Structural Decompositions of Epistemic Logic Programs

Markus Hecher[1], Michael Morak[2], and Stefan Woltran[1]

[1] Institute of Logic and Computation, TU Wien, `lastname@dbai.tuwien.ac.at`
[2] Alpen-Adria Universität Klagenfurt, `michael.morak@aau.at`

Abstract. Epistemic logic programs (ELPs) are a popular generalization of standard Answer Set Programming (ASP) providing means for reasoning over answer sets within the language. This richer formalism comes at the price of higher computational complexity reaching up to the fourth level of the polynomial hierarchy. However, in contrast to standard ASP, dedicated investigations towards tractability have not been undertaken yet. In this paper, we give first results in this direction and show that central ELP problems can be solved in linear time for ELPs exhibiting structural properties in terms of bounded treewidth. We also provide a full dynamic programming algorithm that adheres to these bounds. Finally, we show that applying treewidth to a novel dependency structure—given in terms of epistemic literals—allows to bound the number of ASP solver calls in typical ELP solving procedures.

1 Introduction

Epistemic logic programs (ELPs) [32], also referred to as the language of Epistemic Specifications [17], have received renewed attention in the research community as of late. ELPs are an extension of the language of Answer Set Programming (ASP) [5, 31] with epistemic operators. Gelfond [17] introduced the operators **K** and **M** in order to represent the concepts of *known to be true* and *may be true*, and defined an initial semantics. Several improvements to the semantics have since been proposed in the literature [18, 34, 24, 10]. Shen and Eiter [32] realized that these two operators can be represented via a single negation-type operator that they called *epistemic negation*, denoted **not**, and gave a new semantics based on this operator. Morak [28] proposed a novel characterization of the central construct of the ELP semantics: the *worldview*. While a recent analysis [6] has shown that this semantics still does not eliminate all existing flaws, we will make use of it in this paper, since no clear "winner" semantics has as of yet emerged, and our approach should be equally applicable to other existing semantics that have been proposed.

Evaluating ELPs is a computationally hard task. The central decision problem, checking whether an ELP has a worldview, is Σ_3^P-complete, and problems even higher on the polynomial hierarchy exist [32, 28]. In order to deal with this high complexity efficiently, we propose to use a method from the field of parameterized complexity, namely, investigate how the runtime behaves when looking at

different structural parameters of the problem. For standard ASP, this topic has received considerable interest [27, 21, 15, 3, 14, 12] and even systems exhibiting the parameter treewidth have been proposed [13]. However, the parameterized complexity of epistemic ASP has remained largely unexplored so far. In this paper, we will investigate, in particular, whether ELPs can be solved efficiently if their treewidth (i.e., a measure for the tree-likeness of graphs) is bounded.

It turns out that this question can be answered in the affirmative: the main decision problems become tractable. In practice, a dynamic programming algorithm on tree decompositions can be used to exploit this directly. However, we also aim to investigate a more interesting angle. Many ELP solvers today work by making (up to exponentially many) calls to an underlying ASP solving system in order to check worldview existence. Being able to find a bound on the number of these ASP solver calls would be very useful. Using so-called epistemic (primal) graphs of ELPs that focus on epistemically negated literals only, we can again employ treewidth to establish such bounds. This novel use of structural decomposition intuitively works well in some interesting cases including instances of the scholarship eligibility (SE) benchmark set[3], provided with the "EP-ASP" system [33]. Using the epistemic primal graph representation, these instances naturally decompose into their individual sub-problems, that is, one sub-instance of the SE problem for each student within the original instance.

Contributions. Our contributions are summarized below:

- We investigate the complexity of the ELP worldview existence problem when parameterized by the treewidth of the ELP instance. We establish that this problem is fixed-parameter tractable in this setting, and, in fact, can be solved in linear time if the treewidth is bounded from above by a constant. The same holds for the even more complex problem of worldview formula evaluation.
- Then, we propose a novel graph representation of ELPs, namely, the epistemic primal graph and show how this can be exploited to bound the number of calls to an underlying ASP solver in a classical ELP solver setting. It turns out that the number of calls is bounded in case the epistemic primal graph has bounded treewidth.
- Finally, we provide a full dynamic programming algorithm that could be used in practice to directly exploit the tractability result above. We also show that the worst-case runtime of this algoritm cannot be significantly improved under reasonable complexity-theoretic assumptions.

2   Preliminaries

Answer Set Programming (ASP). A *ground logic program* with nested negation (also called answer set program, ASP program, or, simply, logic program) is a pair $\mathscr{P} = (\mathcal{A}, \mathcal{R})$, where $\mathcal{A}$ is a set of propositional (i.e., ground) atoms and $\mathcal{R}$

---

[3] The Scholarship Eligibility problem was a prime motivator for ELPs [17].

is a set of rules of the form $a_1 \vee \cdots \vee a_l \leftarrow a_{l+1}, \ldots, a_m, \neg\ell_1, \ldots, \neg\ell_n$, where the comma symbol stands for conjunction, $0 \le l \le m$, $0 \le n$, $a_i \in \mathcal{A}$ for all $1 \le i \le m$, and each ℓ_i is a *literal*, that is, either an atom a or its (default) negation $\neg a$ for any atom $a \in \mathcal{A}$.[4] Note that, therefore, doubly negated atoms may occur. We will sometimes refer to such rules as *standard rules*. Each rule $r \in \mathcal{R}$ consists of a *head* $H(r) = \{a_1, \ldots, a_l\}$ and a *body* $B(r) = \{a_{l+1}, \ldots, a_m, \neg\ell_1, \ldots, \neg\ell_n\}$, and is alternatively denoted by $H(r) \leftarrow B(r)$. The *positive* body is given by $B^+(r) = \{a_{l+1}, \ldots, a_m\}$. Sometimes, we add a set of rules $\mathcal{R}'$ to a logic program $\mathscr{P} = (\mathcal{A}, \mathcal{R})$. By some abuse of notation, let $\mathscr{P} \cup \mathcal{R}'$ denote the logic program $(\mathcal{A} \cup \mathcal{A}', \mathcal{R} \cup \mathcal{R}')$, where $\mathcal{A}'$ is the set of atoms occurring in the rules of $\mathcal{R}'$.

An *interpretation* I (over $\mathcal{A}$) is a set of atoms, that is, $I \subseteq \mathcal{A}$. A literal ℓ is true in an interpretation $I \subseteq \mathcal{A}$, denoted $I \vDash \ell$, iff $a \in I$ and $\ell = a$, or $a \notin I$ and $\ell = \neg a$; otherwise ℓ is false in I, denoted $I \nvDash \ell$. Finally, for some literal ℓ, we define that $I \vDash \neg\ell$ if $I \nvDash \ell$. This notation naturally extends to sets of literals. An interpretation M is called a *model* of r, denoted $M \vDash r$, if, whenever $M \vDash B(r)$, it holds that $M \vDash H(r)$. We denote the set of models of r by $mods(r)$; the models of a logic program $\mathscr{P} = (\mathcal{A}, \mathcal{R})$ are given by $mods(\mathscr{P}) = \bigcap_{r \in \mathcal{R}} mods(r)$. We also write $I \vDash r$ (resp. $I \vDash \mathscr{P}$) if $I \in mods(r)$ (resp. $I \in mods(\mathscr{P})$). The *GL-reduct* of a logic program $\mathscr{P} = (\mathcal{A}, \mathcal{R})$ with respect to an interpretation I is given by $\mathscr{P}^I = (\mathcal{A}, \mathcal{R}^I)$ with $\mathcal{R}^I = \{H(r) \leftarrow B^+(r) \mid r \in \mathcal{R}, \forall(\neg\ell) \in B(r) : I \nvDash \ell\}$.

Definition 1 ([19, 20, 25]). $M \subseteq \mathcal{A}$ *is an* answer set *of a logic program* $\mathscr{P}$ *if (1)* $M \in mods(\mathscr{P})$ *and (2) there is no subset* $M' \subset M$ *such that* $M' \in mods(\mathscr{P}^M)$.

The set of answer sets of a logic program $\mathscr{P}$ is denoted by $AS(\mathscr{P})$. The *consistency problem* of ASP, that is, to decide whether for a given logic program $\mathscr{P}$ it holds that $AS(\mathscr{P}) \ne \emptyset$, is Σ_P^2-complete [9], and remains so also in the case where doubly negated atoms are allowed in rule bodies [29].

Epistemic Logic Programs. An *epistemic literal* is a formula **not** ℓ, where ℓ is a literal and **not** is the epistemic negation operator. A *ground epistemic logic program (ELP)* is a pair $\Pi = (\mathcal{A}, \mathcal{R})$, where $\mathcal{A}$ is a set of atoms and $\mathcal{R}$ is a set of *ELP rules*, which are implications of the form $a_1 \vee \cdots \vee a_k \leftarrow \ell_1, \ldots, \ell_m, \xi_1, \ldots, \xi_j, \neg\xi_{j+1}, \ldots, \neg\xi_n$, where each a_i is an atom from $\mathcal{A}$, each ℓ_i is a literal over $\mathcal{A}$, and each ξ_i is an *epistemic literal* of the form **not** ℓ, where ℓ is a literal over $\mathcal{A}$. Similarly to logic programs, let $H(r) = \{a_1, \ldots, a_k\}$, and let $B(r) = \{\ell_1, \ldots, \ell_m, \xi_1, \ldots, \xi_j, \neg\xi_{j+1}, \ldots, \neg\xi_n\}$. Further, $at(r) \subseteq \mathcal{A}$ denotes the set of atoms ocurring in ELP rule r, and $atel(r) \subseteq at(r)$ denotes the set of atoms used in epistemic literals of r. These notions naturally extend to sets of rules.

In order to define the semantics of an ELP, we will use the approach by Morak [28], which follows the semantics defined in [32], but uses a different formal representation. Given an ELP $\Pi = (\mathcal{A}, \mathcal{R})$, a *candidate world interpretation (CWI)* I for Π is a consistent subset $I \subseteq \mathcal{L}$, where $\mathcal{L}$ is the set of all literals that can be built from atoms in $\mathcal{A}$. Note that a CWI I naturally gives rise to a three-valued truth assignment to all the atoms in $\mathcal{A}$; hence, we will sometimes treat a CWI I as a triple of disjoint sets $\langle I^P, I^N, I^U \rangle$, where $I^P = \{a \mid a \in I\}$,

---

[4] In this case, we say that it is a literal *over* $\mathcal{A}$.

$I^N = \{a \mid \neg a \in I\}$ and $I^U = (\mathcal{A} \setminus I^P) \setminus I^N$, with the intended meaning that all atoms in I^P, I^N, and I^U are assigned the truth values "always true," "always false," and "unknown," respectively.

With the above definition in mind, we now define when a CWI is compatible with a given set of interpretations.

Definition 2. *Let $\mathcal{I}$ be a set of interpretations over a set of atoms $\mathcal{A}$. Then, a CWI I is* compatible *with $\mathcal{I}$ iff the following conditions hold:*

1. *$\mathcal{I} \neq \emptyset$;*
2. *for each atom $a \in I^P$, it holds that for each $J \in \mathcal{I}$, $a \in J$;*
3. *for each atom $a \in I^N$, it holds that for each $J \in \mathcal{I}$, $a \notin J$;*
4. *for each atom $a \in I^U$, there are $J, J' \in \mathcal{I}$, such that $a \in J$, but $a \notin J'$.*

The *epistemic reduct* [32, 28] of program $\Pi = (\mathcal{A}, \mathcal{R})$ w.r.t. a CWI I, denoted Π^I, is defined as $\Pi^I = (\mathcal{A}, \{r^I \mid r \in \mathcal{R}\})$ where r^I denotes rule r where each epistemic literal **not** ℓ is replaced by $\neg\ell$ if $\ell \in I$, and by $\top$ otherwise. Note that, hence, Π^I is a plain logic program without epistemic negation.[5]

Now, a CWI I is a *candidate worldview (CWV)* of Π iff the set $AS(\Pi^I)$ of answer sets is compatible with I. The set of CWVs of an ELP Π is denoted $cwv(\Pi)$. Following the principle of knowledge minimization, furthermore I is a *worldview (WV)* iff it is a CWV and there is no proper subset $J \subset I$ such that $J \in cwv(\Pi)$. The set of WVs of an ELP Π is denoted $wv(\Pi)$.

One of the main reasoning tasks regarding ELPs is the *worldview existence problem*, that is, given an ELP Π, decide whether $wv(\Pi) \neq \emptyset$ (or, equivalently, whether $cwv(\Pi) \neq \emptyset$). This problem is known to be Σ_3^P-complete [32, 28]. Another interesting reasoning task is deciding, given an ELP $\Pi = (\mathcal{A}, \mathcal{R})$ and an arbitrary propositional formula φ over $\mathcal{A}$, whether φ holds in some WV, that is, whether there exists $W \in wv(\Pi)$ such that $W \vDash \varphi$. This *formula evaluation problem* is even harder, namely Σ_4^P-complete [32].

Tree Decompositions and Treewidth. We assume that graphs are undirected, simple, and free of self-loops. Let $G = (V, E)$ be a graph, T a rooted tree, and χ a labeling function that maps every node t of T to a subset $\chi(t) \subseteq V$ called the *bag* of t. The pair $\mathcal{T} = (T, \chi)$ is called a *tree decomposition (TD)* [30] of G iff (i) for each $v \in V$, there exists a t in T, such that $v \in \chi(t)$; (ii) for each $\{v, w\} \in E$, there exists t in T, such that $\{v, w\} \subseteq \chi(t)$; and (iii) for each r, s, t of T, such that s lies on the unique path from r to t, we have $\chi(r) \cap \chi(t) \subseteq \chi(s)$. For a node t of T, we say that type(t) is leaf if t has no children and $\chi(t) = \emptyset$; join if t has children t' and t'' with $t' \neq t''$ and $\chi(t) = \chi(t') = \chi(t'')$; intr ("introduce") if t has a single child t', $\chi(t') \subseteq \chi(t)$ and $|\chi(t)| = |\chi(t')| + 1$; rem ("removal") if t has a single child t', $\chi(t') \supseteq \chi(t)$ and $|\chi(t')| = |\chi(t)| + 1$. If for every node $t \in T$, type$(t) \in \{$leaf, join, intr, rem$\}$, then (T, χ) is called *nice*. The *width* of a TD is defined as the cardinality of its largest bag minus one. The *treewidth* of a graph G, denoted by $tw(G)$, is the minimum width over all TDs of G.

---

[5] In fact, Π^I may contain triple-negated atoms $\neg\neg\neg a$. According to [25], such formulas are equivalent to simple negated atoms $\neg a$, and we treat them as such.

Monadic Second Order Logic and Courcelle's Theorem. Monadic Second Order logic (MSO) extends First Order logic (FO) with set variables that range over sets of domain elements. Atomic MSO-formulas over a signature σ are either (1) atoms over some predicate in σ; (2) equality atoms; or (3) atoms of the form $x \in S$, where x is a FO variable, and S is a set variable. MSO-formulas are closed under FO operators. It is convenient to use symbols like $\notin$, $\subseteq$, $\subset$, $\cap$, or $\cup$, with the obvious meanings as abbreviations for the corresponding MSO (sub-)formulas.

MSO formulas are important in the context of parameterized complexity in order to establish running time bounds, as the following landmark theorem by Courcelle [7] shows:

Theorem 1 ([7]). *Let φ be a fixed MSO formula over signature σ and let $\mathfrak{A}$ be a σ-structure with $tw(\mathfrak{A}) \leqslant k$, for some integer k. Then, evaluating φ over $\mathfrak{A}$ can be done in time $O(f(k) \cdot |\mathfrak{A}|)$, for some function f not depending on $|\mathfrak{A}|$.*

Problems with a parameter k that can be solved in time $O(f(k) \cdot n^c)$ for constant c, where f does not depend on n are called *fixed-parameter tractable (FPT)* [8].

3   An MSO Encoding for ELPs

The main objective in this section is to investigate how the semantics of ELPs can be encoded in terms of an MSO formula and thereby investigate, from a theoretical perspective, the time complexity of evaluating ELPs, specifically looking at tree-like instances. In order to start, we first need to define how logical structures can be represented as graphs, and how their tree-likeness is defined: Given a structure $\mathfrak{A}$ over some logical signature σ of arity at most two and with domain $dom(\mathfrak{A})$, we say that the treewidth of $\mathfrak{A}$ equals $tw(G_\mathfrak{A})$, where $G_\mathfrak{A} = (V, E)$ is a graph with $V = dom(\mathfrak{A})$ and edge $\{a, b\} \in E$ iff $r(a, b) \in \mathfrak{A}$, where r is some relation in σ.

Now, our goal will be to offer a fixed MSO encoding, in the spirit of [21], that is able to solve the worldview existence problem for an ELP by evaluating it over a suitable logical structure representing the ELP. In order to begin the construction of this, we first need to fix the signature over which our MSO encoding will be expressed. To this end, let signature

$$\sigma = \{atom, rule, h, b, b^\neg, b^{\mathbf{not}}, b^{\mathbf{not}\,\neg}, b^{\neg\mathbf{not}}, b^{\neg\mathbf{not}\,\neg}\},$$

where $atom(a)$ and $rule(r)$ represent the fact that domain elements a and r are an atom and a rule, respectively; where $h(a, r)$ represents that atom a appears in the head of rule r; and where $b^\square(a, r)$, with $\square \in \{\varepsilon, \neg, \mathbf{not}, \neg\mathbf{not}, \mathbf{not}\,\neg, \neg\mathbf{not}\,\neg\}$ and ε the empty word, represents that fact that the sub-formula $\square a$, with a an atom, appears in the body of rule r. We now construct our MSO encoding.

Lemma 1. *Consider the signature σ above. WV existence can be expressed by means of a fixed MSO formula over σ.*

Proof. Recall that, in order to check the existence of a WV, it suffices to check the existence of a CWV (since WVs are simply subset-minimal CWVs). We

will construct an MSO formula $cwv(\mathbf{P}, \mathbf{N}, \mathbf{U})$ with the intended meaning that it evaluates to true iff the input set variables $\mathbf{P}$, $\mathbf{N}$, and $\mathbf{U}$ represent a CWV W with $W^P = \mathbf{P}$, $W^N = \mathbf{N}$, and $W^U = \mathbf{U}$. The form of our formula will be:

$$cwv(\mathbf{P}, \mathbf{N}, \mathbf{U}) \equiv cwi(\mathbf{P}, \mathbf{N}, \mathbf{U}) \wedge \bigwedge_{i=1}^{4} chk_i(\mathbf{P}, \mathbf{N}, \mathbf{U})$$

where cwi ensures that $\mathbf{P}$, $\mathbf{N}$, and $\mathbf{U}$ indeed encode a valid CWI (i.e., a three-partition of the set of atoms stored in $atom$), and chk_i verifies that Condition i of Definition 2 holds. We will now give the construction of these checks.

First, the check for a valid CWI is expressed as follows:

$$cwi(\mathbf{P}, \mathbf{N}, \mathbf{U}) \equiv \forall X \, (atom(X) \Leftrightarrow X \in \mathbf{P} \cup \mathbf{N} \cup \mathbf{U}) \wedge$$
$$\neg \exists X \, ((X \in \mathbf{P} \cap \mathbf{N}) \vee (X \in \mathbf{N} \cap \mathbf{U}) \vee (X \in \mathbf{P} \cap \mathbf{U}))$$

The four remaining checks have a similar structure:

$$chk_1(\mathbf{P}, \mathbf{N}, \mathbf{U}) \equiv \exists \mathbf{X} \, as(\mathbf{X}, \mathbf{P}, \mathbf{N}, \mathbf{U});$$
$$chk_2(\mathbf{P}, \mathbf{N}, \mathbf{U}) \equiv \forall X \, (X \in \mathbf{P} \Rightarrow$$
$$\forall \mathbf{X} \, (as(\mathbf{X}, \mathbf{P}, \mathbf{N}, \mathbf{U}) \Rightarrow X \in \mathbf{X}));$$
$$chk_3(\mathbf{P}, \mathbf{N}, \mathbf{U}) \equiv \forall X \, (X \in \mathbf{N} \Rightarrow$$
$$\forall \mathbf{X} \, (as(\mathbf{X}, \mathbf{P}, \mathbf{N}, \mathbf{U}) \Rightarrow X \notin \mathbf{X}));$$
$$chk_4(\mathbf{P}, \mathbf{N}, \mathbf{U}) \equiv \forall X \, (X \in \mathbf{U} \Rightarrow$$
$$(\exists \mathbf{X} \, (as(\mathbf{X}, \mathbf{P}, \mathbf{N}, \mathbf{U}) \wedge X \in \mathbf{X}) \wedge$$
$$\exists \mathbf{X} \, (as(\mathbf{X}, \mathbf{P}, \mathbf{N}, \mathbf{U}) \wedge X \notin \mathbf{X}))).$$

The four checks encode precisely the conditions of Definition 2, where $as(\mathbf{X}, \mathbf{P}, \mathbf{N}, \mathbf{U})$ is a sub-formula, to be defined below, that expresses that $\mathbf{X}$ is an answer set of the epistemic reduct w.r.t. the CWI represented by the sets $\mathbf{P}$, $\mathbf{N}$, and $\mathbf{U}$. For example, chk_3 encodes that for each atom X that is set to "always false" in the CWI (i.e., $X \in \mathbf{N}$), it must hold that for every stable model $\mathbf{X}$ of the epistemic reduct, X must not be true in that stable model (i.e., $X \notin \mathbf{X}$).

It now remains to define the sub-formula for checking answer sets. This construction is based on the one presented in [21, Theorem 3.5], but adapted to take the computation of the epistemic reduct into account. Firstly, a set of atoms M is an answer set if it is a model and no proper subset of M is a model of the GL-reduct w.r.t. M. This is expressed as follows (note that any model M is also a model of its GL-reduct):

$$as(\mathbf{X}, \mathbf{P}, \mathbf{N}, \mathbf{U}) \equiv gl(\mathbf{X}, \mathbf{X}, \mathbf{P}, \mathbf{N}, \mathbf{U}) \wedge \forall \mathbf{Y} \, (\mathbf{Y} \subset \mathbf{X} \Rightarrow \neg gl(\mathbf{X}, \mathbf{Y}, \mathbf{P}, \mathbf{N}, \mathbf{U})).$$

Intuitively, $gl(\mathbf{X}, \mathbf{Y}, \mathbf{P}, \mathbf{N}, \mathbf{U})$ shall hold iff $\mathbf{Y}$ is a model of the GL-reduct w.r.t. $\mathbf{X}$ of the epistemic reduct w.r.t. the CWI represented by $\mathbf{P}$, $\mathbf{N}$, and $\mathbf{U}$:

$$gl(\mathbf{X}, \mathbf{Y}, \mathbf{P}, \mathbf{N}, \mathbf{U}) \equiv \forall R \, (rule(R) \Rightarrow \exists Z \, (\, (h(Z, R) \wedge Z \in \mathbf{Y})$$
$$\vee (b(Z, R) \wedge Z \notin \mathbf{Y}) \vee (b^{\neg}(Z, R) \wedge Z \in \mathbf{X})$$
$$\vee (b^{\mathbf{not}}(Z, R) \wedge Z \in \mathbf{P} \wedge Z \in \mathbf{X})$$
$$\vee (b^{\mathbf{not} \, \neg}(Z, R) \wedge Z \in \mathbf{N} \wedge Z \notin \mathbf{X})$$
$$\vee (b^{\neg \mathbf{not}}(Z, R) \wedge ((Z \in \mathbf{N} \cup \mathbf{U}) \vee (Z \in \mathbf{P} \wedge Z \notin \mathbf{X}))$$
$$\vee (b^{\neg \mathbf{not} \, \neg}(Z, R) \wedge ((Z \in \mathbf{P} \cup \mathbf{U}) \vee (Z \in \mathbf{N} \wedge Z \in \mathbf{X})))).$$

Note that the definition of the *gl*-relation is such that it precisely mirrors the definition of both the epistemic reduct and the GL-reduct. It amounts to checking that every rule in the GL-reduct is satisfied, and amounts to a large case distinction, dealing with all seven cases of how atoms can appear in a rule (that is, either in the head, or nested under six combinations of default and epistemic negation in the body). For example, in line three, the first disjunct says that rule R is satisfied if there is an atom Z in the positive body, but this atom is not present in the reduct model $\mathbf{Y}$ (satisfying rule R by not satisfying the body). Line four treats the case of an epistemically negated atom Z in the body. Such a rule is satisfied iff Z is set to "always true" in the CWI (since otherwise the epistemic literal is replaced by $\top$ in the epistemic reduct, and the rule cannot be satisfied solely by this body element in this case), and Z is false in the original model $\mathbf{X}$ (since in the epistemic reduct, the epistemic negation turns into default negation in this case, and default-negated atoms are evaluated over $\mathbf{X}$).

   This completes our MSO encoding. Correctness follows by construction, as explained above. In order to solve the WV existence problem via this encoding, we simply have to quantify the relevant set variables:

$$\varphi = \exists\mathbf{P}\exists\mathbf{N}\exists\mathbf{U}\; cwv(\mathbf{P},\mathbf{N},\mathbf{U}).$$

Evaluating this formula over a σ-structure $\mathfrak{P}$ that represents an ELP Π, we get that Π has a CWV iff $\mathfrak{P} \vDash \varphi$.

   With the above reduction in mind, lets take a closer look at what worst-case solving time guarantees we can give for solving ELPs, in particular w.r.t. structural properties. Let $\Pi = (\mathcal{A}, \mathcal{R})$ be an ELP, and let $\mathfrak{P}$ be the σ-structure that represents it. Recall that $tw(\mathfrak{P}) = tw(G_{\mathfrak{P}})$. In our case, $G_{\mathfrak{P}}$ coincides with the so-called *incidence graph* of the ELP Π, a graph representation that is well-known and studied in the literature for a wide range of logic-based formalisms, and, in particular, for ASP [23, 13]. The incidence graph of an ELP Π is the graph $G = (V, E)$ with $V = \mathcal{A} \cup \mathcal{R}$ and $\{a, r\} \in E$ iff atom $a \in \mathcal{A}$ occurs in rule $r \in \mathcal{R}$ in Π. It is not difficult to verify that the σ-structure $\mathfrak{P}$, when represented as a graph, mirrors the incidence graph of Π precisely. From this correspondence, Theorem 1, and Lemma 1, we thus obtain the following, fundamental parameterized complexity result:

Theorem 2. *Let Π be an ELP, let G be its incidence graph, and let $tw(G) \leqslant k$, for some integer k. Then, checking whether $cwv(\Pi) \neq \emptyset$ can be done in time $O(f(k) \cdot |\Pi|)$, for some function f that does not depend on $|\Pi|$.*

   Using a simple extension of the MSO construction in the proof of Lemma 1, we can state a similar result for the formula evaluation problem. The MSO formula

$$\exists\mathbf{P}\exists\mathbf{N}\exists\mathbf{U}\; cwv(\mathbf{P},\mathbf{N},\mathbf{U}) \wedge entails(\mathbf{P},\mathbf{N},\mathbf{U},\varphi)$$
$$\wedge\, \neg\,(\exists\mathbf{P}'\exists\mathbf{N}'\exists\mathbf{U}'\,(\mathbf{P}' \subset \mathbf{P} \vee \mathbf{N}' \subset \mathbf{N}) \wedge cwv(\mathbf{P}',\mathbf{N}',\mathbf{U}'))$$

checks whether there is at least one WV that satisfies formula φ, where the atom $entails(\mathbf{P},\mathbf{N},\mathbf{U},\varphi)$ encodes the check that the WV represented by $\mathbf{P}$, $\mathbf{N}$, and $\mathbf{U}$ cautiously entails formula φ, a straightforward model-checking construction left to the interested reader. We obtain the following (using argument of Theorem 2):

Theorem 3. *Let Π, G, and k be as in Theorem 2, and let φ be a propositional formula. Then, checking whether Π has a WV that cautiously entails φ can be done in time $O(f(k) \cdot |\Pi|)$, for some function f that does not depend on $|\Pi|$.*

From the above theorems we immediately obtain the fact that ELPs of bounded treewidth can be solved in linear time in the size of the ELP. We will investigate how to exploit this result and pinpoint function $f(k)$ in the next sections.

4   Bounding Calls to Standard ASP Solvers

Before providing a concrete algorithm for the FPT result in Theorem 2 we will investigate a more abstract approach. Many ELP solvers today make use of standard ASP solvers to check the compatibility of a CWI with the set of answer sets of its epistemic reduct. However, the number of calls to such an ASP solver can be at worst exponential. In this section, we will propose an algorithm that makes use of the structural relationships between the epistemic literals in an ELP in order to control the number of ASP solver calls needed and give finer-grained worst-case bounds on this number. In the next section, we will then extend these concepts to a full dynamic programming algorithm using the result in Theorem 2.

We first need to define the structural relationship between atoms occurring in epistemic literals in an ELP. To this end, let $\Pi = (\mathcal{A}, \mathcal{R})$ be an ELP. Then, the *primal graph* $\mathbb{P}_\Pi = (V, E)$ of Π is a graph with $V = \mathcal{A}$ and $\{a, b\} \in E$ iff atoms a and b with $a \neq b$ appear together in a rule in $\mathcal{R}$, that is, iff $\exists r \in \mathcal{R} : \{a, b\} \subseteq at(r)$. Two vertices a, b in the primal graph are *non-epistemically connected* iff there is a path $\langle a, v_1, \ldots, v_n, b \rangle$ with $n \geq 0$ in $\mathbb{P}_\Pi$, such that each vertex v_i, $1 \leq i \leq n$, belongs to the set $\mathcal{A} \setminus atel(\Pi)$. Now, the *epistemic primal graph* $\mathbb{E}_\Pi = (V, E)$ of Π is a graph with the vertex set $V = atel(\Pi)$ being the set of atoms appearing in epistemic literals in Π, and an edge $\{a, b\} \in E$ iff $a \neq b$ and vertices a, b are non-epistemically connected in $\mathbb{P}_\Pi$. Intuitively, two atoms from $atel(\Pi)$ form an edge in $\mathbb{E}_\Pi$ iff they are connected in $\mathbb{P}_\Pi$ via atoms that *do not* appear in epistemic literals. The concept of epistemic primal graph is inspired by the notion of the *torso graph* [16], which is used in parameterized complexity to decompose certain abstraction graphs.

Example 1. Consider the classic *scholarship eligibility problem* encoding, first investigated by Gelfond [17]:

$eligible(X) \leftarrow highGPA(X)$
$ineligible(X) \leftarrow lowGPA(X)$
$\bot \leftarrow eligible(X), ineligible(X)$
$interview(X) \leftarrow \textbf{not } eligible(X), \textbf{not } ineligible(X).$

Now, assume the above abstract (non-ground) program is instantiated with two students (assume that it is copied twice and *mike* and *mark* are substituted for X), and that we add the following rules, resulting in ELP Π:

$lowGPA(mike) \vee highGPA(mike)$

$lowGPA(mark) \vee highGPA(mark)$

Epistemic primal graph $\mathbb{E}_\Pi$ contains four nodes: $eligible(mike)$, $ineligible(mike)$, and the same two for $mark$. Further, $\mathbb{E}_\Pi$ does not have any edges except an edge between the two $mike$-atoms, and the same for $mark$. Since $\mathbb{E}_\Pi$ forms a forest, the treewidth of $\mathbb{E}_\Pi$ is 1.

While the epistemic primal graph does not directly provide new complexity results, it will allow us to give firm guarantees on the number of ASP solver calls needed. As a side-effect, this algorithm is conceptually simpler than the one of the next section, but prepares ideas for later.

Algorithms that solve problems of bounded treewidth typically proceed by *dynamic programming (DP)*, bottom-up, along a TD where at each node t of the TD information is gathered [4] in a table $\tau(t)$. A *table* $\tau(t)$ is a set of rows, where a *row* in $\tau(t)$ is a fixed-length sequence of elements. Tables are derived by an algorithm executed in each bag, called *bag algorithm*, which determine the actual content and meaning of the rows. Then, the DP approach $\mathsf{DP_B}$ for an epistemic logic program Π and a given bag algorithm B performs the following steps:

1. Construct graph representation G of Π that is used by B.
2. Heuristically compute a (nice) TD $\mathcal{T} = (T, \chi)$ of G.
3. Execute B for every node t in TD $\mathcal{T}$ in post-order. As input, B takes a node t, a bag $\chi(t)$, a *solving program* (depending on $\chi(t)$ and B), which is the part of Π currently visible in t, and the tables computed at children of t. Bag algorithm B outputs a table $\tau(t)$.
4. Print the result by interpreting the table for root n of T.

Next, we define a bag algorithm EPRIM for the epistemic primal graph representation of Π. To this end, let $\Pi = (\mathcal{A}, \mathcal{R})$ be the given input epistemic program, $\mathcal{T} = (T, \chi)$ be a nice TD of $\mathbb{E}_\Pi$, t be a node of $\mathcal{T}$, and $\prec$ be any arbitrary total ordering among the nodes in $\mathcal{T}$. To ease notation, for some set $X \subseteq atel(\mathcal{R})$, let $conn(X)$ be the set of vertices (i.e., atoms) from $\mathbb{P}_\Pi$ that lie on a path that non-epistemically connects any two vertices a and b in X.[6] We now define the *induced bag rules* for node t of $\mathcal{T}$, denoted by $\mathcal{R}_t^\mathbb{E}$, as follows. For every rule $r \in \mathcal{R}$, r is *compatible* with node t of $\mathcal{T}$ iff (a) $at(r) \cap conn(atel(\mathcal{R})) \subseteq conn(\chi(t))$, and (b) $\chi(t)$ is subset-maximal among all nodes of $\mathcal{T}$. Now, $r \in \mathcal{R}_t^\mathbb{E}$ iff t is the $\prec$-minimal node among all nodes t' in $\mathcal{T}$ with type(t') = intr compatible with r. The *induced bag program* for node t is the ELP $\Pi_t^\mathbb{E} = (at(\mathcal{R}_t^\mathbb{E}), \mathcal{R}_t^\mathbb{E})$.

Observe that any set of vertices that form a clique within $\mathbb{E}_\Pi$ will appear together in some node t of $\mathcal{T}$. Note that for each node t of $\mathcal{T}$ that has not a non-subset-maximal bag, or has one, but is not $\prec$-minimal for any compatible $r \in \mathcal{R}$, the induced bag program is empty. Therefore, we have that for each rule $r \in \mathcal{R}$ there is exactly one node t of $\mathcal{T}$ where $r \in \mathcal{R}_t^\mathbb{E}$, and, even more stringent, that each atom $a \in at(r) \setminus atel(r)$ appears only in the induced bag program of t, but not in any other node. The bag algorithm EPRIM uses the induced bag program

---

[6] Note that we may have that $a = b$, and hence, $conn(\{a\})$ contains all those vertices from $\mathcal{A} \setminus atel(\Pi)$ connected to atom a.

Listing 1: Bag algorithm $\mathsf{EPRIM}(t, \chi_t, \Pi_t^{\mathbb{E}}, \langle \tau_1, \ldots \rangle)$ for nice TDs of the epistemic primal graph representation.

In: Node t, bag χ_t, induced bag program $\Pi_t^{\mathbb{E}}$, tables $\langle \tau_1, \ldots \rangle$ for child
     nodes $\langle t_1, \ldots \rangle$ of t. **Out:** Table $\tau(t)$.
1 if $\text{type}(t) = \textit{leaf}$ then $\tau(t) = \{\langle \emptyset \rangle\}$/* Abbrevs. below.   */
2 else if $\text{type}(t) = \textit{intr}$, and $a \in \chi_t$ is introduced then
3 $\quad \tau(t) = \{\langle J \rangle \mid \langle I \rangle \in \tau_1, J \in \{I_a^+, I_{\neg a}^+, I\}, \mathscr{P} = (\Pi_t^{\mathbb{E}})^J, \mathcal{R}_t^{\mathbb{E}} = \emptyset$ or
4 $\qquad [\ AS(\mathscr{P}) \neq \emptyset, AS(\mathscr{P} \cup \{\bot \leftarrow (J^P \cup \overline{J^N})\}) = \emptyset,$
$\qquad\quad$ for every $b \in J^U : AS(\mathscr{P} \cup \{\bot \leftarrow b\}) \neq \emptyset, AS(\mathscr{P} \cup \{\bot \leftarrow \neg b\}) \neq \emptyset\ ]\}$
5 else if $\text{type}(t) = \textit{rem}$, and $a \notin \chi_t$ is removed then
6 $\quad \tau(t) = \{\langle I_a^- \rangle \qquad\qquad\qquad\qquad\qquad\qquad\quad \mid \langle I \rangle \in \tau_1\}$
7 else if $\text{type}(t) = \textit{join}$ then $\tau(t) = \tau_1 \cap \tau_2$

$S_e^+ := S \cup \{e\}$, $S_e^- := S \setminus \{e, \neg e\}$, $\overline{S} := \{\neg s \mid s \in S\}$.

as its solving program, and, following the argument above, can check all rules containing atoms from $\mathcal{A} \setminus atel(\Pi)$ in a single node. Hence, during its traversal of the tree decomposition, it does not need to store anything about these atoms. Instead, every row computed as part of a table by EPRIM for a node t, called an *epistemic row*, is of the form $\langle I \rangle$, where $I \subseteq 2^{\{a, \neg a \mid a \in \chi(t)\}}$ is a *partial CWI* (that is, a CWI restricted to and defined w.r.t. $\chi(t))$[7].

Listing 1 presents algorithm EPRIM. For the ease of presentation, it deals with nice TDs only, but can be generalized to arbitrary TDs, requiring a more involved case distinction. Intuitively, since for each leaf node t we have $\chi(t) = \emptyset$, bag algorithm EPRIM ensures in Line 1 that $\tau(t)$ consists only of the empty epistemic row. Then, when an atom a appears in bag $\chi(t)$ for a node t, but does not occur in child bags, CWI J, with either $a \in J^P$, $a \in J^N$, or $a \in J^U$, is computed in Line 3. Further, if the solving program with rules $\mathcal{R}_t^{\mathbb{E}}$ is not empty, i.e., t is the unique node responsible for evaluating all the rules in $\mathcal{R}_t^{\mathbb{E}}$, the four conditions of Definition 2 are checked in Line 4. These checks can be performed by calling a black-box ASP solver a limited number of times for each row in t:

Proposition 1. *To compute a row in a table of* EPRIM*, an ASP solver needs to be called at most* $2 + 2 \cdot |\chi(t)|$ *times.*

On an abstract level, bag algorithm EPRIM hence provides a method for solving epistemic logic programs Π by means of plain ASP solvers based on the structure of the epistemic literals of Π. Whenever an epistemic atom a is removed in node t, indicating that a does not occur in any ancestor bag of t, information about the "role" of a in any CWI is not needed anymore. Finally, for join nodes, Line 7 ensures that the CWIs in $\chi(t)$ coincide with the ones that both child bags have in common. The last step is the evaluation of the root node. If in the root a non-empty table is computed by EPRIM, then the input ELP Π has a CWV.

Example 2. The epistemic primal graph from Example 1 is a "best case"-scenario for using our TD-based approach: the TD naturally separates the ELP into

---

[7] This also means that J^P, J^N, and J^U are defined w.r.t. $\chi(t)$.

one part each for the two students, and algorithm EPRIM would evaluate the two completely separately, which is exactly what intuition would tell us to do. However, standard ELP solvers seem to struggle in this setting when the number of students increases; cf. e.g. [2].

From Proposition 1, and the facts that there are only linearly many TD nodes in the size of the input ELP Π, and that the number of rows per tree node is at most exponential in the treewidth of $\mathbb{E}_\Pi$, we obtain the following statement:

Theorem 4. *Given an input ELP Π of size n, algorithm* $\mathsf{DP_{EPRIM}}$ *makes at most $\mathcal{O}(2^k \cdot n)$ calls to an underlying ASP solver, where $k = tw(\mathbb{E}_\Pi)$.*

Correctness of the algorithm presented above can be established along the same lines as for established TD-based dynamic programming algorithms for ASP [23, 13]. A more formal correctness argument will be given in the next section.

5 A Full Dynamic Programming Algorithm

In this section, we will extend the EPRIM algorithm in such a way that it no longer relies on an underlying ASP solver, but solves an ELP completely on its own, using dynamic programming. This new algorithm, PRIM, will operate on the primal graph, instead of on the epistemic primal graph, and makes use of features of the entire ELP structure.

Recall that the primal graph is defined on all atoms of an ELP, instead of just on the ones appearing in epistemic literals. As a result, we need to define a different solving program for TD nodes. To this end, for the remainder of the section, assume we are a given ELP $\Pi = (\mathcal{A}, \mathcal{R})$ to solve. Further, let $\mathcal{T} = (T, \chi)$ be a nice TD of the primal graph $\mathbb{P}_\Pi$ of Π, and t a node of $\mathcal{T}$. Then, the *bag rules* for t, denoted $\mathcal{R}_t$ are defined as the set $\{r \mid r \in \mathcal{R}, at(r) \subseteq \chi(t)\}$, that is, all the rules of Π that are completely "covered" by $\chi(t)$. Further, the *bag program* of t is defined as $\Pi_t = (\mathcal{A} \cap \chi(t), \mathcal{R}_t)$.

In order to define PRIM, we need to define what a row of a table for a TD node t looks like. Since PRIM, in contrast to EPRIM, now also needs to compute the answer sets underlying a CWV, we start with the following, preliminary definition. Let $M \subseteq \chi(t)$ be an interpretation and $\mathcal{C} \subseteq 2^{\chi(t)}$ a set of interpretations w.r.t. $\chi(t)$. Then, we refer to a tuple $\langle M, \mathcal{C} \rangle$ as an *answer set tuple*. This construct, proposed in [13], directly follows the definition of answer sets as in Definition 1, namely, (1) set M, called a *witness*, is used for storing (parts of) an answer set candidate of some ASP program, and (2) set $\mathcal{C}$, called *counterwitnesses*, holds a set of (partial) models of the GL-reduct w.r.t. M that are potential subsets of M, and hence may be counter-examples to M being extendable to an answer set. An answer set tuple with an empty set of counterwitnesses is referred to as *proving answer set tuple*, which, vaguely speaking, proves that M can be indeed extended to an answer set of some ASP program, which the tuple was constructed for. Answer set tuples are used by algorithm PRIM in order to "transport" information—in the form of parts of models restricted to

---

Listing 2: Bag algorithm $\mathsf{PRIM}(t, \chi_t, \Pi_t, \langle \tau_1, \ldots \rangle)$ for nice TDs of the primal graph.

In: Node t, bag χ_t, bag program Π_t, $\langle \tau_1, \ldots \rangle$ is the seq. of tables for child nodes $\langle t_1, \ldots \rangle$ of t. **Out:** Table $\tau(t)$.

1 **if** $\mathrm{type}(t) = $ *leaf* **then** $\tau(t) = \{\langle \emptyset, \{\langle \emptyset, \emptyset \rangle\}, \emptyset, \emptyset \rangle\}$

2 **else if** $\mathrm{type}(t) = $ *intr*, $a \in \chi_t$ *introduced*, $a \notin atel(\Pi_t)$ **then**

3 $\quad \tau(t) = \{\langle I, \mathcal{M}', \mathcal{K}', \mathcal{S}' \rangle \qquad | \langle I, \mathcal{M}, \mathcal{K}, \mathcal{S} \rangle \in \tau_1, \mathscr{P} = (\Pi_t)^I,$
$\qquad \mathcal{M}' = \mathrm{intTs}(a, \mathcal{M}, \mathscr{P}), \mathcal{K}' = \mathrm{intTs}(a, \mathcal{K}, \mathscr{P}),$

4 $\qquad \mathcal{S}' \in \mathrm{succS}(a, \mathcal{M}', \mathcal{S}, \mathscr{P})\}$

5 **else if** $\mathrm{type}(t) = $ *intr*, $a \in \chi_t$ *introduced*, $a \in atel(\Pi_t)$ **then**

6 $\quad \tau(t) = \{\langle J, \mathcal{M}', \mathcal{K}' \cup \mathcal{K}'', \mathcal{S}' \cup \mathcal{S}'' \rangle | \langle I, \mathcal{M}, \mathcal{K}, \mathcal{S} \rangle \in \tau_1, \mathscr{P} =$
$\qquad (\Pi_t)^J, \mathcal{M}' = \mathrm{intTs}(a, \mathcal{M}, \mathscr{P}), \mathcal{K}' = \mathrm{intTs}(a, \mathcal{K}, \mathscr{P}),$

7 $\qquad \mathcal{S}' \in \mathrm{succS}(a, \mathcal{M}', \mathcal{S}, \mathscr{P}), J \in \{I_a^+, I_{\neg a}^+, I\},$

8 $\qquad \mathcal{K}'' = \mathrm{intTs}(a, \bigcup_{S \in \mathcal{M}} \{S \mid a \in J^P\}, \mathscr{P} \cup \{\bot \leftarrow a\}) \cup$
$\qquad \quad \mathrm{intTs}(a, \bigcup_{S \in \mathcal{M}} \{S \mid a \in J^N\}, \mathscr{P} \cup \{\bot \leftarrow \neg a\}),$

9 $\qquad \{M', M''\} \subseteq \mathcal{M}', \mathcal{S}'' = \{M', M'' \mid a \in J^U\} =$
$\qquad \qquad \{M', M'' \mid a \in J^U, M' \vDash \mathscr{P} \cup \{\bot \leftarrow a\}, M'' \vDash \mathscr{P} \cup \{\bot \leftarrow \neg a\}\}\}$

10 **else if** $\mathrm{type}(t) = $ *rem*, *and* $a \notin \chi_t$ *is removed* **then**

11 $\quad \tau(t) = \{\langle I_a^-, \mathcal{M}_a^\sim, \mathcal{K}_a^\sim, \mathcal{S}_a^\sim \rangle \qquad | \langle I, \mathcal{M}, \mathcal{K}, \mathcal{S} \rangle \in \tau_1\}$

12 **else if** $\mathrm{type}(t) = $ *join* **then**

13 $\quad \tau(t) = \{\langle I, \mathcal{M}_1 \sqcap \mathcal{M}_2, [\mathcal{K}_1 \sqcap (\mathcal{K}_2 \cup \mathcal{M}_2)] \cup [\mathcal{K}_2 \sqcap (\mathcal{K}_1 \cup \mathcal{M}_1)],$
$\qquad [\mathcal{S}_1 \sqcap (\mathcal{S}_2 \cup \mathcal{M}_2)] \cup [\mathcal{S}_2 \sqcap (\mathcal{S}_1 \cup \mathcal{M}_1)] \rangle \mid$
$\qquad \qquad \langle I, \mathcal{M}_1, \mathcal{K}_1, \mathcal{S}_1 \rangle \in \tau_1, \langle I, \mathcal{M}_2, \mathcal{K}_2, \mathcal{S}_2 \rangle \in \tau_2,$

14 $\qquad \qquad |\mathcal{S}_1| = |\mathcal{S}_1 \sqcap (\mathcal{S}_2 \cup \mathcal{M}_2)|, |\mathcal{S}_2| = |\mathcal{S}_2 \sqcap (\mathcal{S}_1 \cup \mathcal{M}_1)|\}$

$S_e^+ := S \cup \{e\}$, $S_e^- := S \setminus \{e, \neg e\}$, $S_e^\sim := \{\langle M_e^-, \{C_e^- \mid C \in \mathcal{C}\} \rangle \mid \langle M, \mathcal{C} \rangle \in S\}$, $\mathcal{M}_1 \sqcap \mathcal{M}_2 := \{\langle M, (\mathcal{C}_M^+ \cap \mathcal{D}) \cup (\mathcal{C} \cap \mathcal{D}_M^+) \rangle \mid \langle M, \mathcal{C} \rangle \in \mathcal{M}_1, \langle M, \mathcal{D} \rangle \in \mathcal{M}_2\}$.

the respective bags—of already evaluated rules of the ASP program from the leaves towards the root during TD traversal.

With this definition in mind, we are now ready to define a row for node t used in algorithm PRIM. Such a row, called *primal row*, is of the form $\langle I, \mathcal{M}, \mathcal{K}, \mathcal{S} \rangle$, where I corresponds to a CWI restricted to $\chi(t)$ as in EPRIM, and sets $\mathcal{M}, \mathcal{K}, \mathcal{S}$ consist of answer set tuples. In the root node n of the TD, a specific primal row $\boldsymbol{u} \in \tau(n)$ is required in table $\tau(n)$ to answer the question of WV existence of Π positively, and PRIM is designed to maintain primal rows accordingly. The set $\mathcal{M}$ of answer set tuples in $\boldsymbol{u}$ is used for ensuring Condition (1) of Definition 2, where a proving answer set tuple in $\mathcal{M}$ gives rise to an answer set of some ASP program $\Pi^{I'}$ of some extension $I' \supseteq I$ of I. For ensuring Conditions (2) and (3), the set $\mathcal{K}$ in $\boldsymbol{u}$ shall not contain any proving program tuples, i.e., proving program tuples of $\mathcal{K}$ are required to vanish (get "killed") during the TD traversal, otherwise Conditions (2) or (3) would be violated. Finally, $\mathcal{S}$ in $\boldsymbol{u}$ serves to establish Condition (4), where no non-proving answer set tuple is allowed, that is, each answer set tuple needs to "survive". The discussed properties of such a primal row $\boldsymbol{u}$ are formalized as follows.

Definition 3. *A primal row* $\langle I, \mathcal{M}, \mathcal{K}, \mathcal{S} \rangle$ *is proving if (1) there is a* $\langle M, \emptyset \rangle \in \mathcal{M}$, *(2) there is no* $\langle M, \mathcal{C} \rangle \in \mathcal{K}$ *with* $\mathcal{C} = \emptyset$, *and (3) there is no* $\langle M, \mathcal{C} \rangle \in \mathcal{S}$ *with* $\mathcal{C} \neq \emptyset$.

Algorithm PRIM is designed to ensure existence of such a proving primal row in $\tau(n)$ of root node n of the TD, iff a WV exists. PRIM uses the following constructs, assuming an answer set tuple $\langle M, \mathcal{C} \rangle$, an atom $a \in \mathcal{A}$, and a program $\mathscr{P}$. For updating an answer set tuple, we let $\mathrm{updT}(M, \mathcal{C}, \mathscr{P}) = \{\langle M, \mathcal{C} \cap mods(\mathscr{P}^M) \rangle \mid M \vDash \mathscr{P}\}$. When some atom a is introduced in an intr-type node, we need to distinguish between a already being in the interpretation, or not. We define $\mathrm{intT}(a, M, \mathcal{C}, \mathscr{P}) = \mathrm{updT}(M_a^+, \bigcup_{C \in \mathcal{C}} \{M, C, C_a^+\}, \mathscr{P}) \cup \mathrm{updT}(M, \mathcal{C}, \mathscr{P})$, which is generalized to sets $\mathcal{M}$ of answer set tuples: $\mathrm{intTs}(a, \mathcal{M}, \mathscr{P}) = \bigcup_{\langle M, \mathcal{C} \rangle \in \mathcal{M}} \mathrm{intT}(a, M, \mathcal{C}, \mathscr{P})$. Finally, to obtain good runtime bounds later and at the same time still ensure Condition (4) of Definition 2 using a set $\mathcal{S}$ of answer set tuples, we need to find, for each answer set tuple in $\mathcal{S}$, exactly one "succeeding" answer set tuple among the set $\mathcal{M}$ of answer set tuples. We define $\mathrm{succS}(a, \mathcal{M}, \mathcal{C}, \mathscr{P}) = \{\mathcal{S}' \mid \mathcal{S}' \subseteq \mathcal{M}, |\mathcal{S}'| = |\mathcal{S}|, \text{ for every } \langle M, \mathcal{C} \rangle \in \mathcal{S} : \mathrm{intT}(a, M, \mathcal{C}, \mathscr{P}) \cap \mathcal{S}' \neq \emptyset\}$.

Bag algorithm PRIM, as presented in Listing 2, again distinguishes between different types of tree nodes during the post-order traversal of $\mathcal{T}$. For leafs, Line 1 returns the primal row consisting of the empty CWI, where the second component $\mathcal{M}$ contains only the proving answer set tuple $\langle \emptyset, \emptyset \rangle$ (since $\emptyset$ is the smallest model of the empty program), and $\mathcal{K}, \mathcal{S}$ are both empty as there is no need to remove or create answer set tuples, respectively. If an atom $a \notin atel(\Pi)$ is introduced, Line 3 updates $\mathcal{M}$ and $\mathcal{K}$. Line 4 ensures that each answer set tuple in $\mathcal{S}$ has at least one succeeding answer set tuple in every primal row of table $\tau(t)$. If an atom $a \in atel(\Pi)$ is introduced, the sets $\mathcal{M}$, $\mathcal{K}$, and $\mathcal{S}$ are similarly updated in Line 6, but the three cases (true, false, and unknown) need to be considered when adding a to I. Conditions (2) and (3) are handled in Line 8, where answer set tuples in $\mathcal{M}$ that violate these two conditions (for $a \in J^P$, and $a \in J^N$, respectively) are added to $\mathcal{K}$. For Condition (4), Line 9 ensures that if $a \in J^N$, there is both a succeeding answer set tuple where a is set to true, and one where it is false. If an atom a is removed, a is removed from the primal rows in Line 11, since we have processed every part of Π where a occurs. Finally, for join nodes, we combine only "compatible" primal rows in Line 13. In particular, Line 14 ensures that no answer set tuple is lost in $\mathcal{S}_1$ or $\mathcal{S}_2$ of the child primal rows.

Theorem 5 (Correctness). *Let Π be an ELP Π, and $\mathcal{T} = (T, \chi)$ a TD of $\mathbb{P}_\Pi$. Then there is a proving row in table $\tau(n)$ obtained by $\mathsf{DP_{PRIM}}$ for root n of $\mathcal{T}$*

Proof (Idea). Soundness follows by showing an invariant I for each row in $\tau(t)$, assuming I holds for each table row of the children of t. Completeness holds by establishing all primal rows fulfilling invariant I are in child tables of t, assuming all primal rows fulfilling I are in table $\tau(t)$.

Then, correctness of $\mathsf{DP_{EPRIM}}$, cf., Sec. 4, is a special case.

Corollary 1 (Correctness of $\mathsf{DP_{EPRIM}}$). *Let Π be an ELP Π and $\mathcal{T} = (T, \chi)$ a TD of $\mathbb{E}_\Pi$. Then there is a row in table $\tau(n)$ for root n obtained by $\mathsf{DP_{EPRIM}}$ iff there is a WV for Π.*

Proof (Idea). $\mathcal{T}$ can be turned into a TD $\mathcal{T}'$ of $\mathbb{P}_\Pi$ by adding to each $\chi(t)$ the set $\mathcal{A}$, where $(\Pi_t^{\mathbb{E}}) = (\mathcal{A}, \cdot)$. $\mathsf{DP_{EPRIM}}$ on $\mathcal{T}$ is, therefore, just a simplification of $\mathsf{DP_{PRIM}}$ on $\mathcal{T}'$.

Next, we briefly discuss running time results of $\mathsf{DP_{PRIM}}$.

Theorem 6 (Runtime). $\mathsf{DP_{PRIM}}$ *runs in time* $2^{2^{2^{\mathcal{O}(k)}}} \cdot |\mathcal{A}|$ *for epistemic program* $\Pi = (\mathcal{A}, \mathcal{R})$, *and treewidth* k *of* $\mathbb{P}_\Pi$.

Proof (Idea). We compute [4] TD $\mathcal{T}$ of width k in $2^{\mathcal{O}(k^3)}$. Each table $\tau(t)$ has at most $2^{(3^k \cdot 2^{2^k} \cdot 2^{2^k} \cdot 2^{2^k})}$ rows, and the size of $\Pi_t^{\mathbb{E}}$ is in $\mathcal{O}(2^{k!})$.

Indeed, under reasonable assumptions in computational complexity, that is, the *exponential time hypothesis (ETH)* [22], one cannot significantly improve $\mathsf{DP_{PRIM}}$.

Proposition 2 (cf. [11], Theorem 19). *Let* $\Pi = (\mathcal{A}, \cdot)$ *be an epistemic logic program with* $\mathbb{P}_\Pi$ *of treewidth* k. *Then, unless ETH fails, WV existence for* Π *cannot be decided in time* $2^{2^{2^{o(k)}}} \cdot 2^{o(|\mathcal{A}|)}$.

6   Conclusions

This work provides the first parameterized complexity analysis of ELP solving w.r.t. treewidth. Tree decompositions have been successfully used in the *selp* ELP solver [2], but for a different purpose, namely that ELPs are rewritten into non-ground ASP programs with long rules, which are then split up using *rule decomposition* [1]. Our approach partitions an ELP according to a tree decomposition, and then solves the ELP by evaluating these parts in turn. This is different from ELP splitting [6], which extends splitting theorem for ASP [26].

For future work, we aim to extend our DP algorithm to the formula evaluation problem, which, viz. Theorem 3, should work in a similar fashion to our existing algorithms, given a suitable graph representation. Furthermore, we would like to apply our approach to other ELP semantics; cf. [17, 18, 24]. Looking at our MSO encoding and DP algorithms, we do not anticipate large obstacles there, since the semantics are sufficiently similar.

References

1. Bichler, M., Morak, M., Woltran, S.: The power of non-ground rules in answer set programming. TPLP 16(5-6), 552–569 (2016)
2. Bichler, M., Morak, M., Woltran, S.: Single-shot epistemic logic program solving. In: Proc. IJCAI. pp. 1714–1720 (2018)
3. Bliem, B., Ordyniak, S., Woltran, S.: Clique-width and directed width measures for answer-set programming. In: Proc. ECAI. pp. 1105–1113 (2016)
4. Bodlaender, H.L., Kloks, T.: Efficient and constructive algorithms for the pathwidth and treewidth of graphs. J. Algorithms 21(2), 358–402 (1996)
5. Brewka, G., Eiter, T., Truszczynski, M.: Answer set programming at a glance. Commun. ACM 54(12), 92–103 (2011)
6. Cabalar, P., Fandinno, J., Fariñas del Cerro, L.: Splitting epistemic logic programs. In: Proc. LPNMR. pp. 120–133 (2019)
7. Courcelle, B.: The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. Inf. Comput. 85(1), 12–75 (1990)
8. Downey, R.G., Fellows, M.R.: Parameterized Complexity. Monographs in Computer Science, Springer (1999)

9. Eiter, T., Gottlob, G.: On the computational cost of disjunctive logic programming: Propositional case. Ann. Math. Artif. Intell. 15(3-4), 289–323 (1995)
10. Fariñas del Cerro, L., Herzig, A., Su, E.I.: Epistemic equilibrium logic. In: Proc. IJCAI. pp. 2964–2970 (2015)
11. Fichte, J.K., Hecher, M., Pfandler, A.: TE-ETH: Lower Bounds for QBFs of Bounded Treewidth (2019), preliminary version at `https://tinyurl.com/y7wnvu6w`
12. Fichte, J.K., Hecher, M.: Treewidth and counting projected answer sets. In: Proc. LPNMR. pp. 105–119 (2019)
13. Fichte, J.K., Hecher, M., Morak, M., Woltran, S.: Answer set solving with bounded treewidth revisited. In: Proc. LPNMR. pp. 132–145 (2017)
14. Fichte, J.K., Kronegger, M., Woltran, S.: A multiparametric view on answer set programming. Ann. Math. Artif. Intell. 86(1-3), 121–147 (2019)
15. Fichte, J.K., Szeider, S.: Backdoors to tractable answer set programming. Artif. Intell. 220, 64–103 (2015)
16. Ganian, R., Ramanujan, M.S., Szeider, S.: Combining treewidth and backdoors for CSP. In: Proc. STACS. pp. 36:1–36:17 (2017)
17. Gelfond, M.: Strong introspection. In: Proc. AAAI. pp. 386–391. AAAI Press / The MIT Press (1991)
18. Gelfond, M.: New semantics for epistemic specifications. In: Proc. LPNMR. LNCS, vol. 6645, pp. 260–265. Springer (2011)
19. Gelfond, M., Lifschitz, V.: The stable model semantics for logic programming. In: Proc. ICLP/SLP. pp. 1070–1080. MIT Press (1988)
20. Gelfond, M., Lifschitz, V.: Classical negation in logic programs and disjunctive databases. New Generation Comput. 9(3/4), 365–386 (1991)
21. Gottlob, G., Pichler, R., Wei, F.: Bounded treewidth as a key to tractability of knowledge representation and reasoning. Artif. Intell. 174(1), 105–132 (2010)
22. Impagliazzo, R., Paturi, R., Zane, F.: Which problems have strongly exponential complexity? J. Comput. Syst. Sci. 63(4), 512–530 (2001)
23. Jakl, M., Pichler, R., Woltran, S.: Answer-set programming with bounded treewidth. In: Proc. IJCAI. pp. 816–822 (2009)
24. Kahl, P.T., Watson, R., Balai, E., Gelfond, M., Zhang, Y.: The language of epistemic specifications (refined) including a prototype solver. J. Log. Comput. 25 (2015)
25. Lifschitz, V., Tang, L.R., Turner, H.: Nested expressions in logic programs. Ann. Math. Artif. Intell. 25(3-4), 369–389 (1999)
26. Lifschitz, V., Turner, H.: Splitting a logic program. In: Proc. ICLP. pp. 23–37 (1994)
27. Lonc, Z., Truszczynski, M.: Fixed-parameter complexity of semantics for logic programs. ACM Trans. Comput. Log. 4(1), 91–119 (2003)
28. Morak, M.: Epistemic logic programs: A different world view. In: Proc. ICLP (2019), to appear
29. Pearce, D., Tompits, H., Woltran, S.: Characterising equilibrium logic and nested logic programs: Reductions and complexity. TPLP 9(5), 565–616 (2009)
30. Robertson, N., Seymour, P.D.: Graph minors. III. planar tree-width. J. Comb. Theory, Ser. B 36(1), 49–64 (1984)
31. Schaub, T., Woltran, S.: Special issue on answer set programming. KI 32(2-3) (2018)
32. Shen, Y., Eiter, T.: Evaluating epistemic negation in answer set programming. Artif. Intell. 237, 115–135 (2016)
33. Son, T.C., Le, T., Kahl, P.T., Leclerc, A.P.: On computing world views of epistemic logic programs. In: Proc. IJCAI. pp. 1269–1275 (2017)
34. Truszczynski, M.: Revisiting epistemic specifications. In: Logic Programming, Knowledge Representation, and Nonmonotonic Reasoning - Essays Dedicated to Michael Gelfond on the Occasion of His 65th Birthday (2011)