

The D-FLAT System for Dynamic Programming on Tree Decompositions

Michael Abseher, Bernhard Bliem, Günther Charwat,
Frederico Dusberger, Markus Hecher, and Stefan Woltran

Institute of Information Systems 184/2
Vienna University of Technology
Favoritenstrasse 9–11, 1040 Vienna, Austria
{abseher,bliem,gcharwat,dusberg,hecher,woltran}@dbai.tuwien.ac.at

Abstract. Complex reasoning problems over large amounts of data pose a great challenge for computer science. To overcome the obstacle of high computational complexity, exploiting structure by means of tree decompositions has proved to be effective in many cases. However, the implementation of suitable efficient algorithms is often tedious. D-FLAT is a software system that combines the logic programming language Answer Set Programming with problem solving on tree decompositions and can serve as a rapid prototyping tool for such algorithms. Since we initially proposed D-FLAT, we have made major changes to the system, improving its range of applicability and its usability. In this paper, we present the system resulting from these efforts.

Keywords: Answer Set Programming, tree decompositions, treewidth.

1 Introduction

Complex reasoning problems over large amounts of data arise in many application domains for computer science and pose a great challenge to push methods from Artificial Intelligence toward practical use. For formalizing complex problems, *declarative approaches* often lead to readable and maintainable code. A more and more popular candidate for such an approach is Answer Set Programming (ASP) [11,17], for which highly efficient solvers are available that offer rich languages for modeling the problems at hand.

In order to overcome performance problems on large problem instances, structural features need to be exploited. A prominent means for this endeavor is the concept of tree decompositions (see, e.g., [9]). Many problems can be efficiently solved with dynamic programming (DP) algorithms on tree decompositions if the structural parameter “treewidth” is bounded, which means, roughly, that the graph resembles a tree to a certain extent. With such an approach, the runtime explosion can be confined to only this parameter instead of input size. Consequently, if the treewidth is bounded, even huge instances of many complex problems can efficiently be solved.

We focus here on a *combination of ASP and problem solving via DP on tree decompositions*. For this, we have implemented a free software system called D-FLAT¹ [4], for rapid prototyping of DP algorithms in ASP. Since ASP is well suited for a lot of problems, it is often also well suited for parts of such problems, making it an appealing candidate to work on decomposed problem instances. The key feature of D-FLAT is that the user is only required to write an encoding of the DP algorithm on a tree decomposition in the ASP language, and the system takes care of tedious tasks that are not related to the problem.

The initial prototype of D-FLAT [4] stored partial solutions in tables. It became clear, however, that for problems higher in the polynomial hierarchy than NP a more general data structure is required. We have shown in [5] that using a tree-shaped data structure instead greatly increases applicability.

In this paper we present the D-FLAT system resulting from the major changes since its initial presentation in [4]. Our main contributions are:

1. We introduce *item trees* as the central data structure in D-FLAT algorithms.
2. We show how item trees allow problems to be solved in the style of Alternating Turing Machines while also taking decomposition into account.
3. We present the special predicates used for communication between the system and the user's encoding.
4. Finally, we show how the system interprets the answer sets of the user's program for constructing item trees and eventually solving the problem.

This work is structured as follows. In Section 2, we provide background on Answer Set Programming, tree decompositions and the original prototype of D-FLAT. In Section 3, we present the current version of D-FLAT and describe its components in detail. Finally, we give a conclusion and an outlook in Section 4.

An extended version of this paper can be found in the technical report [1].

2 Background

Answer Set Programming Answer Set Programming (ASP) is a declarative language where a *program* Π is a set of *rules*

$$a_1 \vee \dots \vee a_k \leftarrow b_1, \dots, b_m, \text{not } b_{m+1}, \dots, \text{not } b_n.$$

The constituents of a rule $r \in \Pi$ are $h(r) = \{a_1, \dots, a_k\}$, $b^+(r) = \{b_1, \dots, b_m\}$ and $b^-(r) = \{b_{m+1}, \dots, b_n\}$. We call r a *fact* if $b^+(r) = b^-(r) = \emptyset$, and we omit the \leftarrow symbol in this case. Intuitively, a rule r states that if an answer set contains all of $b^+(r)$ and none of $b^-(r)$, then it contains some element of $h(r)$. A set of atoms I satisfies a rule r if $I \cap h(r) \neq \emptyset$ or $b^-(r) \cap I \neq \emptyset$ or $b^+(r) \setminus I \neq \emptyset$. I is a *model* of a set of rules if it satisfies each rule. I is an *answer set* of a program Π if it is a subset-minimal model of the program $\Pi^I = \{h(r) \leftarrow b^+(r) \mid r \in \Pi, b^-(r) \cap I = \emptyset\}$ [18].

¹ <http://dbai.tuwien.ac.at/research/project/dflat/> (including download link)

ASP programs can be viewed as succinctly representing problem solving specifications following the *Guess & Check* principle. A “guess” can, for example, be performed using disjunctive rules which non-deterministically open up the search space. Constraints (i.e., rules r with $h(r) = \emptyset$), on the other hand, amount to a “check” by imposing restrictions that solutions must obey.

In this paper, we use the language of the grounder *Gringo* [15,16] (version 4) where programs may contain variables that are instantiated by ground terms (elements of the Herbrand universe, i.e., constants and compound terms containing function symbols) before a solver computes answer sets according to the propositional semantics stated above.

Example 1. The following program solves the 3-COLORABILITY problem for graphs that are given as facts using the predicates `vertex` and `edge`.

```
color(red;grn;blu).
1 { map(X,C) : color(C) } 1 ← vertex(X).
← edge(X,Y), map(X,C;Y,C).
```

Informally, the first rule is shorthand for the three facts `color(red)`, `color(grn)` and `color(blue)`. The second rule states that any vertex from the input graph shall be mapped to one color. The colon controls the instantiation of the variable C such that it is only instantiated with arguments of the predicate `color`. The third rule checks that adjacent vertices never receive the same color. In this rule, `map(X,C;Y,C)` stands for `map(X,C)`, `map(Y,C)`.

Tree Decompositions Tree decompositions and treewidth, originally defined in [21], are well-known tools for tackling computationally hard problems. Informally, treewidth is a measure of the cyclicity of a graph, and many NP-hard problems become tractable if the treewidth is bounded. There are several overviews of this topic, such as [8,6,3,20]. The intuition behind tree decompositions is obtaining a tree from a (potentially cyclic) graph by subsuming multiple vertices under one node and thereby isolating the parts responsible for the cyclicity. The definition of tree decompositions can also be extended to hypergraphs, but in this paper we will only consider graphs for the sake of presentation.

Definition 1. A tree decomposition of a graph $G = (V, E)$ is a pair (T, χ) 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 vertex $v \in V$, 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.

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

Figure 1 shows a graph together with a tree decomposition of it that has width 2. This decomposition is optimal because the graph contains a cycle and thus its treewidth is at least 2. (A graph is a tree iff it has treewidth 1.)

Constructing an optimal tree decomposition is intractable [2]. However, when considering treewidth as a parameter, the problem is fixed-parameter tractable (FPT) [14], i.e., solvable in time $\mathcal{O}(f(w) \cdot n^c)$, where c is a constant, n is the size of the input, w is its treewidth and $f(w)$ depends only on w [7]. Moreover, there are efficient heuristics that produce good tree decompositions [10,13,19].

Tree decompositions are prominently used for solving problems with dynamic programming algorithms. These algorithms generally start at the leaf nodes and traverse the tree decomposition to the root. At each node, partial solutions for the subgraph induced by the vertices encountered so far are computed and stored in a data structure corresponding to that tree decomposition node. Typically, the size of each such data structure only depends on the width of the tree decomposition, and the number of tree decomposition nodes is linear in the size of the input graph. Thus, when the width is bounded by a constant, the search space for each subproblem remains constant as well, and the number of subproblems only grows by a linear factor for larger instances.

Example 2. To solve the 3-COLORABILITY problem with a bottom-up traversal of the tree decomposition in Figure 1, we can first compute all proper 3-colorings of the subgraphs induced by $\{a, b, c\}$ and $\{c, e\}$, respectively. In the node with bag $\{b, c, d\}$ we proceed similarly but now only look for colorings (i) which are consistent combinations of some coloring from each child node, and (ii) where the “new” vertex d is colored consistently with respect to its neighbors (b and c). Here we can disregard the colors of a and e because condition 3 from Definition 1 guarantees that these vertices will never appear again in the bottom-up traversal.

The D-FLAT System D-FLAT [1,4] is a framework for solving computational problems by dynamic programming on a tree decomposition of the input. It proceeds as follows.

1. D-FLAT parses a representation of the problem instance and automatically constructs a tree decomposition of it using heuristic methods which are proposed in [13]. For details, we refer to [1].
2. It provides a data structure suitable for storing partial solutions for many problems. The programmer only needs to specify (using ASP) how to populate the data structure associated with a tree decomposition node.
3. D-FLAT traverses the tree decomposition in post-order and calls an ASP system at each tree decomposition node for computing the data structure corresponding to that node by means of the user-specified program.
4. The framework automatically combines the partial solutions and prints all complete solutions. Alternatively, it is also possible to solve decision, counting and optimization problems.

In our presentation of the initial D-FLAT prototype [4] we were able to successfully apply it to several problems, and we showed in [5] which modifications could further extend its applicability. In the current paper we present the new version of D-FLAT that results from these extensions.

3 The Extended D-FLAT System

This section gives an overview of D-FLAT with the emphasis on the extensions made since [4] for solving any problem expressible in monadic second-order logic in FPT time [5]. This includes many problems in PSPACE.

3.1 Item Trees

D-FLAT equips each tree decomposition node with an *item tree*. This is the data structure that shall contain information about (candidates for) partial solutions. At each node during D-FLAT's bottom-up traversal of the tree decomposition, the problem-specific algorithm can store data in the item tree of that node.

Each node in an item tree contains an *item set*. The elements of this set, called *items*, are arbitrary ground ASP terms. An item tree node also contains additional information about the item set as well as data required for putting together complete solutions, which will be described later in this section.

Item trees are similar to computation trees of Alternating Turing Machines (ATMs) [12]. Like in ATMs, a branch can be seen as a computation sequence, and branching amounts to non-deterministic guesses. We will repeatedly come back to the ATM analogy in the course of this section.

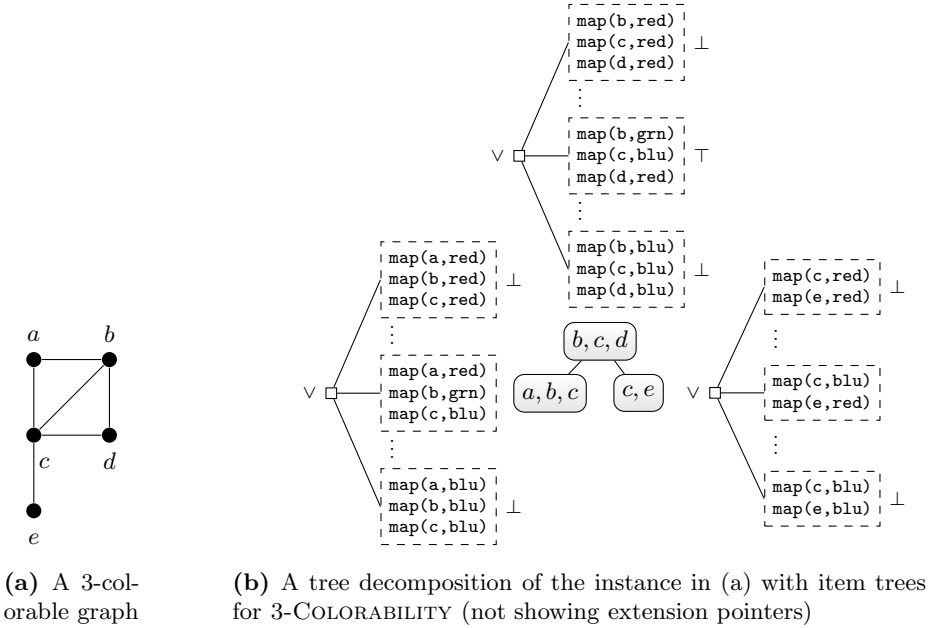
Usually we want to restrict the information within an item tree to information about the current decomposition node's bag elements. The reason is that when the maximum size of an item tree only depends on the bag size, and if the decomposition width is bounded by a constant, the size of each item tree is also bounded. This allows us to achieve FPT algorithms.

Each branch in an item tree may be associated with a cost value, which allows for optimization problems to be solved. If costs are given, D-FLAT automatically only reports optimal solutions. Details on this are found in [1].

Example 3. Figure 1 shows a graph, one of its tree decompositions and, for each decomposition node, the corresponding item tree that could result from an algorithm for 3-COLORABILITY. Each item tree node at depth 1 encodes a coloring of the vertices in the respective bag. The meaning of the symbols \vee , \top and \perp will be explained later in this section.

Extension Pointers. For solving a complete problem instance, it is usually necessary to combine information from different item trees. For example, in order to find out if a proper coloring of a graph exists, we not only have to check if a proper coloring of each subgraph induced by a bag exists but also if, for each bag, we can pick a local coloring in such a way that each vertex is never colored differently by two chosen local colorings.

For this reason each item tree node has a (non-empty) set of *extension pointer tuples*. The elements of such a tuple are called *extension pointers* and reference item tree nodes from children of the respective decomposition node. Roughly, an extension pointer specifies that the information in the source and target nodes can reasonably be combined.



(a) A 3-colorable graph (b) A tree decomposition of the instance in (a) with item trees for 3-COLORABILITY (not showing extension pointers)

Fig. 1. Item trees for a decomposition of a 3-COLORABILITY instance

Example 4. Consider Figure 1b again. In the following examples, let \mathcal{I}_S denote the item tree of the node whose bag is S . In $\mathcal{I}_{\{a,b,c\}}$ and $\mathcal{I}_{\{c,e\}}$, all nodes have the same set of extension pointer tuples: the set consisting of the empty tuple, as those decomposition nodes have no children. The set of extension pointer tuples at the root of $\mathcal{I}_{\{b,c,d\}}$ consists of a single binary tuple – one element references the root of $\mathcal{I}_{\{a,b,c\}}$, the other references the root of $\mathcal{I}_{\{c,e\}}$. For a node ν at depth 1 of $\mathcal{I}_{\{b,c,d\}}$, the set of extension pointer tuples consists of all tuples (ν_1, ν_2) such that ν_1 and ν_2 are nodes at depth 1 of $\mathcal{I}_{\{a,b,c\}}$ and $\mathcal{I}_{\{c,e\}}$, respectively. Moreover, if an element of the current bag $\{b, c, d\}$ is assigned a color in ν_1 or ν_2 , then ν colors it in the same way.

Item Tree Node Types. Like states of ATMs, item tree nodes in D-FLAT can have one of the four types: “or”, “and”, “accept” or “reject”. Unlike ATMs, however, the mapping in D-FLAT is partial, i.e., a node’s type may be undefined. The problem-specific algorithm determines which item tree node is mapped to which type. The following conditions must be fulfilled for an item tree \mathcal{I} .

- If a non-leaf node of \mathcal{I} has been mapped to a type, it is either “or” or “and”.
- If a leaf of \mathcal{I} has been mapped to a type, it is either “accept” or “reject”.
- If a node of \mathcal{I} extends a node with defined type, it must be mapped to the same type.

When D-FLAT has finished processing all answer sets and has constructed the item tree for the current tree decomposition node, it propagates information about the acceptance status of nodes upward in this item tree depending on the node types. These types also play a role when solving optimization problems – roughly, when something is an “or” node, we would like to find a child with minimum cost, and if something is an “and” node, we would like to find a child with maximum cost. This is described in Section 3.3.

Example 5. The item trees in Figure 1b all have roots of type “or”, denoted by the symbol \vee . This is because an ATM for deciding graph colorability starts in an “or” state, then guesses a coloring and accepts if this coloring is proper. Therefore, we shall derive the type “reject” in our decomposition-based algorithm whenever we determine that a guessed coloring is not proper, and we derive “accept” once we are sure that a coloring is proper. In $\mathcal{I}_{\{a,b,c\}}$ and $\mathcal{I}_{\{c,e\}}$, for instance, we have marked all leaves representing an improper coloring with \perp . The types of the other leaves are left undefined, as guesses on vertices that only appear later could still lead to an improper coloring. At the root of the tree decomposition however, we mark all item tree leaves having a yet undefined type with \top because all vertices have been encountered.

3.2 D-FLAT’s Interface for ASP

D-FLAT invokes an ASP solver at each node during a bottom-up traversal of the tree decomposition. The user-defined, problem-specific encoding is augmented with input facts describing the current bag as well as the bags and item trees of child nodes. Additionally, the original problem instance is supplied as input. The answer sets of this ASP call specify the item tree that D-FLAT shall construct for the current decomposition node. D-FLAT provides facts about the tree decomposition and child item trees according to Table 1. We have omitted some less frequently used predicates for clarity. A complete list is given in [1].

Each answer set corresponds to a branch in the new item tree. The predicates for specifying this branch are described in Table 2. One should keep in mind, however, that D-FLAT may merge subtrees as described in Section 3.3. Therefore, after merging, one branch in the item tree may comprise information from multiple answer sets.

Example 6. A possible encoding for the 3-COLORABILITY problem is shown in Listing 1.1. We use colors to highlight **input** and **output** predicates. Note that it would be more convenient (and faster) to encode this problem using the simplified ASP interface for problems in NP described in [1], which we omit from this paper for space reasons.

Line 1 specifies that each answer set declares a branch of length 1, whose root node has the type “or”. Line 2 guesses a color for each current vertex. The “reject” node type is derived in line 3 if this guessed coloring is improper. Lines 4 and 5 guess a branch for each child item tree. Due to line 6, the guessed combination of predecessor branches only leads to an answer set if it does not

Table 1. The most commonly used input predicates describing the tree decomposition and item trees of child nodes in the decomposition

Input predicate	Meaning
<code>final</code>	The current tree decomposition node is the root.
<code>current(V)</code>	V is an element of the current bag.
<code>introduced(V)</code>	V is a current vertex but was in no child node's bag.
<code>removed(V)</code>	V was in a child node's bag but is not in the current one.
<code>root(S)</code>	S is the root of an item tree from a child of the current node.
<code>sub(R, S)</code>	R is an item tree node with child S .
<code>childItem(S, I)</code>	The item set of item tree node S contains I .
<code>childCost(S, C)</code>	C is the cost value corresponding to the item tree leaf S .
<code>childOr(S)</code>	The type of the item tree node S is "or".
<code>childAnd(S)</code>	The type of the item tree node S is "and".
<code>childAccept(S)</code>	The type of the item tree leaf S is "accept".
<code>childReject(S)</code>	The type of the item tree leaf S is "reject".

Table 2. The most commonly used output predicates for constructing the item tree of the current decomposition node

Output predicate	Meaning
<code>item(L, I)</code>	The item set of the node at level L of the current branch shall contain the item I .
<code>extend(L, S)</code>	The node at level L of the current branch shall extend the child item tree node S .
<code>cost(C)</code>	The leaf of the current branch shall have a cost value of C .
<code>length(L)</code>	The current branch shall have length L .
<code>or(L)</code>	The node at level L of the current branch shall have type "or".
<code>and(L)</code>	The node at level L of the current branch shall have type "and".
<code>accept</code>	The leaf of the current branch shall have type "accept".
<code>reject</code>	The leaf of the current branch shall have type "reject".

```

1 length(1). or(0).
2 1 { item(1, map(X, red; X, grn; X, blu)) } 1 ← current(X).
3 reject ← edge(X, Y), item(1, map(X, C; Y, C)).
4 extend(0, S) ← root(S).
5 1 { extend(1, S) : sub(R, S) } 1 ← root(R).
6 ← item(1, map(X, C1)), childItem(S, map(X, C2)),
   extend(_, S), C1 ≠ C2.
7 reject ← childReject(S), extend(_, S).
8 accept ← final, not reject.

```

Listing 1.1. D-FLAT encoding for 3-COLORABILITY

```

1  length(2). or(0). and(1).
2  1 { item(L,map(X,red;X,grn;X,blu)) } 1 ← current(X),
   L = 1..2.
3  ← edge(X,Y), item(L,map(X,C;Y,C)).
4  ← item(2,map(V,red)), not item(1,map(V,red)).
5  extend(0,S) ← root(S).
6  1 { extend(L+1,S) : sub(R,S) } 1 ← extend(L,R),
   L = 0..1.
7  ← item(L,map(X,C1)), childItem(S,map(X,C2)),
   extend(L,S), C1 ≠ C2.
8  item(2,fail) ← item(1,map(V,red)),
   not item(2,map(V,red)).
9  item(2,fail) ← extend(2,S), childItem(S,fail).
10 reject ← final, item(2,fail).
11 accept ← final, not reject.

```

Listing 1.2. D-FLAT for a subset-minimization variant of 3-COLORABILITY

contradict the coloring guessed in line 2. This makes sure that only branches are joined that agree on all common vertices, as each vertex occurring in two child nodes must also appear in the current node due to the connectedness condition of tree decompositions. If a guessed predecessor branch has led to a conflict (denoted by a “reject” node type), this information is retained in line 7. Finally, line 8 derives the “accept” node type if no conflict has occurred.

Example 7. For a more involved example, consider a variant of the 3-COLORABILITY problem: Given an input graph, we now want to compute only those proper 3-colorings whose red vertices are minimal (w.r.t. set inclusion) among all proper 3-colorings. A D-FLAT encoding for this problem is shown in Listing 1.2.

Item trees now have height 2, again with the root being an “or” node, but its children now being “and” nodes (line 1). Each item tree node at depth 1 encodes a proper coloring of the current bag elements (lines 2 and 3). Let R denote the set of vertices assigned “red” by this coloring. The children of this item set node encode all proper colorings whose red vertices R' are a subset of R (lines 2–4). If $R' \subset R$, the item “fail” is put into the respective item set (lines 8 and 9). In case a coloring at depth 2 containing “fail” survives until the final decomposition node, the respective leaf of the item tree is set to the type “reject” (line 10). As this leaf witnesses, its parent encodes a proper coloring whose red vertices are not minimal. Since this parent is an “and” node, it becomes rejecting, too. What remains at depth 1 after pruning rejecting nodes are exactly those nodes that can be extended (via the extension pointers) to proper 3-colorings whose red vertices are minimal in the sense we required.

3.3 D-FLAT's Handling of Item Trees

Every time the ASP solver reports an answer set for the current tree decomposition node, D-FLAT creates a new branch in the current item tree, which results in a so-called *uncompressed item tree*. Subsequently D-FLAT prunes subtrees of that tree that can never be part of a solution in order to avoid unnecessary computations in future decomposition nodes. For optimization problems, D-FLAT then propagates information about the optimization values upward in the uncompressed item tree. The item tree so far is called uncompressed because it may contain redundancies that are eliminated in the final step.

Constructing an Uncompressed Item Tree from the Answer Sets. In an answer set, all atoms using `extend`, `item`, `or` and `and` with the same depth argument, as well as `accept` and `reject`, constitute what we call a *node specification*. To determine where branches from different answer sets diverge, D-FLAT uses the following recursive condition: Two node specifications coincide (i.e., describe the same item tree node) iff (1) they are at the same depth in the item tree; (2) their item sets, extension pointers and node types are equal; and (3) both are at depth 0, or their parent node specifications coincide. In this way, an (uncompressed) item tree is obtained from the answer sets.

Propagation of Acceptance Statuses and Pruning of Item Trees. In Section 3.1 we have defined the different node types that an item tree node can have. When D-FLAT has processed all answer sets and constructed the uncompressed item tree, these types come into play. That is to say, D-FLAT then prunes subtrees from the uncompressed item tree.

First of all, if the current tree decomposition node is the root, D-FLAT prunes from the uncompressed item tree any subtree rooted at a node whose type is still undefined. Then, regardless of whether the current decomposition node is the root, D-FLAT prunes subtrees of the uncompressed item tree depending on the *acceptance status* of its nodes. The acceptance status of a node can either be “undefined”, “accepting” or “rejecting”, which we define now.

A node in an item tree is *accepting* if (a) its type is “accept”, (b) its type is “or” and it has an accepting child, or (c) its type is “and” and all children are accepting. A node is *rejecting* if (a) its type is “reject”, (b) its type is “or” and all children are rejecting, or (c) its type is “and” and it has a rejecting child. The acceptance status of nodes that are neither accepting nor rejecting is *undefined*.

After having computed the acceptance status of all nodes in the current item tree, D-FLAT prunes all subtrees rooted at a rejecting node, as we can be sure that these nodes will never be part of a solution.

Note that in case the current decomposition node is the root, there are no nodes with undefined acceptance status because D-FLAT has pruned all subtrees rooted at nodes with undefined type. Therefore, in this case, the remaining tree consists only of accepting nodes. For decision problems, we can thus conclude that the problem instance is positive iff the remaining tree is non-empty. For

enumeration problems, we can follow the extension pointers down to the leaves of the tree decomposition in order to obtain complete solutions by combining all item sets along the way. This is described in more detail in Section 3.4. Recursively extending all item sets in this way would yield a (generally very big) item tree that usually corresponds to the accepting part of a computation tree that an ATM would have when solving the complete problem instance. (But of course D-FLAT does not materialize this entire tree in memory.)

Example 8. Consider again Figure 1b. Because $\mathcal{I}_{\{b,c,d\}}$ is the final item tree in the decomposition traversal, D-FLAT would subsequently remove all nodes with undefined types (but there are none in this case). Then it would prune all rejecting nodes and conclude that the root of $\mathcal{I}_{\{b,c,d\}}$ is accepting because it has an accepting child. Therefore the problem instance is positive. At the decomposition root, we are then left with an item tree having only six leaves, each encoding a proper coloring of the vertices b , c and d , and storing extension pointers that let us extend the respective coloring of these vertices to proper colorings of all the other vertices, too.

Propagation of Optimization Values in Item Trees. For optimization problems, after an uncompressed item tree has been computed, an additional step is done. Each leaf in the item tree stores an optimization value (or “cost”) that has been specified by the user’s program. D-FLAT now propagates these optimization values from the leaves toward the root of the current uncompressed item tree such that the optimization value of a leaf node is its cost, and the optimization value of an “or” or “and” node is the minimum or maximum, respectively, among the optimization values of its children.

Compressing the Item Tree. The uncompressed item tree obtained in the previous step may contain redundancies that must be eliminated in order to avoid an explosion of memory and runtime. The following situations can arise:

- There are two isomorphic sibling subtrees where all corresponding nodes have equal item sets, node types and (when solving an optimization problem) optimization values. In this case, D-FLAT merges these subtrees into one and unifies their sets of extension pointers.
- An optimization problem is being solved and there are two isomorphic sibling subtrees where all corresponding nodes have equal item sets and node types, but the root of one of these subtrees is “better” than the root of the other. In this context, a node n_1 is “better” than one of its siblings, n_2 , if the parent of n_1 and n_2 either has type “or” and the cost of n_1 is less than that of n_2 , or their parent has type “and” and the cost of n_1 is greater than that of n_2 . In this case, D-FLAT retains only the subtree rooted at the “better” node.

For problems in NP, this redundancy elimination can be done on the fly [1].

3.4 Materializing Complete Solutions

After all item trees have been computed, it remains to materialize complete solutions. We first describe how D-FLAT does this for enumeration problems.

In the item tree at the root of the decomposition there are only accepting nodes left after the pruning described in Section 3.3. Starting with the root of this item tree, D-FLAT extends each of the nodes recursively as follows.

To obtain a complete extension of an item tree rooted at a node n , we recursively extend the item set of n by unifying it with all items of nodes reachable via extension pointers. For each child n' of n we also perform this procedure and add each possible extension of n' to the children to the current extension of n . When extending n' in this way, however, D-FLAT takes care to only pick an extension pointer tuple of n' if every node that is being pointed to in this tuple is a child of a node that is used for the current extension of n .

For optimization problems, D-FLAT only materializes optimal solutions. That is, if n is an “or” node with optimization value c , D-FLAT only extends those children of n that actually have the value c . Due to D-FLAT’s propagation of optimization values (cf. Section 3.3), the optimization value of an “or” node is the minimum of the values of its children. For “and” nodes this is symmetric.

D-FLAT allows the depth until which the final item tree is to be extended to be limited by the user. This is useful if, e.g., only existence of a solution shall be determined. In such a case, we could limit the materialization depth to 0. This would lead to only the root of the final item tree being extended. If D-FLAT yields an extension, a solution in fact exists. This is because the final item tree would have no root in case no solutions existed (cf. Section 3.3). Limiting the materialization depth can thus save us from potentially materializing exponentially many solutions for decision problems.

Moreover, limiting the materialization depth is also helpful for counting problems. If the user limits this depth to d and in the final item tree there is a node at depth d having children, D-FLAT prints for each extension of this node how many extended children would have been materialized. For the most common case, where the materialization depth is limited to 0, D-FLAT is able to calculate the number of possible extensions while doing the main bottom-up traversal of the decomposition for computing the item trees. Hence, for classical counting problems, D-FLAT offers quite efficient counting.

3.5 Debugging Support

Since it can be hard to find the cause of erroneous results of D-FLAT encodings, we have developed a debugging tool that visualizes the generated tree decomposition and the computed item trees. It allows to inspect how certain solutions came to be and thus greatly simplifies debugging. We refer to [1] for details.

3.6 Experiments

In [4] we have experimentally evaluated the prototype of D-FLAT. In the current paper, we additionally consider the MAXIMUM INDEPENDENT SET problem with




Fig. 2. Comparison of runtimes on the MAXIMUM INDEPENDENT SET problem

the new version of D-FLAT. For each number of vertices in $\{40, 42, \dots, 80\}$, we have generated 20 random instances that admit a tree decomposition of width 12. On each instance we have executed D-FLAT 1.0.0 (with an encoding using the simplified ASP interface and the default implementation for join nodes of weakly normalized decompositions [1]) as well as the ASP system Clingo 4.3.0. The experiments were run on a system with 48 GB of RAM and 2 Intel Xeon 64-bit CPUs (each with 4 cores) at 2.33 GHz using only a single core.

Figure 2 shows that the runtime of Clingo 4.3.0 increases exponentially, both for finding just a single optimal solution (denoted by “opt” in the figure) as well as counting the number of optimal solutions (“optN”). Note that especially for “optN” the curve seems to flatten toward the right side of the figure, but this is because timeouts (> 10 minutes) were counted as 10 minutes. The runtime of D-FLAT increases linearly and it outperforms Clingo on the larger instances.

4 Conclusion

In this paper we have presented the D-FLAT system for solving problems by means of dynamic programming on a tree decomposition of the instance. The key feature is that D-FLAT allows the user to specify the problem-specific computations in the logic programming language of Answer Set Programming.

We have discussed the most significant changes made since the initial publication of D-FLAT in [4]. In particular, these extensions allow D-FLAT to solve any problem expressible in monadic second-order logic [5]. This significantly extends its range of applicability. For instance, D-FLAT can thus solve Quantified Boolean Formulas or compute subset-minimal models of propositional theories. Furthermore, we provide a debugging tool that facilitates development of algorithms for D-FLAT.

Future work. We will analyze which features of tree decompositions have the most influence on the runtime of D-FLAT. Moreover, we will investigate for which fragments of ASP we can automatically generate a dynamic programming algorithm for tree decompositions from a “traditional” encoding.

Acknowledgments. This work is supported by the Austrian Science Fund (FWF) projects P25518, P25607 and Y698, and by the Vienna University of Technology special fund “Innovative Projekte” (9006.09/008).

References

1. Abseher, M., Bliem, B., Charwat, G., Dusberger, F., Hecher, M., Woltran, S.: D-FLAT: Progress report. Technical Report DBAI-TR-2014-86, Vienna University of Technology (2014)
2. Arnborg, S., Cornel, D.G., Proskurowski, A.: Complexity of finding embeddings in a k -tree. *SIAM J. Algebraic Discrete Methods* 8(2), 277–284 (1987)
3. Aschinger, M., Drescher, C., Gottlob, G., Jeavons, P., Thorstensen, E.: Structural decomposition methods and what they are good for. In: *Proc. STACS. LIPIcs*, vol. 9, pp. 12–28. Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2011)
4. Bliem, B., Morak, M., Woltran, S.: D-FLAT: Declarative problem solving using tree decompositions and answer-set programming. *TPLP* 12(4-5), 445–464 (2012)
5. Bliem, B., Pichler, R., Woltran, S.: Declarative dynamic programming as an alternative realization of courcelle’s theorem. In: Gutin, G., Szeider, S. (eds.) *IPEC 2013. LNCS*, vol. 8246, pp. 28–40. Springer, Heidelberg (2013)
6. Bodlaender, H.L.: A tourist guide through treewidth. *Acta Cybern.* 11(1-2), 1–22 (1993)
7. Bodlaender, H.L.: A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.* 25(6), 1305–1317 (1996)
8. Bodlaender, H.L.: Discovering treewidth. In: Vojtáš, P., Bieliková, M., Charron-Bost, B., Sýkora, O. (eds.) *SOFSEM 2005. LNCS*, vol. 3381, pp. 1–16. Springer, Heidelberg (2005)
9. Bodlaender, H.L., Koster, A.M.C.A.: Combinatorial optimization on graphs of bounded treewidth. *Comput. J.* 51(3), 255–269 (2008)
10. Bodlaender, H.L., Koster, A.M.C.A.: Treewidth computations I. Upper bounds. *Inf. Comput.* 208(3), 259–275 (2010)
11. Brewka, G., Eiter, T., Truszczyński, M.: Answer set programming at a glance. *Commun. ACM* 54(12), 92–103 (2011)
12. Chandra, A.K., Kozen, D., Stockmeyer, L.J.: Alternation. *J. ACM* 28(1), 114–133 (1981)
13. Dermaku, A., Ganzow, T., Gottlob, G., McMahan, B., Musliu, N., Samer, M.: Heuristic methods for hypertree decomposition. In: Gelbukh, A., Morales, E.F. (eds.) *MICAI 2008. LNCS (LNAI)*, vol. 5317, pp. 1–11. Springer, Heidelberg (2008)
14. Downey, R.G., Fellows, M.R.: *Parameterized Complexity. Monographs in Computer Science.* Springer (1999)
15. Gebser, M., Kaminski, R., Kaufmann, B., Ostrowski, M., Schaub, T., Thiele, S.: A user’s guide to gringo, clasp, clingo, and iclingo. Preliminary Draft (2010), <http://potassco.sourceforge.net>
16. Gebser, M., Kaminski, R., Kaufmann, B., Schaub, T.: *Answer Set Solving in Practice. Synthesis Lectures on Artificial Intelligence and Machine Learning.* Morgan & Claypool Publishers (2012)
17. Gelfond, M., Leone, N.: Logic programming and knowledge representation – the A-Prolog perspective. *Artif. Intell.* 138(1-2), 3–38 (2002)
18. Gelfond, M., Lifschitz, V.: Classical negation in logic programs and disjunctive databases. *New Generation Comput.* 9(3/4), 365–386 (1991)

19. Gottlob, G., Leone, N., Scarcello, F.: Hypertree decompositions and tractable queries. *J. Comput. Syst. Sci.* 64(3), 579–627 (2002)
20. Niedermeier, R.: Invitation to Fixed-Parameter Algorithms. Oxford Lecture Series in Mathematics and its Applications. Oxford University Press (2006)
21. Robertson, N., Seymour, P.D.: Graph minors. III. Planar tree-width. *J. Comb. Theory, Ser. B* 36(1), 49–64 (1984)