

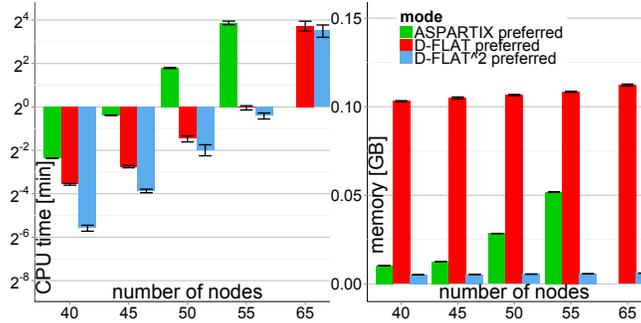


Figure 6: System comparison: Average CPU time (left) and maximum resident set (right).

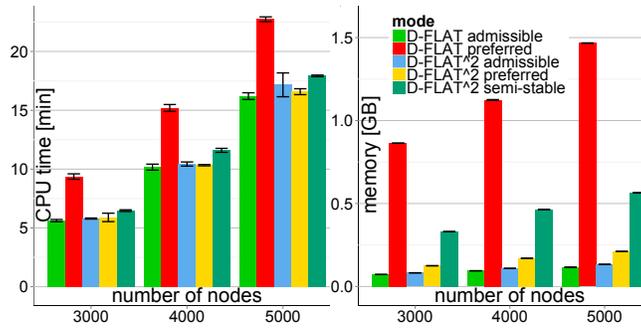


Figure 7: Problem comparison: Average CPU time (left) and maximum resident set (right).

depQBF, a tool for preprocessing QBFs. We considered randomly generated, satisfiable instances with 26 clauses and a varying number of variables. For each number of variables, five instances were generated and run three times. The treewidth of all generated instances did not exceed the value of 9. Figure 8 illustrates the obtained results, including the 95% confidence interval. The number of variables is shown on the x-axis.

Although bloqqer with depQBF is extremely fast, the preprocessing potentially yields simplifications which are not equivalence-preserving. Thus, bloqqer might not be well-suited for some applications. Furthermore, we let depQBF and bloqqer only output the first satisfying assignment (marked with an asterisk in the legend of the plot), since counting solutions is not directly supported by these tools. The actual runtime would, of course, increase in case of counting (via enumeration) since these tools require to incrementally call the solver until the formula (modified by previously computed solutions) is unsatisfiable. Figure 8 also reveals that (for instances of small treewidth) subset optimization in D-FLAT<sup>2</sup> almost comes for free, since SAT and  $\subseteq$ -MINIMAL SAT require approximately the same amount of resources. Note that we are going to conduct further experiments for the most recent version of depQBF released after submission deadline.

**Comparison to ASP.** In addition to before, we also counted minimal models by means of ASP. We used Potassco (Gringo 4.4.0 and Clasp 3.1.1) as a representative candidate among ASP solvers. Here, we again considered satisfiable 3-CNF instances. To guarantee an upper bound of 4 for the treewidth, we generated the instances following a grid-based structure (as above). Each bar in

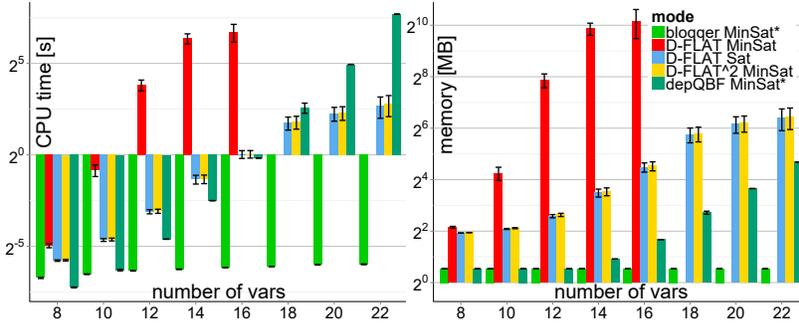


Figure 8: Comparison to QBF: Average CPU time (left) and maximum resident set (right).

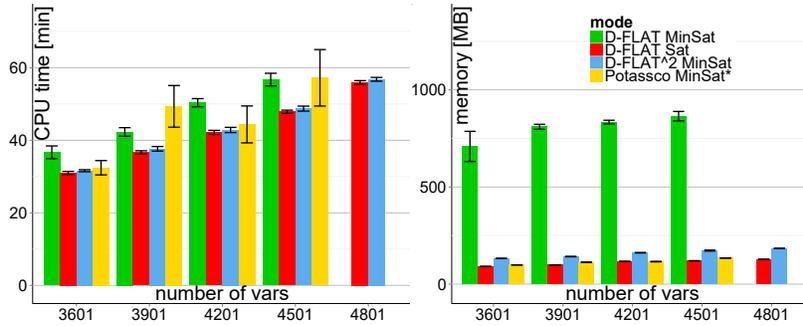


Figure 9: Comparison to ASP: Average CPU time (left) and maximum resident set (right).

Figure 9 presents the average time (or memory) over the five instances per number of variables and three runs per instance assuming a timeout of one hour. Additionally, the 95% confidence interval is given.

Again, Potassco only had to print the first satisfying assignment (marked with an asterisk), i.e. the actual runtime would increase in case of enumeration/counting. Note that the ASP approach is quite fast, but with instances of small treewidth, D-FLAT<sup>2</sup> quickly shows its advantage. Moreover, observe that compared to D-FLAT<sup>2</sup>, Potassco shows potentially increased runtime consumption and a huge runtime deviation.

## 7 Conclusion

To put FPT results for bounded treewidth to use in the AI domain, often DP algorithms for problems involving tasks like subset minimization have to be designed. These algorithms exhibit common properties that are tedious to specify but can be automatically taken care of, as we have shown in this paper. In fact, we have provided a translation that turns a given DP algorithm for computing a set  $S$  of solution candidates (say, models of a formula) into a DP algorithm that computes only the subset-minimal elements of  $S$  (e.g., minimal models). We have shown the translation to be sound and to remain FPT whenever the original DP is. This is indeed superior to a naive way that computes all elements from  $S$  first and then filters out minimal ones in a post-processing step,

which would not yield an FPT algorithm in general.

We presented the D-FLAT<sup>2</sup> system for DP on TDs realizing this idea with further generalizations (e.g., performing minimization only on a given set of atoms). Users of D-FLAT<sup>2</sup> are only required to provide an ASP program that specifies an algorithm for a version of the problem without optimization. Our method then performs the optimization tasks in an automatic and uniform way, thus making the development of such algorithms significantly easier. Preliminary experiments indicate that our new approach brings significant advantages in terms of time and memory compared to previous solutions.

For future work, we would like to investigate the potential of other built-ins for DP algorithms; for instance, checks for connectedness could be treated in a similar way. In the long run, we anticipate a system that facilitates implementing DP algorithms but (in contrast to related systems such as *Sequoia*) keeps the overall design of the concrete DP algorithm in the user's hands.

## References

- [1] Michael Abseher, Bernhard Bliem, Günther Charwat, Frederico Dusberger, Markus Hecher, and Stefan Woltran. D-FLAT: Progress report. Technical Report DBAI-TR-2014-86, TU Wien, 2014.
- [2] Michael Abseher, Bernhard Bliem, Günther Charwat, Frederico Dusberger, Markus Hecher, and Stefan Woltran. The D-FLAT system for dynamic programming on tree decompositions. In *Proc. JELIA*, volume 8761 of *LNCS*, pages 558–572, 2014.
- [3] Stefan Arnborg, Derek G. Corneil, and Andrzej Proskurowski. Complexity of finding embeddings in a k-tree. *SIAM J. Algebraic Discrete Methods*, 8(2):277–284, 1987.
- [4] Bernhard Bliem, Günther Charwat, Markus Hecher, and Stefan Woltran. D-FLAT<sup>2</sup>: Subset minimization in dynamic programming on tree decompositions made easy. In *Proc. ASP-POCP'15*, 2015.
- [5] Hans L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 25(6):1305–1317, 1996.
- [6] Hans L. Bodlaender and Arie M. C. A. Koster. Treewidth computations I. Upper bounds. *Inf. Comput.*, 208(3):259–275, 2010.
- [7] Gerhard Brewka, James P. Delgrande, Javier Romero, and Torsten Schaub. asprin: Customizing answer set preferences without a headache. In *Proc. AAI*, pages 1467–1474. AAAI Press, 2015.
- [8] Gerhard Brewka, Thomas Eiter, and Mirosław Truszczyński. Answer set programming at a glance. *Commun. ACM*, 54(12):92–103, 2011.
- [9] Bruno Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Inf. Comput.*, 85(1):12–75, 1990.

- [10] Bruno Courcelle and Irène Durand. Computations by fly-automata beyond monadic second-order logic. *CoRR*, abs/1305.7120, 2013.
- [11] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshantov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.
- [12] Rina Dechter. *Constraint Processing*. Morgan Kaufmann, 2003.
- [13] Artan Dermaku, Tobias Ganzow, Georg Gottlob, Benjamin J. McMahan, Nysret Musliu, and Marko Samer. Heuristic methods for hypertree decomposition. In *Proc. MICAI*, volume 5317 of *LNCS*, pages 1–11. Springer, 2008.
- [14] Rodney G. Downey and Michael R. Fellows. *Parameterized Complexity*. Monographs in Computer Science. Springer, 1999.
- [15] Phan M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.*, 77(2):321–357, 1995.
- [16] Wolfgang Dvořák, Reinhard Pichler, and Stefan Woltran. Towards fixed-parameter tractable algorithms for abstract argumentation. *Artif. Intell.*, 186:1–37, 2012.
- [17] Uwe Egly, Sarah Alice Gaggl, and Stefan Woltran. Answer-set programming encodings for argumentation frameworks. *Argument and Computation*, 1(2):147–177, 2010.
- [18] Thomas Eiter and Axel Polleres. Towards automated integration of guess and check programs in answer set programming: A meta-interpreter and applications. *TPLP*, 6(1-2):23–60, 2006.
- [19] Martin Gebser, Roland Kaminski, and Torsten Schaub. Complex optimization in answer set programming. *TPLP*, 11(4-5):821–839, 2011.
- [20] Martin Gebser, Benjamin Kaufmann, Roland Kaminski, Max Ostrowski, Torsten Schaub, and Marius Thomas Schneider. Potassco: The potsdam answer set solving collection. *AI Commun.*, 24(2):107–124, 2011.
- [21] Georg Gottlob, Reinhard Pichler, and Fang Wei. Tractable database design and datalog abduction through bounded treewidth. *Inf. Syst.*, 35(3):278–298, 2010.
- [22] Marijn Heule, Martina Seidl, and Armin Biere. A unified proof system for QBF preprocessing. In *Proc. IJCAR*, volume 8562 of *LNCS*, pages 91–106. Springer, 2014.
- [23] Michael Jakl, Reinhard Pichler, Stefan Rümmele, and Stefan Woltran. Fast counting with bounded treewidth. In *Proc. LPAR*, volume 5330 of *LNCS*, pages 436–450. Springer, 2008.
- [24] Michael Jakl, Reinhard Pichler, and Stefan Woltran. Answer-set programming with bounded treewidth. In *Proc. IJCAI*, pages 816–822, 2009.
- [25] Ton Kloks. *Treewidth: Computations and Approximations*, volume 842 of *LNCS*. Springer, 1994.

- [26] Joachim Kneis, Alexander Langer, and Peter Rossmanith. Courcelle’s theorem – a game-theoretic approach. *Discrete Optimization*, 8(4):568–594, 2011.
- [27] Nicola Leone, Gerald Pfeifer, Wolfgang Faber, Thomas Eiter, Georg Gottlob, Simona Perri, and Francesco Scarcello. The DLV system for knowledge representation and reasoning. *ACM Trans. Comput. Log.*, 7(3):499–562, 2006.
- [28] Florian Lonsing and Uwe Egly. Depqbf: An incremental QBF solver based on clause groups. *CoRR*, abs/1502.02484, 2015.
- [29] John McCarthy. Circumscription – a form of non-monotonic reasoning. *Artif. Intell.*, 13(12):27–39, 1980.
- [30] Rolf Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford Lecture Series in Mathematics and its Applications. OUP, 2006.
- [31] David Pearce, Hans Tompits, and Stefan Woltran. Characterising equilibrium logic and nested logic programs: Reductions and complexity. *TPLP*, 9(5):565–616, 2009.
- [32] Neil Robertson and Paul D. Seymour. Graph minors. III. Planar tree-width. *J. Comb. Theory, Ser. B*, 36(1):49–64, 1984.
- [33] Marko Samer and Stefan Szeider. Algorithms for propositional model counting. *J. Discrete Algorithms*, 8(1):50–64, 2010.

## A Appendix: Proofs

**Lemma 1.** *If a table  $R$  is normal, then so is  $\text{compr}(R)$ .*

*Proof.* Let  $R$  be a normal table, let  $R' = \text{compr}(R)$ , and let  $R_1, \dots, R_k$  denote the child tables of  $R$  (and  $R'$ ). As for the first condition of normality of  $R'$ , let  $r' \in R'$ ,  $s' \in S(r')$ ,  $X' \in E(r')$  and  $Y' \in E_{X'}(s')$ . For  $1 \leq i \leq k$ , the set  $X' \cap R_i$  consists of a single element  $r_i$ , as can be seen via Definition 3. Similarly, there is a subrow  $s_i$  such that  $Y' \cap \bigcup_{t \in R_i} S(t) = \{s_i\}$ . As  $Y'$  is an extension relative to  $X'$ ,  $s_i \in S(r_i)$  holds. The fact that  $r_i$  and  $s_i$  are part of extensions of  $r'$  and  $s'$ , respectively, entails that  $(r_1, \dots, r_k) \in P(r')$  and  $(s_1, \dots, s_k) \in P(s')$ . From the definition of  $\text{compr}(\cdot)$  we infer that then there is a row  $r \in R$  having a subrow  $s \in S(r)$  such that  $r \approx r'$ ,  $s \approx s'$  and  $(s_1, \dots, s_k) \in P(s)$ . It must hold by Definition 2 that  $(r_1, \dots, r_k) \in P(r)$  because  $s \in S(r)$ ,  $(s_1, \dots, s_k) \in P(s)$  and  $s_i \in S(r_i)$ . As  $(r_1, \dots, r_k) \in P(r)$  and  $(s_1, \dots, s_k) \in P(s)$ , there are  $X \in E(r)$  and  $Y \in E_X(s)$  with  $X \approx X'$  and  $Y \approx Y'$ . As  $R$  is normal,  $Y \leq X$  holds, which proves that  $Y' \leq X'$ . By definition of  $\text{compr}(\cdot)$ , it holds that  $\text{inc}(s) = \subset$  if and only if  $\text{inc}(s') = \subset$ . Furthermore, due to normality of  $R$ ,  $Y < X$  holds if and only if  $\text{inc}(s) = \subset$ . This proves that  $Y' < X'$  holds if and only if  $\text{inc}(s') = \subset$ .

To prove the second condition, let  $r' \in R'$ ,  $s' \in S(r')$  and  $Y' \in E(s')$ . Using the same reasoning as before, we can see that there is a subrow  $s$  of some row in  $R$  such that  $s \approx s'$ , and there are elements  $s_i \in Y'$  such that  $(s_1, \dots, s_k) \in P(s)$ . Hence there is an extension  $Y \in E(s)$  with  $Y \approx Y'$ . By normality of  $R$ , there is a row  $t \in R$  with  $t \approx s$  and an extension  $T \in E(t)$  with  $T \approx Y$ . Similar to before, then there are elements  $t_i \in T$  such that  $(t_1, \dots, t_k) \in P(t)$ . By definition of  $\text{compr}(\cdot)$ , then there is a row  $t' \in R'$  with  $t' \approx t$  and  $(t_1, \dots, t_k) \in E(t')$ . Furthermore, there are  $T_i \in E(t_i)$  such that  $T = \{t\} \cup T_1 \cup \dots \cup T_k$ . We can construct  $T' = (T \setminus \{t\}) \cup \{t'\}$  and observe that  $T' \in E(t')$  and  $T' \approx T = Y \approx Y'$  as well as  $t' \approx s'$ .

As for the third condition, let  $q', r' \in R'$ ,  $Z' \in E(q')$  and  $X' \in E(r')$  such that  $Z' \leq X'$ . Similar to before, then there are rows  $q, r \in R$  such that  $q \approx q'$  and  $r \approx r'$ , there are elements  $q_i \in Z'$  such that  $(q_1, \dots, q_k) \in P(q)$ , and there are elements  $r_i \in X'$  such that  $(r_1, \dots, r_k) \in P(r)$ . Hence there are  $Z \in E(q)$  and  $X \in E(r)$  with  $Z \approx Z'$  and  $X \approx X'$ , so  $Z \leq X$ . By normality of  $R$ , then there exist  $u \in S(r)$  and  $U \in E_X(u)$  such that  $u \approx q$  and  $U \approx Z$ , and there are  $u_i \in U$  such that  $(u_1, \dots, u_k) \in P(u)$ . Then by definition of  $\text{compr}(\cdot)$ , there is a subrow  $u' \in S(r')$  with  $u' \approx u$  and  $(u_1, \dots, u_k) \in P(u')$ . So  $u' \approx q'$  and there is an extension  $U' \in E_{X'}(u')$  with  $U' \approx U = Z \approx Z'$ , which concludes the proof.  $\square$

Furthermore, we show that the extensions from rows in a table  $R$  are in a one-to-one correspondence to the extensions of rows from  $\text{compr}(R)$  such that the data of corresponding extensions is the same.

**Lemma 4.** *Let  $R$  be a table. Then for any  $r \in R$  and  $X \in E(r)$  there are  $r' \in \text{compr}(R)$  and  $X' \in E(r')$  such that  $r' \approx r$  and  $X' \approx X$ . Also, for any  $r' \in \text{compr}(R)$  and  $X' \in E(r')$  there are  $r \in R$  and  $X \in E(r)$  such that  $r \approx r'$  and  $X \approx X'$ .*

*Proof.* Let  $R$  be a table and  $R' = \text{compr}(R)$ . For the first statement, let  $r \in R$  and  $X \in E(r)$ . Then there are  $(r_1, \dots, r_k) \in P(r)$  and  $X_i$  such that  $X = \{r\} \cup X_1 \cup \dots \cup X_k$ . By definition of

$\text{compr}(\cdot)$ , there is  $r' \in R'$  with  $r' \approx r$  and  $(r_1, \dots, r_k) \in P(r')$ . Then  $X' = \{r'\} \cup X_1 \cup \dots \cup X_k$  is in  $E(r')$  and  $X' \approx X$ . The second statement is proved symmetrically.  $\square$

As explained after our definition of  $\text{aug}(\cdot)$ , each row and subrow in the (compressed) augmentation result has a unique originating row. We formalize this as follows.

**Definition 16.** *Let  $R$  be a table in some augmentable computation, let  $Q = \text{aug}(R)$  and  $Q' = \text{compr}(Q)$ . We define a function  $\text{orig}_R(\cdot)$  that maps each row or subrow in  $Q$  or  $Q'$  to a row in  $R$ . For  $q \in Q \cup Q'$  we define  $\text{orig}_R(q)$  to be the unique  $r \in R$  with  $r \approx q$ . For  $s \in S(q)$  we define  $\text{orig}_R(s)$  to be the unique  $r \in R$  with  $r \approx s$ .*

Let  $R$  be a table from an augmentable computation. The following lemma formalizes that for each (sub)row  $q$  in  $\text{aug}(R)$  with  $(q_1, \dots, q_k) \in P(q)$ , the originating row  $\text{orig}_R(q)$  in  $R$  has a corresponding EPT  $(r_1, \dots, r_k)$  such that each  $q_i$  is resulting from  $r_i$ . This will later be useful to establish that for each extension in  $Q$  we can also obtain one in  $R$  having the same data.

**Lemma 5.** *Let  $R$  be a table from an augmentable computation and let the child tables of  $R$  be called  $R_1, \dots, R_k$ . We have  $(\text{orig}_{R_1}(q_1), \dots, \text{orig}_{R_k}(q_k)) \in P(\text{orig}_R(q))$  for any  $q \in \text{aug}(R)$  and  $(q_1, \dots, q_k) \in P(q)$ . Furthermore, for any  $s \in S(q)$  and  $(s_1, \dots, s_k) \in P(s)$  it holds that  $(\text{orig}_{R_1}(s_1), \dots, \text{orig}_{R_k}(s_k)) \in P(\text{orig}_R(s))$ .*

*Proof.* Let  $R$  be a table in some augmentable computation such that  $R_1, \dots, R_k$  denote the child tables of  $R$ , and let  $Q = \text{aug}(R)$  with child tables  $Q_i = \text{compr}(\text{aug}(R_i))$ . Moreover, let  $q \in Q$  and  $(q_1, \dots, q_k) \in P(q)$ . By construction of  $\text{aug}(R)$ , then there are  $r \in R$  and  $(r_1, \dots, r_k) \in P(r)$  such that  $r \approx q$  and  $q_i \in \text{res}(r_i)$  for each  $1 \leq i \leq k$ . But then  $q_i \approx r_i$ , so  $\text{orig}_{R_i}(q_i) = r_i$  holds. Furthermore,  $\text{orig}_R(q) = r$  holds since  $q \approx r$ . Because of  $(r_1, \dots, r_k) \in P(r)$ , this entails  $(\text{orig}_{R_1}(q_1), \dots, \text{orig}_{R_k}(q_k)) \in P(\text{orig}_R(q))$ .

As for the second statement, let  $s \in S(q)$  and  $(s_1, \dots, s_k) \in P(s)$ . By construction of  $\text{aug}(R)$ , then there are  $q' \in Q$  and  $(q'_1, \dots, q'_k) \in P(q')$  such that  $q' \approx s$  and  $q'_i \approx s_i$  for each  $1 \leq i \leq k$ . We know that this entails  $(\text{orig}_{R_1}(q'_1), \dots, \text{orig}_{R_k}(q'_k)) \in P(\text{orig}_R(q'))$  by the first statement of this lemma. As we have seen that  $q'_i \approx s_i$  and  $q' \approx s$ , it must hold that  $\text{orig}_R(s) = \text{orig}_R(q')$  and  $\text{orig}_{R_i}(s_i) = \text{orig}_{R_i}(q'_i)$ , which proves the second statement.  $\square$

**Lemma 2.** *Let  $R$  be a table from an augmentable computation and  $Q = \text{aug}(R)$ . Then for any  $r \in R$  and  $Z \in E(r)$  there are  $q \in Q$  and  $X \in E(q)$  such that  $r \approx q$  and  $Z \approx X$ . Also, for any  $q \in Q$  and  $X \in E(q)$  there are  $r \in R$  and  $Z \in E(r)$  such that  $q \approx r$  and  $X \approx Z$ .*

*Proof.* Let  $R$  be a table in some augmentable computation such that  $R_1, \dots, R_k$  denote the child tables of  $R$ , and let  $Q = \text{aug}(R)$  with child tables  $Q_i = \text{compr}(\text{aug}(R_i))$ . We prove the lemma by induction. First suppose  $R$  and  $Q$  are leaf tables. Then it can be easily verified using the definition of  $Q$  that the rows in  $Q$  and  $R$  are in a one-to-one correspondence. Formally,  $|Q| = |R|$  holds and for each  $r \in R$  there is a  $q \in Q$  with  $q \approx r$ . As the EPTs in  $R$  and  $Q$  always consist of just the empty tuple, the data of any extension of a row in  $R$  or  $Q$  coincides with the data of that row.

Now suppose  $R$  and  $Q$  are arbitrary tables and both statements from the lemma hold for all  $R_i$ . We first attend to the first statement. Let  $r \in R$  and  $Z \in E(r)$ . Then there are  $(r_1, \dots, r_k) \in P(r)$

and  $Z_i \in E(r_i)$  such that  $Z = \{r\} \cup Z_1 \cup \dots \cup Z_k$ . By the induction hypothesis, there are  $q_i \in Q_i$  and  $X_i \in E(q_i)$  such that  $r_i \approx q_i$  and  $X_i \approx Z_i$ . As  $r_i \approx q_i$ , it holds that  $q_i \in \text{res}(r_i)$ . Hence condition 1 from the construction of  $Q$  ensures that there is a row  $q \in Q$  with  $q \approx r$  and  $P(q) = \{(q_1, \dots, q_k)\}$ . Now clearly  $X = \{q\} \cup X_1 \cup \dots \cup X_k$  is contained in  $E(q)$ , and  $X \approx Z$ .

As for the second statement, let  $q \in Q$  and  $X \in E(q)$ . Then there are  $(q_1, \dots, q_k) \in P(q)$  and  $X_i \in E(q_i)$  such that  $X = \{q\} \cup X_1 \cup \dots \cup X_k$ . By the induction hypothesis, there are  $r_i \in R_i$  and  $Z_i \in E(r_i)$  such that  $q_i \approx r_i$  and  $X_i \approx Z_i$ . Then  $\text{orig}_{R_i}(q_i) = r_i$  holds because of  $q_i \approx r_i$ . Let  $r = \text{orig}_R(q)$ , so obviously  $q \approx r$ . By Lemma 5,  $(q_1, \dots, q_k) \in P(q)$  entails  $(\text{orig}_{R_1}(q_1), \dots, \text{orig}_{R_k}(q_k)) \in P(\text{orig}_R(q))$ , i.e.,  $(r_1, \dots, r_k) \in P(r)$ . Now clearly  $Z = \{r\} \cup Z_1 \cup \dots \cup Z_k$  is contained in  $E(r)$ , and  $X \approx Z$ .  $\square$

For subrows we show a result similar to the second statement from Lemma 2.

**Lemma 6.** *Let  $R$  be a table from an augmentable computation and  $Q = \text{aug}(R)$ . Then for any  $q \in Q$ ,  $s \in S(q)$  and  $Y \in E(s)$  there are  $r \in R$  and  $Z \in E(r)$  such that  $s \approx r$  and  $Y \approx Z$ .*

*Proof.* Let  $R$  be a table in some augmentable computation such that  $R_1, \dots, R_k$  denote the child tables of  $R$ , and let  $Q = \text{aug}(R)$  with child tables  $Q_i = \text{compr}(\text{aug}(R_i))$ . We prove the lemma by induction. First suppose  $R$  and  $Q$  are leaf tables. By definition of  $Q$ , for each subrow  $s$  in  $Q$  there is some  $q' \in Q$  with  $s \approx q'$ . Obviously,  $q' \approx \text{orig}_R(q')$  and  $\text{orig}_R(s) = \text{orig}_R(q')$ , so  $s \approx \text{orig}_R(s)$ , which proves the base case.

Now suppose  $R$  and  $Q$  are arbitrary tables and the statement holds for all  $R_i$ . Let  $q \in Q$ ,  $s \in S(q)$ ,  $Y \in E(s)$  and  $r = \text{orig}_R(s)$ . Then there are  $(s_1, \dots, s_k) \in P(s)$  and  $Y_i \in E(s_i)$  such that  $Y = \{s\} \cup Y_1 \cup \dots \cup Y_k$ . By the induction hypothesis, there are  $r_i \in R_i$  and  $Z_i \in E(r_i)$  such that  $s_i \approx r_i$  and  $Y_i \approx Z_i$ . Then  $\text{orig}_{R_i}(s_i) = r_i$ . By Lemma 5,  $(s_1, \dots, s_k) \in P(s)$  entails  $(\text{orig}_{R_1}(s_1), \dots, \text{orig}_{R_k}(s_k)) \in P(\text{orig}_R(s))$ , so  $(r_1, \dots, r_k) \in P(r)$ . Now clearly  $Z = \{r\} \cup Z_1 \cup \dots \cup Z_k$  is contained in  $E(r)$ , and  $Y \approx Z$ .  $\square$

**Lemma 3.** *Let  $R$  be a table from an augmentable computation. Then the table  $\text{aug}(R)$  is normal.*

*Proof.* Let  $R$  be a table in some augmentable computation such that  $R_1, \dots, R_k$  denote the child tables of  $R$  and let  $Q = \text{aug}(R)$ . We inductively show that  $Q$  is normal if each  $\text{aug}(R_i)$  is normal.

Suppose  $R$  and  $Q$  are leaf tables. Then the EPTs of each row in  $R$  consist of just the empty tuple, which makes the definition of  $Q$  in this case equivalent to the following: For any  $r \in R$ , there is a row  $q \in Q$  with  $q \approx r$  and  $P(q) = \{()\}$ . Furthermore, for each  $q, q' \in Q$  with  $q' \leq q$ , there is a subrow  $s \in S(q)$  with  $s \approx q'$  and  $P(s) = \{()\}$ . Moreover,  $\text{inc}(s) = \subset$  if  $q' < q$ , otherwise  $\text{inc}(s) = \text{eq}$ . It can now easily be verified using Definition 6 that  $Q$  is normal.

Now suppose  $R$  and  $Q$  are arbitrary tables and all  $\text{aug}(R_i)$  are normal. Let also  $Q_i = \text{compr}(\text{aug}(R_i))$ . By Lemma 1, all  $Q_i$  are normal. Let  $q \in Q$ ,  $s \in S(q)$ ,  $X \in E(q)$  and  $Y \in E_X(s)$  be arbitrary, and let  $P(q) = \{(q_1, \dots, q_k)\}$  and  $P(s) = \{(s_1, \dots, s_k)\}$ . It follows that there are extensions  $X_i \in E(q_i)$  such that  $X = \{q\} \cup X_1 \cup \dots \cup X_k$ , and there are extensions  $Y_i \in E_{X_i}(s_i)$  such that  $Y = \{s\} \cup Y_1 \cup \dots \cup Y_k$ . To introduce auxiliary terminology, for any  $x \in Q$  and  $y \in S(q)$ , we write  $\bar{x}$  to denote  $\text{orig}_R(x)$  and  $\bar{y}$  for  $\text{orig}_R(y)$ . Furthermore, for any  $x \in Q_i$  and  $y \in S(x)$ , we write  $\bar{x}$  to denote  $\text{orig}_{R_i}(x)$  and  $\bar{y}$  to denote  $\text{orig}_{R_i}(y)$ .

To prove the first part of condition 1 of normality for  $Q$ , recall that property 2 from the definition of  $Q$  states that  $s \approx q'$  for some  $q' \in Q$  with  $q' \leq q$ . Hence  $s \leq q$ . By normality of  $Q_i$  it holds that  $Y_i \leq X_i$ . This proves  $Y \leq X$ .

For the second part of condition 1 of normality for  $Q$ , first suppose  $Y < X$  while  $\text{inc}(s) = \text{eq}$ . The reason for the latter is that  $\text{inc}(s_i) = \text{eq}$  for all  $s_i$ . For all  $1 \leq i \leq k$ , this entails that  $Y_i \approx X_i$  due to the normality of  $Q_i$ . So  $Y < X$  must be caused by  $s < q$ , but this would entail  $\text{inc}(s) = \subset$ , which is a contradiction.

On the other hand, suppose  $Y \approx X$  while  $\text{inc}(s) = \subset$ . If there is some  $d \in D(q) \setminus D(s)$ , then for  $Y \approx X$  to hold also  $d \in D(Y_j)$  for some  $1 \leq j \leq k$ . But then  $d$  has been illegally removed at  $\bar{s}$ , which contradicts  $R$  being augmentable. More precisely,  $d \in D(\bar{q}) \setminus D(\bar{s})$  holds, and due to  $Y \in E(s)$  and Lemma 6 there is some  $Z \in E(\bar{s})$  such that  $d \in D(Z)$ .

So  $q \approx s$  and for some  $1 \leq j \leq k$  there is a  $d \in D(X_j) \setminus D(Y_j)$ . For  $Y \approx X$  to hold,  $d \in D(s)$  or there is some  $1 \leq h \leq k$  such that  $d \in D(Y_h)$ .

In case  $d \in D(s)$ ,  $d$  has been illegally introduced at  $\bar{s}$ , which contradicts  $R$  being augmentable. More precisely,  $d \in D(\bar{s})$ ,  $(\bar{s}_1, \dots, \bar{s}_k) \in P(\bar{s})$  (by Lemma 5) and  $d \notin D(\bar{s}_j)$  hold while  $d \in D^*(R_j)$  due to  $d \in D(X_j)$ , which amounts to  $d \in D(Z)$  for some  $Z \in E(\bar{q}_j)$ .

In case  $d \in D(Y_h)$  for some  $1 \leq h \leq k$ , we have  $h \neq j$  since  $d \notin D(Y_j)$ . Then also  $d \in D(X_h)$  because  $Y_h \leq X_h$  by normality of  $Q_h$ . Then, by Lemma 2, there are  $H \in E(\bar{q}_h)$  and  $J \in E(\bar{q}_j)$  such that  $d \in D(H)$  and  $d \in D(J)$ . Because  $R$  is augmentable, also  $d \in D(\bar{q})$  and thus  $d \in D(q)$ . As  $q \approx s$ , also  $d \in D(s)$ . But then  $d$  has been illegally introduced at  $\bar{s}$ , as we have already seen. This concludes the proof of the second part of condition 1 of normality for  $Q$ .

Now we show condition 2 of normality for  $Q$ . The reason for  $s \in S(q)$  is that there is some  $q' \in Q$  with  $P(q') = \{(q'_1, \dots, q'_k)\}$  and, for all  $1 \leq i \leq k$ , there is some  $s_i \in S(q_i)$  such that  $s_i \approx q'_i$ . Let  $r = \text{orig}_R(q')$  and  $r_i = \text{orig}_{R_i}(q'_i)$  and  $Z \in E(s)$  be arbitrary. Then there are extensions  $Z_i \in E(s_i)$  such that  $Z = \{s\} \cup Z_1 \cup \dots \cup Z_k$ . By the induction hypothesis, for any  $1 \leq i \leq k$  there is a row  $q''_i \in Q_i$  and an extension  $X''_i \in E(q''_i)$  such that  $q''_i \approx s_i \approx q'_i$  and  $X''_i \approx Z_i$ . As  $q'_i$  and  $q''_i$  have the same data and augmentable tables have at most one row having a certain data, also  $(q''_1, \dots, q''_k) \in \text{res}(r_1) \times \dots \times \text{res}(r_k)$  holds. By construction of  $Q$ , this entails that there is some  $q'' \in Q$  with  $P(q'') = \{(q''_1, \dots, q''_k)\}$  and  $q'' \approx r = q' \approx s$ . We can construct  $X'' = \{q''\} \cup X''_1 \cup \dots \cup X''_k$  and observe that  $X'' \in E(q'')$ . From  $Z_i \approx X''_i$ , for all  $1 \leq i \leq k$ , and  $s \approx q''$  it follows that  $Z \approx X''$ , which proves condition 2 of normality for  $Q$ .

Finally we show condition 3 of normality for  $Q$ . Suppose there are  $q, q' \in Q$  with  $P(q) = \{(q_1, \dots, q_k)\}$  and  $P(q') = \{(q'_1, \dots, q'_k)\}$ , and there are  $X \in E(q)$  and  $X' \in E(q')$  such that  $X' \leq X$ .

Now for all  $1 \leq i \leq k$  let  $X_i \in E(q_i)$  and  $X'_i \in E(q'_i)$ . We first show that  $X'_i \leq X_i$ . Suppose to the contrary that for some  $1 \leq j \leq k$  there is a  $d \in D(X'_j) \setminus D(X_j)$ . For  $X' \leq X$  to hold,  $d \in D(q)$  or there is some  $1 \leq h \leq k$  such that  $d \in D(X_h)$ .

In case  $d \in D(q)$ ,  $d$  has been illegally introduced at  $\bar{q}$ , which contradicts  $R$  being augmentable. More precisely,  $d \in D(\bar{q})$ ,  $(\bar{q}_1, \dots, \bar{q}_k) \in P(\bar{q})$  and  $d \notin D(\bar{q}_j)$  hold while  $d \in D^*(R_j)$  due to  $d \in D(X'_j)$ , which amounts to  $d \in D(Z)$  for some  $Z \in E(\text{orig}_{R_j}(q'_j))$ .

In case  $d \in D(X_h)$  for some  $1 \leq h \leq k$ , we have  $h \neq j$  since  $d \notin D(X_j)$ . Then there are  $H \in E(\bar{q}_h)$  and  $J \in E(\bar{q}_j)$  such that  $d \in D(H)$  and  $d \in D(J)$ . Because  $R$  is augmentable, also

$d \in D(\bar{q})$  and thus  $d \in D(q)$ . But then  $d$  has been illegally introduced at  $\bar{q}$ , as we have already seen.

Hence  $X'_i \leq X_i$  for any  $1 \leq i \leq k$ . By the induction hypothesis, this entails that there are  $s_i \in S(q_i)$  and  $Y_i \in E_{X_i}(s_i)$  with  $s_i \approx q'_i$  and  $Y_i \approx X'_i$ . By construction of  $Q$ , then there is a subrow  $s \in S(q)$  with  $s \approx q'$  and  $P(s) = \{(s_1, \dots, s_k)\}$ . We can now construct  $Y = \{s\} \cup Y_1 \cup \dots \cup Y_k$  and observe that  $Y \in E_X(s)$  and  $Y \approx X'$ , which concludes the proof.  $\square$