Leipzig University
Faculty of Mathematics and Computer Science
Computer Science Institute

Augustusplatz 10
04103 Leipzig
Germany

UNIVERSITÄT LEIPZIG

Analyzing the Computational Complexity of Abstract Dialectical Frameworks via Approximation Fixpoint Theory

Technical Report 2 (2013)

Hannes Strass
Computer Science Institute
Leipzig University
Germany

Johannes Peter Wallner
Institute of Information Systems
Vienna University of Technology
Austria

ISSN 1430-3701

Analyzing the Computational Complexity of Abstract Dialectical Frameworks via Approximation Fixpoint Theory

Hannes Strass
Computer Science Institute
Leipzig University
Germany

Johannes Peter Wallner
Institute of Information Systems
Vienna University of Technology
Austria

27th November 2013

Abstract

Abstract dialectical frameworks (ADFs) have recently been proposed as a versatile generalization of Dung's abstract argumentation frameworks (AFs). In this paper, we present a comprehensive analysis of the computational complexity of ADFs. Our results show that while ADFs are one level up in the polynomial hierarchy compared to AFs, there is a useful subclass of ADFs which is as complex as AFs while arguably offering more modeling capacities. As a technical vehicle, we employ the approximation fixpoint theory of Denecker, Marek and Truszczyński, thus showing that it is also a useful tool for complexity analysis of operator-based semantics.

1 Introduction

Formal models of argumentation are increasingly being recognized as viable tools in knowledge representation and reasoning [Bench-Capon and Dunne, 2007]. A particularly successful formalism are Dung's abstract argumentation frameworks (AFs) [1995]. AFs treat arguments as abstract entities and natively represent only attacks between them using a binary relation. Typically, abstract argumentation frameworks are used as a target language for translations from more concrete languages. For example, the Carneades formalism for structured argumentation [Gordon et al., 2007] has been translated to AFs [Van Gijzel and Prakken, 2011]; Caminada and Amgoud [2007] and Wyner et al. [2013] translate rule-based defeasible theories into AFs. Despite their popularity, abstract argumentation frameworks have limitations. Most significantly, their limited expressiveness is a notable obstacle for applications: AF arguments can only attack one another. Furthermore, Caminada and Amgoud [2007] observed how AFs that arise as translations of defeasible theories sometimes lead to unintuitive conclusions.

To address the limitations of abstract argumentation frameworks, researchers have proposed quite a number of generalizations of AFs [Brewka et al., 2013b]. Among the most general of those are Brewka and Woltran's *abstract dialectical frameworks (ADFs)* [2010]. ADFs are even more abstract than AFs: while in AFs arguments are abstract and the relation between arguments is fixed to attack, in ADFs also the relations are abstract (and called *links*). The relationship between different arguments (called *statements* in ADFs) is specified by *acceptance conditions*. These are Boolean functions indicating the conditions under which a statement s

can be accepted when given the acceptance status of all statements with a direct link to s (its *parents*). ADFs have been successfully employed to address the shortcomings of AFs: Brewka and Gordon [2010] translated Carneades to ADFs and for the first time allowed cyclic dependencies amongst arguments; for rule-based defeasible theories we [Strass, 2013b] showed how to deal with the problems observed by Caminada and Amgoud [2007].

There is a great number of semantics for AFs already, and many of them have been generalized to ADFs. Thus it might not be clear to potential ADF users which semantics are adequate for a particular application domain. In this regard, knowing the computational complexity of semantics can be a valuable guide. However, existing complexity results for ADFs are scattered over different papers, miss several semantics and some of them present upper bounds only. In this paper, we provide a comprehensive complexity analysis for ADFs. In line with the literature, we represent acceptance conditions by propositional formulas as they provide a compact and elegant way to represent Boolean functions.

Technically, we base our complexity analysis on the approximation fixpoint theory (AFT) by Denecker et al. [2000, 2003, 2004]. This powerful framework provides an algebraic account of how monotone and nonmonotone two-valued operators can be approximated by monotone three- or four-valued operators. (As an example of an operator to be approximated, think of the two-valued van Emden-Kowalski consequence operator from logic programming.) AFT embodies the intuitions of decades of KR research; we believe that this is very valuable also for relatively recent languages (such as ADFs), because we get the enormously influential formalizations of intuitions of Reiter and others for free. (As a liberal variation on Newton, we could say that approximation fixpoint theory allows us to take the elevator up to the shoulders of giants instead of walking up the stairs.) In fact, approximation fixpoint theory can be and partially has already been used to define some of the semantics of ADFs [Brewka et al., 2013a; Strass, 2013a]. There, we generalized various AF and logic programming semantics to ADFs using AFT, which has provided us with two families of semantics, that we call – for reasons that will become clear later – *approximate* and *ultimate*, respectively. Intuitively speaking, both families approximate the original two-valued model semantics of ADFs, where the ultimate family is more *precise* in a formally defined sense. The present paper employs approximating operators for complexity analysis and thus shows that AFT is also well-suited for studying the computational complexity of formalisms.

Along with providing a comparison of the approximate and ultimate families of semantics, our main results can be summarized as follows. We show that: (1) the computational complexity of ADF decision problems is one level up in the polynomial hierarchy from their AF counterparts; (2) the ultimate semantics are as complex as the approximate semantics, with the notable exception of two-valued stable models; (3) there is a certain subclass of ADFs, called *bipolar* ADFs (BADFs), which is of the same complexity as AFs. Intuitively, in bipolar ADFs all links between statements are supporting or attacking. To formalize these notions, Brewka and Woltran [2010] gave a precise semantical definition of support and attack. In our work, we assume that the link types are specified by the user along with the ADF. We consider this a harmless assumption since the existing applications of ADFs produce bipolar ADFs where the link types are known [Brewka and Gordon, 2010; Strass, 2013b]. This attractiveness of bipolar ADFs from a KR point of view is the most significant result of the paper: it shows that BADFs offer – in addition to AF-like and more general notions of attack – also a syntactical notion of support *without any increase in computational cost.* Support for a statement s, in this case, can be anything among "set support" (all statements in a certain set must be accepted for the support to become active), "individual support" (at least one statement supporting s must be accepted for the support to become active). In the same vein, BADFs offer "set attack" (all statements in a certain set must be accepted for the attack to become active) and the traditional AF-like "individual attack" (at least one statement attacking s must be accepted for the attack to become active). Naturally, these notions of support and attack can be freely combined.

Previously, Brewka et al. [2011] translated BADFs into AFs and suggested indirectly that their complexities align, albeit restricted to two-valued semantics. Here we go a direct route, which has more practical relevance since it directly affects algorithm design. Our work was also inspired by the complexity analysis of assumption-based argumentation by Dimopoulos et al. [2002] – they derived generic results in a way similar to ours.

The paper proceeds as follows. We first provide the background on approximation fixpoint theory, abstract dialectical frameworks and the necessary elements of complexity theory. In the section afterwards, we define the relevant decision problems, survey existing complexity results, use examples to illustrate how operators revise ADF interpretations and show generic upper complexity bounds. In the main section on complexity results for general ADFs, we back up the upper bounds with matching lower bounds; the section afterwards does the same for bipolar ADFs. We round up with a brief discussion of related and future work.

2 Background

A *complete lattice* is a partially ordered set $(A, \sqsubseteq)$ where every subset of A has a least upper and a greatest lower bound. In particular, a complete lattice has a least and a greatest element. An operator $O : A \to A$ is *monotone* if for all $x \sqsubseteq y$ we find $O(x) \sqsubseteq O(y)$. An $x \in A$ is a *fixpoint* of O if $O(x) = x$; an $x \in A$ is a *prefixpoint* of O if $O(x) \sqsubseteq x$ and a *postfixpoint* of O if $x \sqsubseteq O(x)$. Due to a fundamental result by Tarski and Knaster, for any monotone operator O on a complete lattice, the set of its fixpoints forms a complete lattice itself [Davey and Priestley, 2002, Theorem 2.35]. In particular, its least fixpoint $\mathit{lfp}(O)$ exists.

In this paper, we will be concerned with more general algebraic structures: complete partially ordered sets (CPOs). A CPO is a partially ordered set with a least element where each directed subset has a least upper bound. A set is directed iff it is nonempty and each pair of elements has an upper bound in the set. Clearly every complete lattice is a complete partially ordered set, but not necessarily vice versa. Fortunately, complete partially ordered sets still guarantee the existence of (least) fixpoints for monotone operators.

Theorem 1 ([Davey and Priestley, 2002, Theorem 8.22])**.** *In a complete partially ordered set* $(A, \sqsubseteq)$*, any* $\sqsubseteq$*-monotone operator* $O : A \to A$ *has a least fixpoint.*

2.1 Approximation Fixpoint Theory

Denecker et al. [2000] introduce the important concept of an approximation of an operator. In the study of semantics of knowledge representation formalisms, elements of lattices represent objects of interest. Operators on lattices transform such objects into others according to the contents of some knowledge base. Consequently, fixpoints of such operators are then objects that are fully updated – informally, the knowledge base can neither increase nor decrease the amount of information in a fixpoint.

To study fixpoints of operators O, DMT study their *approximation operators* $\mathcal{O}$.[1] When O operates on a set A, its approximation $\mathcal{O}$ operates on pairs $(x, y) \in A \times A$. Such a pair (x, y) can be seen as representing a *set* of lattice elements by providing a lower bound x and an upper bound y. Consequently, (x, y) approximates all $z \in A$ such that $x \sqsubseteq z \sqsubseteq y$. We will restrict our attention to *consistent* pairs – those where $x \sqsubseteq y$, that is, the set of approximated elements is

[1] The approximation of an operator O is typographically indicated by a calligraphic $\mathcal{O}$.

Kripke-Kleene semantics	$lfp(\mathcal{O})$	grounded pair
admissible/reliable pair (x, y)	$(x, y) \leq_i \mathcal{O}(x, y)$	admissible pair
three-valued supported model (x, y)	$(x, y) = \mathcal{O}(x, y)$	complete pair
M-supported model (x, y)	$(x, y) \leq_i \mathcal{O}(x, y)$ and (x, y) is $\leq_i$-maximal	preferred pair
two-valued supported model (x, x)	$(x, x) = \mathcal{O}(x, x)$	model
two-valued stable model (x, x)	$x = lfp(\mathcal{O}'(\cdot, x))$	stable model

Table 1: Operator-based semantical notions (and their argumentation names on the right) for a complete lattice $(A, \sqsubseteq)$ and an approximating operator $\mathcal{O} : A^c \to A^c$ on the consistent CPO. While an approximating operator always possesses three-valued (post-)fixpoints, two-valued fixpoints need not exist. Clearly, any two-valued stable model is a two-valued supported model is a preferred pair is a complete pair is an admissible pair; furthermore the grounded semantics is a complete pair.

nonempty; we denote the set of all consistent pairs over A by A^c. A pair (x, y) with $x = y$ is called *exact* – it "approximates" a single element of the original lattice.

It is natural to order approximating pairs according to their information content. Formally, for $x_1, x_2, y_1, y_2 \in A$ define the *information ordering* $(x_1, y_1) \leq_i (x_2, y_2)$ iff $x_1 \sqsubseteq x_2$ and $y_2 \sqsubseteq y_1$. This ordering and the restriction to consistent pairs leads to a complete partially ordered set $(A^c, \leq_i)$, the *consistent CPO*. For example, the *trivial pair* $(\bot, \top)$ consisting of $\sqsubseteq$-least $\bot$ and $\sqsubseteq$-greatest lattice element $\top$ approximates all lattice elements and thus contains no information – it is the least element of the CPO $(A^c, \leq_i)$; exact pairs (x, x) are the maximal elements of $(A^c, \leq_i)$.

To define an approximation operator $\mathcal{O} : A^c \to A^c$, one essentially has to define two functions: a function $\mathcal{O}' : A^c \to A$ that yields a revised *lower* bound (first component) for a given pair; and a function $\mathcal{O}'' : A^c \to A$ that yields a revised *upper* bound (second component) for a given pair. Accordingly, the overall approximation is then given by $\mathcal{O}(x, y) = (\mathcal{O}'(x, y), \mathcal{O}''(x, y))$ for $(x, y) \in A^c$. The operator $\mathcal{O} : A^c \to A^c$ is *approximating* iff it is $\leq_i$-monotone and it satisfies $\mathcal{O}'(x, x) = \mathcal{O}''(x, x)$ for all $x \in A$, that is, $\mathcal{O}$ assigns exact pairs to exact pairs. Such an $\mathcal{O}$ then *approximates* an operator $O : A \to A$ on the original lattice iff $\mathcal{O}'(x, x) = O(x)$ for all $x \in A$.

The main contribution of Denecker et al. [2000] was the association of the *stable operator* to an approximating operator. Their original definition was four-valued; in this paper we are only interested in two-valued stable models and simplified the definitions. For an approximating operator $\mathcal{O}$ on a consistent CPO, a (two-valued) pair $(x, x) \in A^c$ is a (two-valued) *stable model of* $\mathcal{O}$ iff x is the least fixpoint of the operator $\mathcal{O}'(\cdot, x)$ defined by $w \mapsto \mathcal{O}'(w, x)$ for $w \sqsubseteq x$. This general, lattice-theoretic approach yields a uniform treatment of the standard semantics of the major nonmonotonic knowledge representation formalisms – logic programming, default logic and autoepistemic logic [Denecker et al., 2003].

In subsequent work, Denecker et al. [2004] presented a general, abstract way to define the most precise – called the *ultimate* – approximation of a given operator O. Most precise here refers to a generalisation of $\leq_i$ to operators, where for $\mathcal{O}_1, \mathcal{O}_2$, they define $\mathcal{O}_1 \leq_i \mathcal{O}_2$ iff for all $x \sqsubseteq y \in A$ it holds that $\mathcal{O}_1(x, y) \leq_i \mathcal{O}_2(x, y)$.

Denecker et al. [2004] show that the most precise approximation of O is $\mathcal{U}_O : A^c \to A^c$ with

$$(x, y) \mapsto \left(\bigsqcap \{O(z) \mid x \sqsubseteq z \sqsubseteq y\}, \bigsqcup \{O(z) \mid x \sqsubseteq z \sqsubseteq y\} \right)$$

where $\sqcap$ denotes the greatest lower bound and $\sqcup$ the least upper bound in the complete lattice $(A, \sqsubseteq)$.

In recent work, we defined new operator-based semantics inspired by semantics from logic programming and abstract argumentation [Strass, 2013a].[2] An overview is in Table 1.

2.2 Abstract Dialectical Frameworks

An abstract dialectical framework (ADF) is a directed graph whose nodes represent statements or positions which can be accepted or not. The links represent dependencies: the status of a node s only depends on the status of its parents (denoted $par(s)$), that is, the nodes with a direct link to s. In addition, each node s has an associated acceptance condition C_s specifying the exact conditions under which s is accepted. C_s is a function assigning to each subset of $par(s)$ one of the truth values $\mathbf{t}$, $\mathbf{f}$. Intuitively, if for some $R \subseteq par(s)$ we have $C_s(R) = \mathbf{t}$, then s will be accepted provided the nodes in R are accepted and those in $par(s) \setminus R$ are not accepted.

Definition 1. An *abstract dialectical framework* is a tuple $\Xi = (S, L, C)$ where

- S is a set of statements (positions, nodes),
- $L \subseteq S \times S$ is a set of links,
- $C = \{C_s\}_{s \in S}$ is a collection of total functions $C_s : 2^{par(s)} \to \{\mathbf{t}, \mathbf{f}\}$, one for each statement s. The function C_s is called *acceptance condition of* s.

It is often convenient to represent acceptance conditions by propositional formulas. In particular, we will do so for the complexity results of this paper. There, each C_s is represented by a propositional formula φ_s over $par(s)$. Then, clearly, $C_s(R \cap par(s)) = \mathbf{t}$ iff $R \models \varphi_s$. Furthermore, throughout the paper we will denote ADFs by Ξ and tacitly assume that $\Xi = (S, L, C)$ unless stated otherwise.

Brewka and Woltran [2010] introduced a useful subclass of ADFs called *bipolar*: Intuitively, in bipolar ADFs (BADFs) each link is supporting or attacking (or both). Formally, a link $(r, s) \in L$ is *supporting in* Ξ iff for all $R \subseteq par(s)$, we have that $C_s(R) = \mathbf{t}$ implies $C_s(R \cup \{r\}) = \mathbf{t}$; symmetrically, a link $(r, s) \in L$ is *attacking in* Ξ iff for all $R \subseteq par(s)$, we have that $C_s(R \cup \{r\}) = \mathbf{t}$ implies $C_s(R) = \mathbf{t}$. An ADF $\Xi = (S, L, C)$ is *bipolar* iff all links in L are supporting or attacking; we use L^+ to denote all supporting and L^- to denote all attacking links of L in Ξ. For an $s \in S$ we define $att_\Xi(s) = \{x \mid (x, s) \in L^-\}$ and $supp_\Xi(s) = \{x \mid (x, s) \in L^+\}$.

The semantics of ADFs can be defined using approximating operators. For two-valued semantics of ADFs we are interested in sets of statements, that is, we work in the complete lattice $(A, \sqsubseteq) = (2^S, \subseteq)$. To approximate elements of this lattice, we use consistent pairs of sets of statements and the associated consistent CPO $(A^c, \leq_i)$ – the *consistent CPO over S-subset pairs*. Such a pair $(X, Y) \in A^c$ can be regarded as a three-valued interpretation where all elements in X are true, those in $Y \setminus X$ are unknown and those in $S \setminus Y$ are false. (This allows us to use "pair" and "interpretation" synonymously from now on.) The following definition specifies how to revise a given three-valued interpretation.

Definition 2 ([Strass, 2013a, Definition 3.1]). Let Ξ be an ADF. Define the following operator $\mathcal{G}_\Xi : 2^S \times 2^S \to 2^S \times 2^S$ by

$$\begin{aligned}
\mathcal{G}_\Xi(X, Y) &= (\mathcal{G}'_\Xi(X, Y), \mathcal{G}'_\Xi(Y, X)) \\
\mathcal{G}'_\Xi(X, Y) &= \{s \in S \mid B \subseteq par(s), C_s(B) = \mathbf{t}, B \subseteq X, \\
&\qquad (par(s) \setminus B) \cap Y = \emptyset\}
\end{aligned}$$

[2] To be precise, we used a slightly different technical setting there. The results can however be transferred to the present setting [Denecker et al., 2004, Theorem 4.2].

Intuitively, statement s is included in the revised lower bound iff the input pair provides sufficient reason to do so, given acceptance condition C_s. Although the operator is defined for all pairs (including inconsistent ones), its restriction to consistent pairs is well-defined since it maps consistent pairs to consistent pairs. This operator defines the *approximate* family of ADF semantics according to Table 1. Based on the three-valued operator $\mathcal{G}_\Xi$, a two-valued one-step consequence operator for ADFs can be defined by $G_\Xi(X) = \mathcal{G}'_\Xi(X, X)$. The general result of Denecker et al. [2004] (Theorem 5.6) then immediately defines the ultimate approximation of G_Ξ as the operator $\mathcal{U}_\Xi$ given by $\mathcal{U}_\Xi(X, Y) = (\mathcal{U}'_\Xi(X, Y), \mathcal{U}''_\Xi(X, Y))$ with

- $\mathcal{U}'_\Xi(X, Y) = \{s \in S \mid \text{ for all } X \subseteq Z \subseteq Y, Z \models \varphi_s\}$ and
- $\mathcal{U}''_\Xi(X, Y) = \{s \in S \mid \text{ for some } X \subseteq Z \subseteq Y, Z \models \varphi_s\}$.

Incidentally, Brewka and Woltran [2010] already defined this operator, which was later used to define the *ultimate* family of ADF semantics according to Table 1 [Brewka et al., 2013a].[3] In this paper, we will refer to the two families of three-valued semantics as "approximate σ" and "ultimate σ" for σ among admissible, grounded, complete, preferred and stable. For two-valued supported models (or simply models), approximate and ultimate semantics coincide.

Although Table 1 defines two-valued stable models also for the ultimate operator, Brewka et al. [2013a] have their own tailor-made definition of two-valued stable models. There, a two-valued pair (M, M) is a stable model of an ADF $\Xi = (S, L, C)$ iff M is the lower bound of the ultimate grounded semantics of the reduced ADF $\Xi^M = (M, L \cap (M \times M), C^M)$ where the reduced acceptance formula for an $s \in S$ is given by the partial evaluation $\varphi_s^{(\emptyset,M)}$: For a propositional formula φ over vocabulary P and $X \subseteq Y \subseteq P$ we define the *partial valuation of* φ *by* (X, Y) as $\varphi^{(X,Y)} = \varphi[p/\mathbf{t} : p \in X][p/\mathbf{f} : p \in P \setminus Y]$. This partial evaluation takes the two-valued part of (X, Y) and replaces the evaluated variables by their truth values. Naturally, $\varphi^{(X,Y)}$ is a formula over the vocabulary $Y \setminus X$.

It is not hard to prove that the definition of two-valued stable models given by Brewka et al. [2013a] coincides with ultimate two-valued stable models. We start with an easy observation.

Lemma. *Let φ be a propositional formula over vocabulary S, and let A, B, C, D be sets with $A \subseteq B \subseteq S$ and $C \subseteq D \subseteq S$.*

$$\left(\varphi^{(A,B)}\right)^{(C,D)} = \varphi^{(A \cup C, B \cap D)}$$

Next, it is easy to see that M is the lower bound of the ultimate grounded semantics of the reduced ADF $\Xi^M = (M, L \cap (M \times M), C^M)$ if and only if (M, M) is the ultimate grounded semantics of Ξ^M. Furthermore, M is a model of Ξ, whence it is a model of Ξ^M. Thus all acceptance formulas in Ξ^M are satisfiable and for any $X \subseteq M$ we get $\mathcal{U}''_{\Xi^M}(X, M) = M$. That is, during computation of the least fixpoint of $\mathcal{U}_{\Xi^M}$, the upper bound remains constant at M. Now for any $X \subseteq M$ and $s \in S$, we have $s \in \mathcal{U}'_\Xi(X, M)$ iff $\varphi_s^{(X,M)}$ is a tautology iff $\left(\varphi_s^{(\emptyset,M)}\right)^{(X,M)}$ is a tautology iff $s \in \mathcal{U}'_{\Xi^M}(X, M)$. So in the complete lattice $(2^M, \subseteq)$, the operators $\mathcal{U}'_\Xi(\cdot, M)$ and $\mathcal{U}'_{\Xi^M}(\cdot, M)$ coincide. Therefore, their least fixpoints coincide.

2.3 Complexity theory

We assume familiarity with the complexity classes P, NP and coNP, as well as with polynomial reductions and hardness and completeness for these classes. We also make use of the polynomial

[3] Technically, Brewka et al. [2013a] represented interpretations not by pairs $(X, Y) \in A^c$ but by mappings $v : S \to \{\mathbf{t}, \mathbf{f}, \mathbf{u}\}$ into the set of truth values $\mathbf{t}$ (true), $\mathbf{f}$ (false) and $\mathbf{u}$ (unknown or undecided). Clearly the two representations are interchangeable.

hierarchy, that can be defined (using oracle Turing machines) as follows: $\Sigma_0^P = \Pi_0^P = \Delta_0^P = \mathsf{P}$, $\Sigma_{i+1}^P = \mathsf{NP}^{\Sigma_i^P}$, $\Pi_{i+1}^P = \mathsf{coNP}^{\Sigma_i^P}$, $\Delta_{i+1}^P = \mathsf{P}^{\Sigma_i^P}$ for $i \geq 0$.

As a somewhat non-standard polynomial hierarchy complexity class, we use D_i^P, a generalisation of the complexity class DP to the polynomial hierarchy. A language is in DP iff it is the intersection of a language in NP and a language in coNP. Generally, a language is in D_{i+1}^P iff it is the intersection of a language in Σ_i^P and a language in Π_i^P. The canonical problem of $\mathsf{DP} = \mathrm{D}_2^P$ is SAT-UNSAT, the problem to decide for a given pair (ψ_1, ψ_2) of propositional formulas whether ψ_1 is satisfiable and ψ_2 is unsatisfiable. Obviously, by definition $\Sigma_i^P, \Pi_i^P \subseteq \mathrm{D}_{i+1}^P \subseteq \Delta_{i+1}^P$ for all $i \geq 0$.

3 Preparatory Considerations

We first introduce some notation to make precise what decision problems we will analyze. For a set S, let

- $(A^c, \leq_i)$ be the consistent CPO of S-subset pairs,
- $\mathcal{O}$ an approximating operator on $(A^c, \leq_i)$,
- $\sigma \in \{adm, com, grd, pre, 2su, 2st\}$ a semantics among admissible, complete, grounded, preferred, two-valued supported and two-valued stable semantics, respectively.

In the *verification* problem we decide whether $(X, Y) \in A^c$ is a σ-model/pair of $\mathcal{O}$, denoted by $\mathsf{Ver}_\sigma^{\mathcal{O}}(X, Y)$. In the *existence* problem we ask whether there exists a σ-model/pair of $\mathcal{O}$ which is non-trivial, i.e. different to $(\emptyset, S)$, denoted by $\mathsf{Exists}_\sigma^{\mathcal{O}}$. For query reasoning and $s \in S$ we consider the problem of deciding whether there exists a σ-model/pair (X, Y) of $\mathcal{O}$ s.t. $s \in X$, denoted by $\mathsf{Cred}_\sigma^{\mathcal{O}}(s)$ (*credulous* reasoning) and the problem of deciding whether in all σ-models/pairs (X, Y) of $\mathcal{O}$ we have $s \in X$, denoted by $\mathsf{Skept}_\sigma^{\mathcal{O}}(s)$ (*skeptical* reasoning). Note that it is no restriction to check only for truth, since checking for falsity of an $s \in S$ can be modeled by introducing a new statement s' that behaves like the logical negation of s, by setting its acceptance condition to $\varphi_{s'} = \neg s$.

3.1 Existing results

We briefly survey – to the best of our knowledge – all existing complexity results for abstract dialectical frameworks. For general ADFs Ξ and the ultimate family of semantics, Brewka et al. [2013a] have shown the following:

- $\mathsf{Ver}_{\mathsf{2su}}^{\mathcal{U}_\Xi}$ is in P, $\mathsf{Exists}_{\mathsf{2su}}^{\mathcal{U}_\Xi}$ is NP-complete (Proposition 5)
- $\mathsf{Ver}_{\mathsf{adm}}^{\mathcal{U}_\Xi}$ is coNP-complete (Proposition 10)
- $\mathsf{Ver}_{\mathsf{grd}}^{\mathcal{U}_\Xi}$ and $\mathsf{Ver}_{\mathsf{com}}^{\mathcal{U}_\Xi}$ are D_2^P-complete (Theorem 6, Cor. 7)
- $\mathsf{Ver}_{\mathsf{2st}}^{\mathcal{U}_\Xi}$ is in D_2^P (Proposition 8)
- $\mathsf{Exists}_{\mathsf{2st}}^{\mathcal{U}_\Xi}$ is Σ_2^P-complete (Theorem 9)

For bipolar ADFs, Brewka and Woltran [2010] showed that $\mathsf{Ver}_{\mathsf{grd}}^{\mathcal{U}_\Xi}$ is in P (Proposition 15). So particularly for BADFs, this paper will greatly illuminate the complexity landscape.

3.2 Relationship between the operators

Since $\mathcal{U}_\Xi$ is the ultimate approximation of G_Ξ it is clear that for any $X \subseteq Y \subseteq S$ we have $\mathcal{G}_\Xi(X,Y) \leq_i \mathcal{U}_\Xi(X,Y)$. In other words, the ultimate revision operator produces new bounds that are at least as tight as those of the approximate operator. More explicitly, the ultimate new lower bound always contains the approximate new lower bound: $\mathcal{G}'_\Xi(X,Y) \subseteq \mathcal{U}'_\Xi(X,Y)$; conversely, the ultimate new upper bound is contained in the approximate new upper bound: $\mathcal{U}''_\Xi(X,Y) \subseteq \mathcal{G}''_\Xi(X,Y)$. Somewhat surprisingly, it turns out that the revision operators for the upper bound coincide.

Lemma 2. *Let $\Xi = (S,L,C)$ be an ADF and $X \subseteq Y \subseteq S$.*

$$\mathcal{G}''_\Xi(X,Y) = \mathcal{U}''_\Xi(X,Y)$$

Proof. Let $s \in S$. We will use that for all $B,X,P \subseteq S$, we find $(P \setminus B) \cap X = \emptyset$ iff $P \cap X \subseteq B$. Now $s \in \mathcal{G}''_\Xi(X,Y)$

$$\begin{aligned} &\text{iff } \exists B : B \subseteq par(s) \cap Y \text{ and } C_s(B) = \mathbf{t} \text{ and} \\ &\qquad\qquad (par(s) \setminus B) \cap X = \emptyset \\ &\text{iff } \exists B : par(s) \cap X \subseteq B \subseteq par(s) \cap Y \text{ and } C_s(B) = \mathbf{t} \\ &\text{iff } \exists Z : X \subseteq Z \subseteq Y \text{ and } C_s(Z \cap par(s)) = \mathbf{t} \\ &\text{iff } s \in \mathcal{U}''_\Xi(X,Y) \end{aligned}$$

□

The operators for computing a new lower bound are demonstrably different, since we can find Ξ and (X,Y) with $\mathcal{U}'_\Xi(X,Y) \not\subseteq \mathcal{G}'_\Xi(X,Y)$, as the following ADF shows.

Example 1. Consider the ADF $D = (\{a\}, \{(a,a)\}, \{\varphi_a\})$ with one self-dependent statement a that has acceptance formula $\varphi_a = a \vee \neg a$. In Figure 1, we show the relevant CPO and the behavior of approximate and ultimate operators: we see that $\mathcal{G}_D(\emptyset, \{a\}) <_i \mathcal{U}_D(\emptyset, \{a\})$, which shows that in some cases the ultimate operator is strictly more precise.

So in a sense the approximate operator cannot see beyond the case distinction $a \vee \neg a$. As we will see shortly, this difference really amounts to the capability of tautology checking.

Example 2. ADF $E = (\{a,b\}, \{(b,a),(b,b)\}, \{\varphi_a, \varphi_b\})$ has acceptance formulas $\varphi_a = b \vee \neg b$ and $\varphi_b = \neg b$. So b is self-attacking and the link from b to a is redundant. In Figure 1, we show the relevant CPO and the behavior of the operators $\mathcal{U}_E$ and $\mathcal{G}_E$ on this CPO.

The examples show that the approximate and ultimate families of semantics really are different, save for one straightforward inclusion relation in case of admissible.

Corollary 3. *For any ADF Ξ, we have the following:*

1. *An approximate admissible pair is an ultimate admissible pair, but not vice versa.*
2. *With respect to their sets of pairs, the approximate and ultimate versions of preferred, complete and grounded semantics are $\subseteq$-incomparable.*

Proof.
1. The inclusion follows from $\mathcal{G}_\Xi \leq_i \mathcal{U}_\Xi$. In Example 2, $(\{a\}, \{a,b\})$ is ultimate admissible but not approximate admissible.
2. In Example 2, we have: (1) approximate grounded, preferred and complete semantics coincide; (2) ultimate grounded, preferred and complete semantics coincide; (3) approximate grounded and ultimate grounded semantics are different with no subset relation either way. □

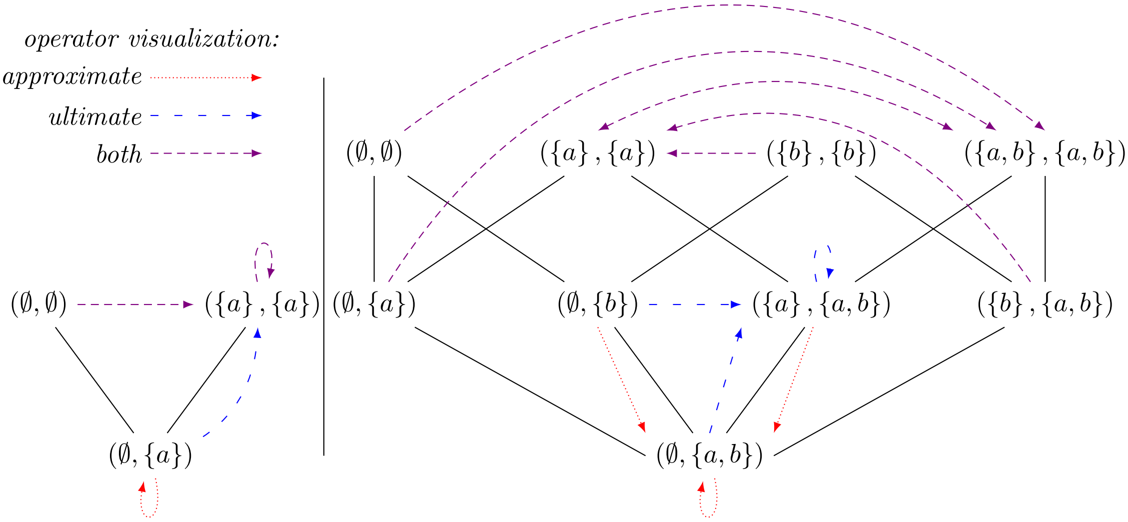

Figure 1: Hasse diagrams of consistent CPOs for the ADFs from Example 1 (left) and Example 2 (right). Solid lines represent the information ordering $\leq_i$. Directed arrows express how revision operators map pairs to other pairs. For pairs where the revisions coincide, the arrows are densely dashed and violet. When the operators revise a pair differently, we use a dotted red arrow for the ultimate and a loosely dashed blue arrow for the approximate operator. Exact (two-valued) pairs are the $\leq_i$-maximal elements. For those pairs, (and any ADF Ξ) it is clear that the operators $\mathcal{U}_\Xi$ and $\mathcal{G}_\Xi$ coincide since they approximate the same two-valued operator G_Ξ. In Example 1 on the left, we can see that the ultimate operator maps all pairs to its only fixpoint $(\{a\}, \{a\})$ where a is true. The approximate operator has an additional fixpoint, $(\emptyset, \{a\})$, where a is unknown. In Example 2 on the right, the major difference between the operators is whether statement a can be derived given that b has truth value unknown. This is the case for the ultimate, but not for the approximate operator. Since there is no fixpoint in the upper row (showing the two-valued operator G_E), the ADF E does not have a two-valued model. Each of the revision operators has however exactly one three-valued fixpoint, which thus constitutes the respective grounded, preferred and complete semantics.

3.3 Operator complexities

We next analyze the computational complexity of deciding whether a single statement is contained in the lower or upper bound of the revision of a given pair. This then leads to the complexity of checking whether current lower/upper bounds are pre- or postfixpoints of the revision operators for computing new lower/upper bounds, that is, whether the revisions represent improvements in terms of the information ordering. Intuitively, these results describe how hard it is to "use" the operators and lay the foundation for the rest of the complexity results.

Proposition 4. *Let Ξ be an ADF, $s \in S$ and $X \subseteq Y \subseteq S$.*

1. *Deciding $s \in \mathcal{G}'_\Xi(X, Y)$ is in* P.
2. *Deciding $\mathcal{G}'_\Xi(X, Y) \subseteq X$ is in* P.
3. *Deciding $X \subseteq \mathcal{G}'_\Xi(X, Y)$ is in* P.

Now let $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$.

4. *Deciding* $s \in \mathcal{O}''(X,Y)$ *is* NP*-complete.*

5. *Deciding* $\mathcal{O}''(X,Y) \subseteq Y$ *is* coNP*-complete.*

6. *Deciding* $Y \subseteq \mathcal{O}''(X,Y)$ *is* NP*-complete.*

Proof. 1. *Since* $X \subseteq Y$*, we have that whenever there exists a* $B \subseteq X \cap par(s)$ *with* $C_s(B) = \mathbf{t}$ *and* $par(s) \setminus B \subseteq S \setminus Y$*, we know that* $B = X \cap par(s)$*: Assume there is an* $r \in (X \cap par(s)) \setminus B$*. Then* $r \in par(s)$ *and* $r \notin B$*, whence* $r \in par(s) \setminus B \subseteq S \setminus Y$*. By* $r \in X \subseteq Y$ *we get* $r \notin S \setminus Y$*, contradiction. Thus* $B = X \cap par(s)$*. Now*

$$\begin{aligned} s \in \mathcal{G}'_\Xi(X,Y) &\text{ iff there exists } B \subseteq X \cap par(s) \text{ with } C_s(B) = \mathbf{t} \text{ and } par(s) \setminus B \subseteq S \setminus Y \\ &\text{ iff } C_s(X \cap par(s)) = \mathbf{t} \text{ and } par(s) \setminus X \subseteq S \setminus Y \\ &\text{ iff } C_s(X \cap par(s)) = \mathbf{t} \text{ and } (Y \setminus X) \cap par(s) = \emptyset \end{aligned}$$

For acceptance functions represented by propositional formulas, $C_s(X \cap par(s)) = \mathbf{t}$ *can be decided in polynomial time, since we only have to check whether* $X \models \varphi_s$*. It can be decided in quadratic time whether there is an undecided parent* $r \in par(s)$ *with* $r \in Y \setminus X$*.*

2. *Deciding* $s \in \mathcal{G}''_\Xi(X,Y)$ *is* NP*-complete:*

   in NP: *By definition,* $\mathcal{G}''_\Xi(X,Y) = \mathcal{G}'_\Xi(Y,X)$*. To verify* $s \in \mathcal{G}'_\Xi(Y,X)$*, we can guess a set* $M \subseteq S$ *and verify that* $M \subseteq Y$*,* $par(s) \setminus M \subseteq S \setminus X$ *and* $M \models \varphi_s$*.*

   NP-hard: *For hardness, we provide a reduction from SAT. Let* ψ *be a propositional formula over vocabulary* P*. Define an ADF* $\Xi = (S,L,C)$ *as follows. Set* $S = P \cup \{z\}$ *where* $z \notin P$*,* $\varphi_z = \psi$ *and* $\varphi_p = p$ *for all* $p \in P$*. Observe that* $par(z) = P$*, and set* $X = \emptyset$ *and* $Y = P$*. Now* $z \in \mathcal{G}''_\Xi(X,Y)$ *iff* $z \in \mathcal{G}'_\Xi(Y,X)$ *iff* $z \in \mathcal{G}'_\Xi(P,\emptyset)$ *iff there is an* $M \subseteq P$ *with* $P \setminus M \cap \emptyset = \emptyset$ *and* $M \models \varphi_z$ *iff there is an* $M \subseteq P$ *with* $M \models \psi$ *iff* ψ *is satisfiable.*

3. *Deciding* $\mathcal{G}'_\Xi(X,Y) \subseteq X$ *is in* P*: For each* $s \in S \setminus X$*, we have to check whether* $s \notin \mathcal{G}_\Xi(X,Y)$*. Any one check can be done in polynomial time by item 1, and there are at most linearly many checks.*

4. *Deciding* $X \subseteq \mathcal{G}'_\Xi(X,Y)$ *is in* P*: For each* $s \in X$*, we have to check whether* $s \in \mathcal{G}'_\Xi(X,Y)$*. Again, one check can be done in polynomial time by item 1, and there are at most linearly many checks.*

5. *Deciding* $\mathcal{G}''_\Xi(X,Y) \subseteq Y$ *is* coNP*-complete:*

   in coNP: *To show* $\mathcal{G}'_\Xi(Y,X) \nsubseteq Y$*, we guess an* $s \in S \setminus Y$ *and a set* $M_s \subseteq par(s)$ *that witnesses* $s \in \mathcal{G}'_\Xi(Y,X)$*.*

   coNP-hard: *Set* $X = \emptyset$ *and* $Y = P$*. By definition, we have* $z \notin P$*, thus* $z \in \mathcal{G}'_\Xi(P,\emptyset)$ *implies* $\mathcal{G}'_\Xi(P,\emptyset) \nsubseteq P$*. Conversely,* $P = S \setminus \{z\}$*, by definition* $\mathcal{G}'_\Xi(P,\emptyset) \subseteq S$ *and hence* $z \notin \mathcal{G}'_\Xi(P,\emptyset)$ *implies* $\mathcal{G}'_\Xi(P,\emptyset) \subseteq P$*. In combination,* $\mathcal{G}'_\Xi(P,\emptyset) \subseteq P$ *iff* $z \notin \mathcal{G}'_\Xi(P,\emptyset)$*. With what we inferred above we see that* $\mathcal{G}''_\Xi(X,Y) \subseteq Y$ *iff* $\mathcal{G}'_\Xi(Y,X) \subseteq Y$ *iff* $\mathcal{G}'_\Xi(P,\emptyset) \subseteq P$ *iff* $z \notin \mathcal{G}'_\Xi(P,\emptyset)$ *iff* ψ *is unsatisfiable, and so the claim follows.*

6. *Deciding* $Y \subseteq \mathcal{G}''_\Xi(X,Y)$ *is* NP*-complete:*

   in NP: *We can guess for each* $s \in Y$ *a set* $M_s \subseteq par(s)$ *to witness* $s \in \mathcal{G}'_\Xi(Y,X)$*. Note that the guesses do not depend on each other.*

NP-hard: *For hardness, we first note that* $p \in \mathcal{G}'_\Xi(P, \emptyset)$ *for all* $p \in P$*, hence* $P \subseteq \mathcal{G}'_\Xi(P, \emptyset)$ *and* $S \subseteq \mathcal{G}'_\Xi(P, \emptyset)$ *iff* $z \in \mathcal{G}'_\Xi(P, \emptyset)$*. Furthermore* z *does not occur in any acceptance formula, so* $\mathcal{G}'_\Xi(S, \emptyset) = \mathcal{G}'_\Xi(P, \emptyset)$*. Now set* $X = \emptyset$ *and* $Y = S$*. Using item 1, it follows that*

$$\begin{aligned} Y \subseteq \mathcal{G}''_\Xi(X, Y) & \text{ iff } Y \subseteq \mathcal{G}'_\Xi(Y, X) \\ & \text{ iff } S \subseteq \mathcal{G}'_\Xi(S, \emptyset) \\ & \text{ iff } S \subseteq \mathcal{G}'_\Xi(P, \emptyset) \\ & \text{ iff } z \in \mathcal{G}'_\Xi(P, \emptyset) \\ & \text{ iff } \psi \text{ is satisfiable} \end{aligned}$$

□

These results can also be formulated in terms of partial evaluations of acceptance formulas: We have $s \in \mathcal{G}'_\Xi(X, Y)$ iff the partial evaluation $\varphi_s^{(X,Y)}$ is a formula without variables that has truth value **t**. Similarly, we have $s \in \mathcal{G}''_\Xi(X, Y)$ iff the partial evaluation $\varphi_s^{(X,Y)}$ is satisfiable. Under standard complexity assumptions, computing a new lower bound with the ultimate operator is harder than with the approximate operator. This is because, intuitively, $s \in \mathcal{U}'_\Xi(X, Y)$ iff the partial evaluation $\varphi_s^{(X,Y)}$ is a tautology.

Proposition 5. *Let* Ξ *be an ADF,* $s \in S$ *and* $X \subseteq Y \subseteq S$*.*

1. *Deciding* $s \in \mathcal{U}'_\Xi(X, Y)$ *is* coNP*-complete.*
2. *Deciding* $\mathcal{U}'_\Xi(X, Y) \subseteq X$ *is* NP*-complete.*
3. *Deciding* $X \subseteq \mathcal{U}'_\Xi(X, Y)$ *is* coNP*-complete.*

Proof. The hardness proofs use the same ADF for their reduction (it is the one from Proposition 4): Let ψ *be a propositional formula over vocabulary* P*. Define an ADF* $\Xi = (S, L, C)$ *as follows. Set* $S = P \cup \{z\}$ *where* $z \notin P$*,* $\varphi_z = \psi$ *and* $\varphi_p = p$ *for all* $p \in P$*.*

1. *Deciding* $s \in \mathcal{U}'_\Xi(X, Y)$ *is* coNP*-complete:*

   in coNP: *To decide that* $s \notin \mathcal{U}'_\Xi(X, Y)$*, we guess a* Z *with* $X \subseteq Z \subseteq Y$ *and verify that* $Z \not\models \varphi_s$*.*

   coNP-hard: *In addition to the reduction above, set* $X = \emptyset$ *and* $Y = P$*. Now*

$$\begin{aligned} z \in \mathcal{U}'_\Xi(X, Y) & \text{ iff } z \in \mathcal{U}'_\Xi(\emptyset, P) \\ & \text{ iff for all } Z \subseteq P, \text{ we have } Z \models \varphi_z \\ & \text{ iff for all } Z \subseteq P, \text{ we have } Z \models \psi \\ & \text{ iff } \psi \text{ is a tautology} \end{aligned}$$

2. *Deciding* $\mathcal{U}'_\Xi(X, Y) \subseteq X$ *is* NP*-complete:*

   in NP: *To show that* $S \setminus X \subseteq S \setminus \mathcal{U}'_\Xi(X, Y)$*, for each statement* $s \in S \setminus X$ *we guess a respective* Z_s *with* $X \subseteq Z_s \subseteq Y$ *that witnesses* $s \notin \mathcal{U}'_\Xi(X, Y)$*.*

   NP-hard: *Set* $X = \emptyset$ *and* $Y = P$*. Then*

$$\begin{aligned} \mathcal{U}'_\Xi(X, Y) \subseteq X & \text{ iff } \mathcal{U}'_\Xi(\emptyset, S) \subseteq \emptyset \\ & \text{ iff } z \notin \mathcal{U}'_\Xi(\emptyset, S) \\ & \text{ iff } \psi \text{ is refutable} \end{aligned}$$

3. *Deciding $X \subseteq \mathcal{U}'_{\sqsubseteq}(X,Y)$ is* coNP*-complete:*

   in coNP: *To verify that $X \not\subseteq \mathcal{U}'_{\sqsubseteq}(X,Y)$, we guess an $s \in X$ and the $X \subseteq Z \subseteq Y$ that witnesses $s \notin \mathcal{U}'_{\sqsubseteq}(X,Y)$.*

   coNP-hard: *Set $X = \{z\}$ and $Y = S$. We observe that z does not occur in any acceptance formula and thus $\mathcal{U}'_{\sqsubseteq}(\{z\},S) = \mathcal{U}'_{\sqsubseteq}(\emptyset,P)$. Then*

$$\begin{aligned} X \subseteq \mathcal{U}'_{\sqsubseteq}(X,Y) &\text{ iff } z \in \mathcal{U}'_{\sqsubseteq}(\{z\},S) \\ &\text{ iff } z \in \mathcal{U}'_{\sqsubseteq}(\emptyset,P) \\ &\text{ iff } \psi \text{ is a tautology} \end{aligned}$$

□

The next result considerably simplifies the complexity analysis of deciding the existence of non-trivial pairs.

Lemma 6. *Let $(A,\sqsubseteq)$ be a complete lattice and $\mathcal{O}$ an approximating operator on A^c. The following are equivalent:*

1. *$\mathcal{O}$ has a non-trivial admissible pair.*
2. *$\mathcal{O}$ has a non-trivial preferred pair.*
3. *$\mathcal{O}$ has a non-trivial complete pair.*

Proof. "(1) ⇒ (2)": *Let $(\bot,\top) <_i (x,y) \leq_i \mathcal{O}(x,y)$. We show that there is a preferred pair $(p,q) \geq_i (x,y)$. Define $D = \{(a,b) \mid (x,y) \leq_i (a,b)\}$, then the pair $(D,\leq_i)$ is a CPO on which $\mathcal{O}$ is an approximating operator. (Obviously $(a,b) \in D$ implies $(x,y) \leq_i (a,b)$ whence by presumption and $\leq_i$-monotonicity of $\mathcal{O}$ we get $(x,y) \leq_i \mathcal{O}(x,y) \leq_i \mathcal{O}(a,b)$ and $\mathcal{O}(a,b) \in D$.) Now any sequence $(a,b) \leq_i \mathcal{O}(a,b) \leq_i \mathcal{O}(\mathcal{O}(a,b)) \leq_i \ldots$ is a non-empty chain in D and therefore has an upper bound in D. By Zorn's lemma, the set of all $\mathcal{O}$-admissible pairs in A has a maximal element $(p,q) \geq_i (x,y) >_i (\bot,\top)$.*

"(2) ⇒ (3)": *By [Strass, 2013a, Theorem 3.10], every preferred pair is complete.*

"(3) ⇒ (1)": *Any complete pair is admissible (Table 1).* □

This directly shows the equivalence of the respective decision problems, that is, it holds that $\mathsf{Exists}^{\mathcal{O}}_{\mathsf{adm}} = \mathsf{Exists}^{\mathcal{O}}_{\mathsf{pre}} = \mathsf{Exists}^{\mathcal{O}}_{\mathsf{com}}$.

Regarding decision problems for querying, skeptical reasoning w.r.t. admissibility is trivial, i.e. $(\emptyset,S)$ is always an admissible pair in any ADF. Further credulous reasoning w.r.t. admissibility, complete and preferred semantics coincides.

Lemma 7. *Let Ξ be an ADF, $\mathcal{O} \in \{\mathcal{G}_{\sqsubseteq},\mathcal{U}_{\sqsubseteq}\}$ and $s \in S$. Then $\mathsf{Cred}^{\mathcal{O}}_{\mathsf{adm}}(s)$ iff $\mathsf{Cred}^{\mathcal{O}}_{\mathsf{com}}(s)$ iff $\mathsf{Cred}^{\mathcal{O}}_{\mathsf{pre}}(s)$.*

Proof. Assume (X,Y) with $s \in X$ is admissible w.r.t. $\mathcal{O}$, then there exists a (X',Y') with $(X,Y) \leq_i (X',Y')$ which is preferred w.r.t. $\mathcal{O}$ and $s \in X'$, see proof of Lemma 6. Since any preferred pair is also complete and any complete pair is also admissible the claim follows. □

3.4 Generic upper bounds

We now show generic upper bounds for the computational complexity of the considered problems. This kind of analysis is in the spirit of the results by Dimopoulos et al. [2002, Section 4]. The first item is furthermore a straightforward generalization of [Denecker et al., 2004, Theorem 6.13].

Theorem 8. *Let S be a finite set, define $A = 2^S$ and let $\mathcal{O}$ be an approximating operator on $(A^c, \leq_i)$, the consistent CPO of S-subset pairs. For $(X, Y) \in A^c$ let the problems of deciding whether $X \subseteq \mathcal{O}'(X, Y)$, as well as $\mathcal{O}''(X, Y) \subseteq Y$ be in Π_i^P; let the problems of deciding $\mathcal{O}'(X, Y) \subseteq X$ as well as $Y \subseteq \mathcal{O}''(X, Y)$ be in Σ_i^P. For any pair $(X, Y) \in A^c$ and statement $s \in S$, we have:*

1. *The least fixpoint of $\mathcal{O}$ can be computed in polynomial time with a polynomial number of calls to a Σ_i^P-oracle.*
2. $\mathsf{Ver}_{\mathsf{adm}}^{\mathcal{O}}(X, Y)$ *is in* Π_i^P; $\mathsf{Cred}_{\mathsf{adm}}^{\mathcal{O}}(s)$ *is in* Σ_{i+1}^P;
3. $\mathsf{Ver}_{\mathsf{com}}^{\mathcal{O}}(X, Y)$ *is in* D_i^P; $\mathsf{Cred}_{\mathsf{com}}^{\mathcal{O}}(s)$ *is in* Σ_{i+1}^P;
4. $\mathsf{Ver}_{\mathsf{pre}}^{\mathcal{O}}(X, Y)$ *is in* Π_{i+1}^P; $\mathsf{Cred}_{\mathsf{pre}}^{\mathcal{O}}(s)$ *is in* Σ_{i+1}^P; $\mathsf{Skept}_{\mathsf{pre}}^{\mathcal{O}}(s)$ *is in* Π_{i+2}^P.

Proof.

1. *For any $(X, Y) \in A^c$ we can use the oracle to compute an application of $\mathcal{O}$ by simply asking whether $z \in \mathcal{O}'(X, Y)$ for each $z \in S$. This means we can compute with a linear number of oracle calls the sets $\mathcal{O}'(X, Y)$ and $\mathcal{O}''(X, Y)$, thus the pair $\mathcal{O}(X, Y)$. Hence we can compute the sequence $(\emptyset, S) \leq_i \mathcal{O}(\emptyset, S) \leq_i \mathcal{O}(\mathcal{O}(\emptyset, S)) \leq_i \ldots$ which converges to the least fixpoint after a linear number of operator applications.*
2. $\mathsf{Ver}_{\mathsf{adm}}^{\mathcal{O}}(X, Y)$ *is in Π_i^P by assumption. For* $\mathsf{Cred}_{\mathsf{adm}}^{\mathcal{O}}(s)$*, we guess a pair (X_1, Y_1) with $s \in X_1$ (resp. $s \in S \setminus Y$) and check if it is admissible w.r.t. $\mathcal{O}$, which is in Π_i^P by 2.*
3. $\mathsf{Ver}_{\mathsf{com}}^{\mathcal{O}}(X, Y)$ *is in D_i^P by assumption.* $\mathsf{Cred}_{\mathsf{com}}^{\mathcal{O}}(s) = \mathsf{Cred}_{\mathsf{adm}}^{\mathcal{O}}(s)$ *by Lemma 7.*
4. *For* $\mathsf{Ver}_{\mathsf{pre}}^{\mathcal{O}}(X, Y)$*, consider the co-problem, i.e. deciding whether (X, Y) is not a preferred pair. We first check if (X, Y) is a complete pair w.r.t. $\mathcal{O}$, which is in D_i^P by 3, i.e. can be achieved via two oracle calls as above. If this holds, we guess a (X_1, Y_1) with $(X, Y) <_i (X_1, Y_1)$ and check if it is a complete pair w.r.t. $\mathcal{O}$.*
   $\mathsf{Cred}_{\mathsf{pre}}^{\mathcal{O}}(X, Y)$*: coincides with credulous reasoning w.r.t. admissibility, see Lemma 7;*
   $\mathsf{Skept}_{\mathsf{pre}}^{\mathcal{O}}(s)$*: Consider the co-problem, i.e. deciding whether there exists a preferred pair (X_1, Y_1) with $X_1 \cap \{a\} = \emptyset$. We guess such a pair (X_1, Y_1) and check if it is a preferred pair w.r.t. $\mathcal{O}$.* □

Naturally, the capability of solving the functional problem of *computing* the grounded semantics allows us to solve the associated decision problems.

Corollary 9. *Under the assumptions of Theorem 8, the problems* $\mathsf{Ver}_{\mathsf{grd}}^{\mathcal{O}}$ *and* $\mathsf{Exists}_{\mathsf{grd}}^{\mathcal{O}}$ *are in* Δ_{i+1}^P.

4 Complexity of General ADFs

Due to the coincidence of $\mathcal{G}_{\Xi}''$ and $\mathcal{U}_{\Xi}''$, the computational complexities of decision problems that concern only the upper bound operator also coincide. This will save both work and space in the subsequent developments. Additionally, for all containment results (except for the grounded semantics), we can use Theorem 8 and need only show hardness.

Proposition 10. *Let Ξ be an ADF, $X, Y \subseteq S$ and consider any $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$.* $\mathsf{Ver}^{\mathcal{O}}_{\mathsf{adm}}(X, Y)$ *is* coNP*-complete.*

Proof. *Hardness follows from Proposition 4, item 5.* □

Recall that a pair (X, Y) is an approximate/ultimate complete pair iff it is a fixpoint of the corresponding (approximate/ultimate) operator. Given the complexities of operator computation, it is straightforward to show the following.

Proposition 11. *Let Ξ be an ADF, $X \subseteq Y \subseteq S$ and consider any $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$.* $\mathsf{Ver}^{\mathcal{O}}_{\mathsf{com}}(X, Y)$ *is* D_2^P*-complete.*

Proof. $\mathcal{O} = \mathcal{G}_\Xi$: **Claim.** *Let Ξ be an ADF and $X, Y \subseteq S$.*

1. *Deciding whether $\mathcal{G}'_\Xi(X, Y) = X$ is D_2^P-complete.*
2. *Deciding whether $\mathcal{G}'_\Xi(X, Y) = Y$ is D_2^P-complete.*
3. *Deciding whether (X, Y) is an approximate four-valued supported model is D_2^P-complete.*
4. *Deciding whether (X, Y) is an approximate three-valued supported model is D_2^P-complete.*

Proof of the claim. All hardness proofs use the same reduction from SAT-UNSAT, the problem to decide for a given pair (ψ_1, ψ_2) of propositional formulas whether ψ_1 is satisfiable and ψ_2 is unsatisfiable. (We can use the techniques from the proof of Proposition 4.) Let (ψ_1, ψ_2) be an instance of SAT-UNSAT. For convenience, we assume w.l.o.g. that ψ_1 uses vocabulary P_1 and formula ψ_2 uses vocabulary P_2 with $P_1 \cap P_2 = \emptyset$. We construct an ADF $\Xi = (S, L, C)$ as follows:

- *$S = P_1 \cup P_2 \cup \{z_1, z_2\}$ (where $z_1, z_2 \notin P_1 \cup P_2$),*
- *$\varphi_p = p$ for all $p \in P_1 \cup P_2$,*
- *$\varphi_{z_1} = \psi_1$ and $\varphi_{z_2} = \psi_2$.*

1. *Deciding whether $\mathcal{G}'_\Xi(X, Y) = X$ is D_2^P-complete:*
   in D_2^P: *We have to decide whether $X \subseteq \mathcal{G}'_\Xi(X, Y)$ (in* NP*) and $\mathcal{G}'_\Xi(X, Y) \subseteq X$ (in* coNP*).*
   D_2^P-complete: *Set $X = S \setminus \{z_2\}$ and $Y = \emptyset$. We have the following:*
   - *$X \subseteq \mathcal{G}'_\Xi(X, Y)$ iff $X \subseteq \mathcal{G}'_\Xi(X, \emptyset)$ iff $z_1 \in \mathcal{G}'_\Xi(X, \emptyset)$ iff ψ_1 is satisfiable.*
   - *$\mathcal{G}'_\Xi(X, Y) \subseteq X$ iff $\mathcal{G}'_\Xi(X, \emptyset) \subseteq X$ iff $z_2 \notin \mathcal{G}'_\Xi(X, \emptyset)$ iff ψ_2 is unsatisfiable.*

   This shows that (ψ_1, ψ_2) is a positive instance of SAT-UNSAT iff $\mathcal{G}'_\Xi(X, Y) = X$.
2. *Deciding whether $\mathcal{G}'_\Xi(X, Y) = Y$ is D_2^P-complete:*
   in D_2^P: *We have to decide whether $Y \subseteq \mathcal{G}'_\Xi(X, Y)$ (in* NP*) and $\mathcal{G}'_\Xi(X, Y) \subseteq Y$ (in* coNP*).*
   D_2^P-hard: *Modify the ADF such that $\varphi_p = \bot$ for all $p \in P_1 \cup P_2$. Set $X = S$ and $Y = \{z_1\}$. We have the following:*
   - *$Y \subseteq \mathcal{G}'_\Xi(X, Y)$ iff $\{z_1\} \subseteq \mathcal{G}'_\Xi(X, \{z_1\})$ iff $z_1 \in \mathcal{G}'_\Xi(X, \{z_1\})$ iff ψ_1 is satisfiable.*
   - *$\mathcal{G}'_\Xi(X, Y) \subseteq Y$ iff $\mathcal{G}'_\Xi(X, \{z_1\}) \subseteq \{z_1\}$ iff $z_2 \notin \mathcal{G}'_\Xi(X, \{z_1\})$ iff ψ_2 is unsatisfiable.*

   Combining these two yields that (ψ_1, ψ_2) is a positive instance of SAT-UNSAT iff $\mathcal{G}'_\Xi(X, Y) = Y$.

3. *Deciding whether (X, Y) is a four-valued supported model is* D_2^P*-complete:* *First of all (X, Y) is a four-valued supported model iff $\mathcal{G}_\Xi(X, Y) = (X, Y)$ iff $\mathcal{G}'_\Xi(X, Y) = X$ and $\mathcal{G}'_\Xi(Y, X) = Y$.*

   in D_2^P: *The guesses for $X \subseteq \mathcal{G}'_\Xi(X, Y)$ and $Y \subseteq \mathcal{G}'_\Xi(Y, X)$ are independent of each other and can be combined. The same holds for the guesses for $\mathcal{G}'_\Xi(X, Y) \subseteq X$ and $\mathcal{G}'_\Xi(Y, X) \subseteq Y$.*

   D_2^P-hard: *For technical reasons, we additionally assume w.l.o.g. that $P_1, P_2 \neq \emptyset$.*[4] *It follows that $\emptyset \neq par(s) \subseteq P_1 \cup P_2$ for all $s \in S$.*
   Now set $X = \emptyset$ and $Y = S \setminus \{z_2\}$ (and note for the next item that $X \subseteq Y$). We observe the following:

   (a) *ψ_1 is satisfiable iff $X \subseteq \mathcal{G}'_\Xi(X, Y)$ and $Y \subseteq \mathcal{G}'_\Xi(Y, X)$:*
      i. *Let ψ_1 be satisfiable. Then $z_1 \in \mathcal{G}'_\Xi(Y, X)$. By definition of the acceptance conditions, $P_1 \cup P_2 \subseteq \mathcal{G}'_\Xi(Y, X)$. Hence $Y \subseteq \mathcal{G}'_\Xi(Y, X)$. $X = \emptyset \subseteq \mathcal{G}'_\Xi(X, Y)$ is trivial.*
      ii. *Let ψ_1 be unsatisfiable. Then $z_1 \notin \mathcal{G}'_\Xi(Y, X)$ and $Y \not\subseteq \mathcal{G}'_\Xi(Y, X)$.*

   (b) *ψ_2 is unsatisfiable iff $\mathcal{G}'_\Xi(Y, X) \subseteq Y$ and $\mathcal{G}'_\Xi(X, Y) \subseteq X$:*
      i. *Let ψ_2 be unsatisfiable. Then $z_2 \notin \mathcal{G}'_\Xi(Y, X)$, and $\mathcal{G}'_\Xi(Y, X) \subseteq Y$. Furthermore $\mathcal{G}'_\Xi(X, Y) = X = \emptyset$. (Assume to the contrary that there is some $s \in \mathcal{G}'_\Xi(X, Y)$. Then there is an $M \subseteq X = \emptyset$ with $M \models \varphi_s$ and $(par(s) \setminus M) \cap Y = \emptyset$. Hence $par(s) \cap Y = \emptyset$. Contradiction, since $P_1 \cup P_2 \subseteq Y$ and all statements in s have non-empty parents among $P_1 \cup P_2$.)*
      ii. *Let ψ_2 be satisfiable. Then $z_2 \in \mathcal{G}'_\Xi(Y, X)$ and $\mathcal{G}'_\Xi(Y, X) \not\subseteq Y$ since $z_2 \notin Y$.*

   We conclude that (ψ_1, ψ_2) is a positive instance of SAT-UNSAT iff ψ_1 is satisfiable and ψ_2 is unsatisfiable iff $X \subseteq \mathcal{G}'_\Xi(X, Y)$ and $Y \subseteq \mathcal{G}'_\Xi(Y, X)$ and $\mathcal{G}'_\Xi(Y, X) \subseteq Y$ and $\mathcal{G}'_\Xi(X, Y) \subseteq X$ iff $\mathcal{G}'_\Xi(X, Y) = X$ and $\mathcal{G}'_\Xi(Y, X) = Y$ iff (X, Y) is a four-valued supported model of Ξ.

4. *Note that in the hardness proof of item 3, the constructed pair was a three-valued supported model.* ◊

$\mathcal{O} = \mathcal{U}_\Xi$: *Let (ψ_1, ψ_2) be an instance of SAT-UNSAT. For convenience, we again assume w.l.o.g. that ψ_1 uses vocabulary P_1 and formula ψ_2 uses vocabulary P_2 with $P_1 \cap P_2 = \emptyset$. We construct an ADF $\Xi = (S, L, C)$ as follows:*

- *$S = P_1 \cup P_2 \cup \{z_1, z_2, d\}$ (where $z_1, z_2, d \notin P_1 \cup P_2$),*
- *$\varphi_p = p$ for all $p \in P_1 \cup P_2$,*
- *$\varphi_d = d$,*
- *$\varphi_{z_1} = \psi_1 \wedge d$ and $\varphi_{z_2} = \psi_2$.*

Now we show that for $X = \emptyset$ and $Y = S \setminus \{z_2\}$ we have $(X, Y) = \mathcal{U}_\Xi(X, Y)$ iff (ψ_1, ψ_2) is a yes instance of the SAT-UNSAT problem.

"only if": *Assume $\mathsf{Ver}_{\mathsf{com}}^{\mathcal{U}_\Xi}(X, Y)$ is true. Then due to the fact that $z_2 \notin \mathcal{U}_\Xi(X, Y)$, we have for each V with $\emptyset \subseteq V \subseteq Y$ that $V \not\models \varphi_{z_2} = \psi_2$ and since $P_2 \subseteq V$ it follows that ψ_2 is unsatisfiable. Since $z_1 \in Y$ it holds that φ_{z_1} is satisfiable and φ_{z_1} is satisfiable iff ψ_1 is satisfiable.*

[4]There are only two formulas for an empty vocabulary ($\top$ and $\bot$), and those two can be equivalently formulated with a non-empty vocabulary ($p \vee \neg p$ and $p \wedge \neg p$).

"if": *Assume ψ_1 is satisfiable and ψ_2 is unsatisfiable. Clearly then for V with $\emptyset \subseteq V \subseteq P_2$ we have $V \not\models \psi_2 = \varphi_{z_2}$, thus $z_2 \notin \mathcal{U}''_\Xi(X,Y)$. Further φ_{z_1} is satisfiable. In particular there exists a $V \subseteq P_1$ such that $V \models \psi_1$ and thus $V \cup \{d\} \models \varphi_{z_1}$ and since $X \subseteq (V \cup \{d\}) \subseteq Y$ it follows that $z_1 \in Y$. Since φ_{z_1} is not a tautology (e.g. $V \not\models \varphi_{z_1}$) we have that $z_1 \notin \mathcal{U}'_\Xi(X,Y)$. All remaining statements similarly have satisfiable but not tautological acceptance conditions, thus are in $\mathcal{U}''_\Xi(X,Y)$ but not in $\mathcal{U}'_\Xi(X,Y)$. This implies that (X,Y) is complete w.r.t. $\mathcal{U}_\Xi$.* □

Next, we analyze the complexity of verifying that a given pair is the approximate (ultimate) Kripke-Kleene semantics of an ADF Ξ, that is, the least fixpoint of $\mathcal{G}_\Xi$ ($\mathcal{U}_\Xi$). Interestingly, the membership part is the tricky one, where we encode the steps of the operator computation into propositional logic.

Theorem 12. *Let Ξ be an ADF, $X \subseteq Y \subseteq S$ and consider any operator $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$. Deciding $\mathsf{Ver}^{\mathcal{O}}_{\mathsf{grd}}(X,Y)$ is D_2^P-complete.*

Proof. *We begin the proof for $\mathcal{O} = \mathcal{G}_\Xi$.*

in D_2^P: *We provide a reduction to SAT-UNSAT. Assume that $S = \{s_1, \dots, s_n\}$ and set $P = \{t_i, u_i, b_{i,j} \mid 1 \le i, j \le n\}$. For each statement s_i, the propositional variable t_i indicates that s_i is true, while u_i indicates that s_i is undefined. Thus the truth values of the t_i and u_i determine a four-valued interpretation (T,U). The $b_{i,j}$ are used to guess parents that are needed to derive the acceptance of statement s_i in one operator application step; more precisely, $b_{i,j}$ indicates that s_j is a parent of s_i that is "needed" to infer u_i. By φ_i we denote the acceptance formula of s_i; by φ_i^t we denote φ_i where each s_j has been replaced by t_j; by φ_i^b we denote φ_i where each s_j has been replaced by $b_{i,j}$. Now define the formulas*

$$\phi_{T \subseteq U} = \bigwedge_{s_i \in S} (t_i \to u_i) \qquad (T,U) \text{ is a consistent pair}$$

$$\phi_{\le_i} = \bigwedge_{s_i \notin X} \neg t_i \wedge \bigwedge_{s_i \in Y} u_i \qquad (T,U) \le_i (X,Y)$$

$$\phi_{\ge_i} = \bigwedge_{s_i \in X} t_i \wedge \bigwedge_{s_i \notin Y} \neg u_i \qquad (T,U) \ge_i (X,Y)$$

$$\phi_= = \phi_{\le_i} \wedge \phi_{\ge_i} \qquad (T,U) = (X,Y)$$

$$\phi_{<_i} = \phi_{\le_i} \wedge \neg\phi_{\ge_i} \qquad (T,U) <_i (X,Y)$$

$$\phi_i^{\mathsf{2v}} = \bigwedge_{r_j \in par(s_i)} (u_j \to t_j) \qquad s_i \text{ has no undecided parents}$$

$$\phi_i^{?} = \bigwedge_{r_j \in par(s_i)} (t_j \to b_{i,j}) \qquad \text{guesses for } s_i \text{ are consistent with } T$$

$$\phi_{\mathsf{fpl}} = \bigwedge_{s_i \in S} (t_i \leftrightarrow (\varphi_i^t \wedge \phi_i^{\mathsf{2v}})) \qquad \mathcal{G}'_\Xi(T,U) = T$$

$$\phi_{\mathsf{fpu}} = \bigwedge_{s_i \in S} (u_i \leftrightarrow (\varphi_i^b \wedge \phi_i^{?})) \qquad \mathcal{G}''_\Xi(T,U) = U$$

$$\phi_{\mathsf{fp}} = \phi_{\mathsf{fpl}} \wedge \phi_{\mathsf{fpu}} \qquad \mathcal{G}_\Xi(T,U) = (T,U)$$

$$\psi_1 = \phi_{\mathsf{fp}} \wedge \phi_= \wedge \phi_{T \subseteq U} \qquad \mathcal{G}_\Xi(T,U) = (T,U) \text{ with } (T,U) = (X,Y) \text{ and } T \subseteq U$$

$$\psi_2 = \phi_{\mathsf{fp}} \wedge \phi_{<_i} \qquad \mathcal{G}_\Xi(T,U) = (T,U) \text{ with } (T,U) <_i (X,Y)$$

We claim that (1) ψ_1 is satisfiable iff (X, Y) is a consistent fixpoint of $\mathcal{G}_\Xi$, and (2) ψ_2 is satisfiable iff there is a fixpoint $(T, U) <_i (X, Y)$ of $\mathcal{G}_\Xi$. From this it follows that (ψ_1, ψ_2) is a positive instance of SAT-UNSAT iff (X, Y) is the Kripke-Kleene semantics of Ξ.

1. *ψ_1 is satisfiable iff (X, Y) is a consistent fixpoint of $\mathcal{G}_\Xi$.*

   "if": *Let $\mathcal{G}_\Xi(X, Y) = (X, Y)$ with $X \subseteq Y$. Define an interpretation $I \subseteq P$ as follows:*
   - *Set $t_i \in I$ iff $s_i \in X$ and $u_i \in I$ iff $s_i \in Y$.*
   - *Since $\mathcal{G}'_\Xi(Y, X) = Y$, we have for each $1 \leq i \leq n$ that $s_i \in Y$ iff there is a $B_i \subseteq par(s_i)$ with $B_i \models \varphi_i$ and $par(s_i) \setminus B_i \subseteq S \setminus X$. Now pick such a B_i for each $s_i \in S$ and set $b_{i,j} \in I$ iff $s_j \in B_i$.*

   We have to show $I \models \psi_1$. By definition $(T, U) = (X, Y)$ whence $I \models \phi_=$. Adding $X \subseteq Y$, it becomes clear that $I \models \phi_{T \subseteq U}$.
   For any $s_i \in S$, we have $I \models t_i$ iff $s_i \in X$ iff $s_i \in \mathcal{G}'_\Xi(X, Y)$ iff $X \models \varphi_i$ and $par(s_i) \cap Y \subseteq par(s_i) \cap X$ iff $I \models \varphi_i^t$ and $I \models \phi_i^{\mathsf{2v}}$ iff $I \models \varphi_i^t \wedge \phi_i^{\mathsf{2v}}$. Thus $I \models \phi_{\mathsf{fpl}}$.
   For any $s_i \in S$, we have $I \models u_i$ iff $s_i \in Y$ iff $s_i \in \mathcal{G}'_\Xi(Y, X)$ iff $B_i \models \varphi_i$ and $par(s_i) \setminus B_i \subseteq S \setminus X$ iff $I \models \varphi_i^b$ and $I \models \bigwedge_{s_j \in S} (\neg b_{i,j} \to \neg t_j)$ iff $I \models \varphi_i^b$ and $I \models \bigwedge_{s_j \in S} (t_j \to b_{i,j})$ iff $I \models \phi_{\mathsf{fpu}}$.
   Hence $I \models \phi_= \wedge \phi_{T \subseteq U} \wedge \phi_{\mathsf{fpl}} \wedge \phi_{\mathsf{fpu}}$ whence $I \models \psi_1$ and ψ_1 is satisfiable.

   "only if": *Let $I \subseteq P$ be such that $I \models \psi_1$. Define a three-valued pair (X, Y) and a sequence $B_1, \ldots, B_n$ by setting*
   - *$s_i \in X$ iff $t_i \in I$ and $s_i \in Y$ iff $u_i \in I$, and*
   - *$s_j \in B_i$ iff $b_{i,j} \in I$.*

   We have to show $\mathcal{G}_\Xi(X, Y) = (X, Y)$.
   We have $s_i \in \mathcal{G}'_\Xi(X, Y)$ iff $X \models \varphi_i$ and $par(s_i) \cap Y \subseteq par(s_i) \cap X$ iff $I \models \varphi_i^t$ and $I \models \phi_i^{\mathsf{2v}}$ iff $I \models \varphi_i^t \wedge \phi_i^{\mathsf{2v}}$ iff $I \models t_i$ (since $I \models \phi_{\mathsf{fpl}}$) iff $s_i \in X$. Hence $\mathcal{G}'_\Xi(X, Y) = X$.
   Similarly, we have $s_i \in Y$ iff $I \models u_i$ iff $I \models (\varphi_i^b \wedge \phi_i^?)$ (since $I \models \phi_{\mathsf{fpu}}$) iff $I \models \varphi_i^b$ and $I \models \phi_i^?$ iff $B_i \models \varphi_i$ and $par(s_i) \setminus B_i \subseteq S \setminus X$ iff $s_i \in \mathcal{G}'_\Xi(Y, X)$.
   Hence $Y = \mathcal{G}'_\Xi(Y, X)$ and in combination $\mathcal{G}_\Xi(X, Y) = (X, Y)$.

2. *ψ_2 is satisfiable iff there is a fixpoint $(T, U) <_i (X, Y)$ of $\mathcal{G}_\Xi$.*

   "if": *Let $(T, U) <_i (X, Y)$ with $\mathcal{G}_\Xi(T, U) = (T, U)$. We can define a two-valued interpretation $I \subseteq P$ as above with (T, U) playing the role of (X, Y). In an entirely analogous way we can then prove $I \models \phi_{\mathsf{fp}}$. It is straightforward to show that $(T, U) <_i (X, Y)$ implies $I \models \phi_{<_i}$.*

   "only if": *From an interpretation $I \models \psi_2$ we can define a pair (T, U) as above and can show that it is a fixpoint of $\mathcal{G}_\Xi$ with $(T, U) <_i (X, Y)$.*

D_2^P-hard: *This follows from the proof of Item 4 in Proposition 11: The three-valued supported model to verify there coincides with the Kripke-Kleene semantics of the constructed ADF.*

For $\mathcal{O} = \mathcal{U}_\Xi$ this result was shown already in [Brewka et al., 2013a, Theorem 6], but the proof was omitted due to space limitations. For sake of completeness we will present here an alternative proof which will be re-used later for query based reasoning. We first prove the following useful claim.

Claim. *Let Ξ be an ADF and $X \subseteq Y \subseteq S$. If it holds that*

1. *for each $x \in X$ there exists an I_x s.t. $X \subseteq I_x \subseteq Y$ and $I_x \models \varphi_x$;*
2. *for each $f \in S \setminus Y$ there exists an I_f s.t. $X \subseteq I_f \subseteq Y$ and $I_f \not\models \varphi_f$; and*

3. for each $y \in Y \setminus X$ there exists two I_y, I'_y s.t. $X \subseteq I_y, I'_y \subseteq Y$, $I_y \models \varphi_x$ and $I'_y \not\models \varphi_y$;

then lfp($\mathcal{U}_\Xi$) $\leq_i (X, Y)$. Further the grounded pair lfp($\mathcal{U}_\Xi$) satisfies these three properties.

Proof of the claim. Let (L, U) be the grounded pair of $\mathcal{U}_\Xi$, then it is straightforward to show that this pair satisfies all three properties, just recall that for each $s \in L$ we have $\varphi_s^{(L,U)}$ is tautological, for $s \in S \setminus U$ it is unsatisfiable and otherwise it is neither.

Assume (X, Y) satisfies all three properties, then we show by induction on $n \geq 1$ that $\mathcal{U}_\Xi^n(\emptyset, S) \leq_i (X, Y)$, with the usual meaning of iterative applications of operators, i.e. $\mathcal{U}_\Xi^n(X, Y) = \mathcal{U}_\Xi^{n-1}(\mathcal{U}_\Xi(X, Y))$. For $n = 1$ and $\mathcal{U}_\Xi^1(\emptyset, S) = (L_1, U_1)$ we have that if $s \in L_1$ then φ_s is a tautology, implying that $s \in X$, since otherwise there would exist a two-valued interpretation which does not satisfy φ_s. The case for $s \in S \setminus U_1$ is symmetric. Now assume the induction hypothesis $(L_n, U_n) = \mathcal{U}_\Xi^n(\emptyset, S) \leq_i (X, Y)$ and to show that $\mathcal{U}_\Xi^{n+1}(\emptyset, S) \leq_i (X, Y)$ holds consider $\mathcal{U}_\Xi^{n+1}(\emptyset, S) = (L_{n+1}, U_{n+1})$. If $s \in L_{n+1} \setminus L_n$ then $\varphi_s^{(L_n,U_n)}$ is tautological, which means that s must be in X. Similarly for the arguments set to false. This proves the claim. ◊

Now note that the pair in the claim, as well as all interpretations which show the three properties can be constructed in polynomial time w.r.t. to a given ADF. Additionally the evaluations for each such interpretation are computable in polynomial time w.r.t. the size of the ADF. Using this claim now we can guess a pair (X, Y) with $X \subseteq Y$ and $s \notin X$ and a set of interpretations as defined in the claim to show that s is not true in the grounded pair. This is achieved by checking if (X, Y) satisfies all three properties defined in the claim. If it does then s cannot be true in the grounded pair, since every argument set to true in the grounded pair is also true in the pair (X, Y). Checking if an argument is not false in the grounded pair is symmetric. This solves the complementary problem of deciding whether s is true/false in the grounded pair. To show that s is undefined in the grounded pair, we check if it is not true and not false in the grounded pair. Lastly, the verification problem simply consists of checking for each argument the corresponding truth value, which can be decided in D_2^P (NP for unknown values and otherwise coNP). Note that these checks can be done independently of each other.

For D_2^P-hardness, as for the approximate operator, consider the proof of Proposition 11 and the reduction shown there for $\mathcal{O} = \mathcal{U}_\Xi$. The constructed pair for that ADF is the grounded pair iff ψ_1 is satisfiable and ψ_2 is unsatisfiable. □

We next ask whether there exists a *non-trivial* admissible pair, that is, if at least one statement has a truth value other than unknown. Clearly, we can guess a pair and perform the coNP-check to show that it is admissible. The next result shows that this is also the best we can do.

Theorem 13. *Let Ξ be an ADF and consider any operator $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$. $\mathsf{Exists}^{\mathcal{O}}_{\mathsf{adm}}$ is Σ_2^P-complete.*

Proof. in Σ_2^P: *We guess a pair (X, Y) and verify that $X \subseteq Y$ and $(\emptyset, S) <_i (X, Y)$ in polynomial time, and $(X, Y) \leq_i \mathcal{O}(X, Y)$ using the NP oracle.*

Σ_2^P-hard: *We provide a reduction from the Σ_2^P-hard problem $QBF_{2,\exists}$-TRUTH. We show the proof for $\mathcal{O} = \mathcal{G}_\Xi$, as the proof for the ultimate operator is analogous. Let $\exists P \forall Q \psi$ be a QBF. We define an ADF Ξ as follows:*

- $S = P \cup \text{-}P \cup Q \cup \{z\}$ where $\text{-}P = \{\text{-}p \mid p \in P\}$,
- $\varphi_p = \neg z \wedge \neg\text{-}p$ for $p \in P$,
- $\varphi_{\text{-}p} = \neg z \wedge \neg p$ for $\text{-}p \in \text{-}P$,

- $\varphi_q = \neg q$ *for* $q \in Q$,
- $\varphi_z = \neg z \wedge \neg\psi$.

We show that D *has a non-trivial admissible pair iff* $\exists P \forall Q \psi$ *is true.*

"if": *Let* $M \subseteq P$ *be such that* $\psi^{(M,M)}$ *is a tautology. We define a consistent pair* (X, Y) *by setting*

- $X = M \cup \text{-}P \setminus M$ *and*
- $Y = X \cup Q$.

(X, Y) *is obviously non-trivial, since* $z \notin Y$. *Note that by definition* $p \in X$ *iff* $p \in M$, *and* $p \notin Y$ *iff* $p \notin M$, *whence* $\psi^{(M,M)} = \psi^{(X,Y)}$. *It remains to show that* (X, Y) *is admissible for* Ξ, *that is,* $X \subseteq \mathcal{G}'_\Xi(X, Y)$ *and* $\mathcal{G}'_\Xi(Y, X) \subseteq Y$.

$X \subseteq \mathcal{G}'_\Xi(X, Y)$:
- *Let* $p \in X$ *for* $p \in P$. *Then by definition* $\text{-}p \notin Y$ *and* $\varphi_p^{(X,Y)} = \neg\bot \wedge \neg\bot \equiv \top$ *whence* $p \in \mathcal{G}'_\Xi(X, Y)$.
- *Let* $\text{-}p \in X$ *for* $\text{-}p \in \text{-}P$. *Symmetric.*

$\mathcal{G}'_\Xi(Y, X) \subseteq Y$:
- *Let* $p \notin Y$ *for* $p \in P$. *Then by definition* $\text{-}p \in X$ *and* $\varphi_p^{(X,Y)} = \neg\bot \wedge \neg\top \equiv \bot$ *whence* $p \notin \mathcal{G}'_\Xi(Y, X)$.
- *Let* $\text{-}p \notin Y$ *for* $\text{-}p \in \text{-}P$. *Symmetric.*
- *Finally, we have* $\varphi_z^{(X,Y)} = \neg\bot \wedge \neg\psi^{(X,Y)} \equiv \neg\psi^{(X,Y)} = \neg\psi^{(M,M)}$. *Since* $\psi^{(M,M)}$ *is a tautology,* $\neg\psi^{(M,M)} = \neg\psi^{(X,Y)}$ *is unsatisfiable and* $y \notin \mathcal{G}'_\Xi(Y, X)$.

"only if": *Let* (X, Y) *be a non-trivial admissible pair for* Ξ. *We have to show that* $\exists P \forall Q \psi$ *is true. Define* $M = X \cap P$, *we show that* $\psi^{(M,M)}$ *is a tautology. We first observe that* $Q \subseteq Y \setminus X$ *by their acceptance conditions and since* (X, Y) *is admissible: if* $q \in X$, *then* $q \notin \mathcal{G}'_\Xi(X, Y)$; *if* $q \notin Y$, *then* $q \in \mathcal{G}'_\Xi(Y, X)$. *We next show* $z \notin Y$.

By the presumption that (X, Y) *is non-trivial, we get that either (1)* $X \neq \emptyset$ *or (2)* $Y \subsetneq S$.

1. $X \neq \emptyset$.
   (a) $z \in X$. *Then* $\varphi_z^{(X,Y)} = \neg\top \wedge \psi^{(X,Y)} \equiv \bot$ *and* $z \notin \mathcal{G}'_\Xi(X, Y)$, *that is,* (X, Y) *is not admissible. Contradiction.*
   (b) $p \in X$. *Then by admissibility* $p \in \mathcal{G}'_\Xi(X, Y)$, *that is, we have the equivalence* $\varphi_p^{(X,Y)} = (\neg z \wedge \neg\text{-}p)^{(X,Y)} \equiv \top$, *whence* $z \notin Y$ *and* $\text{-}p \notin Y$.
   (c) $\text{-}p \in X$. *Symmetric.*
2. $Y \subsetneq S$.
   (a) $z \in S \setminus Y$. *This is what we want to show.*
   (b) $p \in S \setminus Y$. *By admissibility* $p \notin \mathcal{G}'_\Xi(Y, X)$ *and the partially evaluated acceptance formula* $\varphi_p^{(X,Y)} = (\neg z \wedge \neg\text{-}p)^{(X,Y)}$ *must be unsatisfiable. Since* $z \notin X$, *we get* $\text{-}p \in X$. *By item 1c above, we get* $z \notin Y$.
   (c) $\text{-}p \in S \setminus Y$. *Symmetric.*

   Hence $z \notin Y$. *Since* (X, Y) *is admissible,* $z \notin \mathcal{G}'_\Xi(Y, X)$. *Thus the partially evaluated acceptance formula* $\varphi_z^{(X,Y)} = \neg\bot \wedge \neg\psi^{(X,Y)}$ *is unsatisfiable, that is,* $\neg\psi^{(X,Y)}$ *is unsatisfiable and* $\psi^{(X,Y)} = \psi^{(M,M)}$ *is a tautology.* □

Lemma 6 implies the same complexity for the existence of non-trivial complete and preferred pairs.

Corollary 14. *Let Ξ be an ADF, $\sigma \in \{com, pre\}$ and consider any operator $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$. $\mathsf{Exists}_\sigma^{\mathcal{O}}$ is Σ_2^P-complete.*

By corollary to Theorem 12, the existence of a non-trivial grounded pair can be decided in D_2^P by testing whether the trivial pair $(\emptyset, S)$ is (not) a fixpoint of the relevant operator. The following result shows that this bound can be improved.

Proposition 15. *Let Ξ be an ADF and consider any operator $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$. $\mathsf{Exists}_{\mathsf{grd}}^{\mathcal{O}}$ is coNP-complete.*

Proof. *Obviously, Ξ has a non-trivial approximate grounded semantics iff the trivial pair $(\emptyset, S)$ is not a fixpoint of $\mathcal{G}_\Xi$, so we show that the co-problem (deciding whether $\mathcal{G}_\Xi(\emptyset, S) = (\emptyset, S)$) is NP-complete.*

in NP: *We have that $\mathcal{G}_\Xi(\emptyset, S) = (\emptyset, S)$ iff $\emptyset \subseteq \mathcal{G}'_\Xi(\emptyset, S) \subseteq \emptyset$ and $S \subseteq \mathcal{G}'_\Xi(S, \emptyset) \subseteq S$. So mainly we have to verify $\mathcal{G}'_\Xi(\emptyset, S) \subseteq \emptyset$ and $S \subseteq \mathcal{G}'_\Xi(S, \emptyset)$. By Proposition 4, the first part can be decided in P (item 2) and the second part in NP (item 6).*

NP-hard: *We give a reduction from SAT. Let ψ be a propositional formula over vocabulary P. Define an ADF $D = (S, L, C)$ with $S = P \cup \{z\}$ for $z \notin P$ and $\varphi_p = p$ for $p \in P$ and $\varphi_z = z \wedge \psi$. It is readily verified that by definition every statement has a parent that is undecided in $(\emptyset, S)$ and thus $\mathcal{G}'_\Xi(\emptyset, S) = \emptyset$. Furthermore, $P \subseteq \mathcal{G}'_\Xi(S, \emptyset)$ is easy to show. Thus $S \subseteq \mathcal{G}'_\Xi(S, \emptyset)$ iff $z \in \mathcal{G}'_\Xi(S, \emptyset)$ iff there is an $M \subseteq S$ with $par(z) \setminus M \subseteq S \setminus \emptyset$ and $M \models \varphi_z$ iff there is an $M \subseteq S$ with $M \models \varphi_z$ iff $\varphi_z = z \wedge \psi$ is satisfiable iff ψ is satisfiable.*

For $\mathcal{O} = \mathcal{U}_\Xi$, the proof is analogous to the one above – we show NP-completeness of the complementary problem.

in NP: *We have to verify $\mathcal{U}'_\Xi(\emptyset, S) \subseteq \emptyset$ and $S \subseteq \mathcal{U}''_\Xi(\emptyset, S)$. By Proposition 5(2) and Proposition 4(6), this can be done in NP.*

NP-hard: *The construction is the same as in Proposition 15.* ☐

Using the result for existence of non-trivial admissible pairs, the verification complexity for the preferred semantics is straightforward to obtain, similarly as in the case of AFs [Dimopoulos and Torres, 1996].

Proposition 16. *Let Ξ be an ADF, $X \subseteq Y \subseteq S$ and consider any $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$. $\mathsf{Ver}_{\mathsf{pre}}^{\mathcal{O}}(X, Y)$ is Π_2^P-complete.*

Proof. in Π_2^P: *To show that (X, Y) is not preferred, we guess a pair (M, N) with $(X, Y) <_i (M, N)$ and use the NP oracle to show that (M, N) is a complete pair (which can be done in D_2^P).*

Π_2^P-hard: *Consider the complementary problem, that is, deciding whether a given pair is not a preferred pair. Even for the special case of the pair $(\emptyset, S)$, Theorem 13 shows that this problem is Σ_2^P-hard.* ☐

Considering query reasoning we now show that on general ADFs credulous reasoning with respect to admissibility is harder than on AFs. By Lemma 7, the same lower bound holds for complete and preferred semantics.

Proposition 17. *Let Ξ be an ADF, $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$ be an operator and $s \in S$. $\mathsf{Cred}^{\mathcal{O}}_{\mathsf{adm}}(s)$ is Σ^P_2-complete.*

Proof. Membership is given by Theorem 8. Hardness is shown by a reduction from the Σ^P_2-hard problem $QBF_{2,\exists}$-TRUTH. Let $\exists P \forall Q \psi$ be a QBF. We define an ADF Ξ as follows:

- $S = P \cup Q \cup \{f\}$,
- $\varphi_p = p$ for $p \in P$,
- $\varphi_q = \neg q$ for $q \in Q$,
- $\varphi_f = \neg f \wedge \neg \psi$.

We now show that there exists an admissible pair (X, Y) in Ξ w.r.t. $\mathcal{O}$ with $f \in S \setminus Y$ iff $\exists P \forall Q \psi$ is true. Note that for any admissible pair we have that f is not set to true, since then its acceptance condition evaluates to false.

"if": Assume the QBF is valid. Then there exists a $P' \subseteq P$ such that for any $Q' \subseteq Q$ we have $P' \cup Q' \models \psi$. We now show that $(P', P' \cup Q)$ is admissible in Ξ for $\mathcal{O}$. Since for any $p \in P'$ we have that $\{p\} \models \varphi_p$ it follows that $P' \subseteq \mathcal{G}'_\Xi(P', P' \cup Q)$ and $P' \subseteq \mathcal{U}'_\Xi(P', P' \cup Q)$. Further $\emptyset \models \varphi_q$ for $q \in Q$, thus $Q \subseteq \mathcal{G}''_\Xi(P', P' \cup Q)$. Lastly, $R \not\models \varphi_f$ for $P' \subseteq R \subseteq P' \cup Q$ since $R \models \psi$, hence $f \notin \mathcal{G}''_\Xi(P', P' \cup Q)$.

"only if": Assume that there exists an admissible pair (X, Y) in Ξ w.r.t. $\mathcal{O}$ such that $f \in S \setminus Y$. First we show that $Q \subseteq (Y \setminus X)$. Suppose the contrary, then (X, Y) would not be admissible w.r.t. $\mathcal{O}$, since if there is a $q \in Q$ and $q \in X$, then $q \notin \mathcal{G}''_\Xi(X, Y)$. Similarly $q \in S \setminus Y$, implies that q is in the new lower bound for both operators in $\mathcal{O}$. We now show that $X \cup Q' \models \psi$ for any $Q' \subseteq Q$, thus implying that ψ is valid (note that $X \subseteq P$). Suppose there exists a $Q' \subseteq Q$ such that $X \cup Q' \not\models \psi$, then $X \cup Q' \models \varphi_f$ and $X \subseteq (X \cup Q') \subseteq Y$ and thus $f \in \mathcal{G}''_\Xi(X, Y)$, implying that (X, Y) is not admissible w.r.t. $\mathcal{O}$, which is a contradiction. □

For credulous and skeptical reasoning with respect to the grounded semantics, we first observe that the two coincide since there is always a unique grounded pair. Furthermore, a statement s is true in the approximate grounded pair iff s is true in the least fixpoint (of $\mathcal{G}_\Xi$) iff s is true in all fixpoints iff there is no fixpoint where s is unknown or false. This condition can be encoded in propositional logic and leads to the next result. For the ultimate operator we can use results for the verification problem [Brewka et al., 2013a, Theorem 6]. Briefly put, the problem is in coNP since the NP hardness comes from verifying that certain arguments are undefined in the ultimate grounded pair, which is not needed for credulous reasoning. For coNP-hardness the proof of [Brewka and Woltran, 2010, Proposition 13] can be easily adapted.

Proposition 18. *Let Ξ be an ADF, $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$ and $s \in S$. Both $\mathsf{Cred}^{\mathcal{O}}_{\mathsf{grd}}(s)$ and $\mathsf{Skept}^{\mathcal{O}}_{\mathsf{grd}}(s)$ are* coNP*-complete.*

Proof. For showing the membership result for $\mathcal{O} = \mathcal{U}_\Xi$, see the proof of Theorem 12, in particular the claim proved there. For hardness the proof for both operators is the same. For showing the results for $\mathcal{O} = \mathcal{G}_\Xi$ consider the following proof.

in coNP: We reduce to unsatisfiability checking in propositional logic. Let $\Xi = (S, L, C)$ be an ADF with $S = \{s_1, \ldots, s_n\}$ and assume we want to verify that s_k is true in the grounded pair of Ξ for some $1 \leq k \leq n$. Set $P = \{t_i, u_i, b_{i,j} \mid 1 \leq i, j \leq n\}$. For each statement s_i, the propositional variable t_i indicates that s_i is true, while u_i indicates that s_i is undefined.

Thus the truth values of the t_i and u_i determine a four-valued interpretation (T, U). The $b_{i,j}$ are used to guess parents that are needed to derive the acceptance of statement s_i in one operator application step; more precisely, $b_{i,j}$ indicates that s_j is a parent of s_i that is "needed" to infer u_i. By φ_i we denote the acceptance formula of s_i; by φ_i^t we denote φ_i where each s_j has been replaced by t_j; by φ_i^b we denote φ_i where each s_j has been replaced by $b_{i,j}$. Now define the formulas

$$
\begin{aligned}
\phi_{T\subseteq U} &= \bigwedge_{s_i\in S} (t_i \to u_i) && (T,U) \text{ is a consistent pair}\\
\phi_i^{\mathsf{2v}} &= \bigwedge_{r_j\in par(s_i)} (u_j \to t_j) && s_i \text{ has no undecided parents}\\
\phi_i^{?} &= \bigwedge_{r_j\in par(s_i)} (t_j \to b_{i,j}) && \text{guesses for } s_i \text{ are consistent with } T\\
\phi_{\mathsf{fpl}} &= \bigwedge_{s_i\in S} (t_i \leftrightarrow (\varphi_i^t \wedge \phi_i^{\mathsf{2v}})) && \mathcal{G}'_\Xi(T,U) = T\\
\phi_{\mathsf{fpu}} &= \bigwedge_{s_i\in S} (u_i \leftrightarrow (\varphi_i^b \wedge \phi_i^{?})) && \mathcal{G}''_\Xi(T,U) = U\\
\phi_{\mathsf{cfp}} &= \phi_{\mathsf{fpl}} \wedge \phi_{\mathsf{fpu}} \wedge \phi_{T\subseteq U} && \mathcal{G}_\Xi(T,U) = (T,U) \text{ and } T\subseteq U\\
\phi'_{\mathsf{fp}} &= \phi_{\mathsf{fp}}[p/p' : p\in P] && \text{renamed copy}\\
\psi &= (\phi_{\mathsf{fp}} \wedge \neg t_k) \vee (\phi'_{\mathsf{fp}} \wedge \neg u'_k)
\end{aligned}
$$

(Observe that this is basically the construction from Theorem 12.) We claim that ψ is unsatisfiable iff there is no consistent fixpoint where s_k is unknown or false.

1. *$\phi_{\mathsf{fp}} \wedge \neg t_k$ is unsatisfiable iff there is no consistent fixpoint where s_k is unknown:*

   "if": *Let $\phi_{\mathsf{fp}} \wedge \neg t_k$ be satisfiable. Then there is an interpretation $I \subseteq P$ such that $I \models \phi_{\mathsf{fp}}$ and $I \not\models t_k$. As in the proof of Theorem 12 we can construct a consistent pair (X, Y) and show that it is a fixpoint of $\mathcal{G}_\Xi$ with $s_k \in Y \setminus X$.*

   "only if": *Let $X \subseteq Y \subseteq S$ such that $\mathcal{G}_\Xi(X, Y) = (X, Y)$ and $s_k \in Y \setminus X$. As in the proof of Theorem 12 we can construct an interpretation $I \subseteq P$ such that $I \models \phi_{\mathsf{fp}}$ and $I \not\models t_k$.*

2. *$\phi_{\mathsf{fp}} \wedge \neg u_k$ is unsatisfiable iff there is no consistent fixpoint where s_k is false: Similar.*

coNP-hard: *Let ψ be a propositional formula over vocabulary P. Define the ADF $D = (S, L, C)$ with $S = P \cup \{z\}$, $\varphi_p = \neg\psi$ for $p \in P$, and $\varphi_z = \bigwedge_{p\in P} \neg p$. We show that z is true in the grounded semantics of $\mathcal{G}_D$ iff ψ is a tautology.*

"if": *Let ψ be a tautology. Then $\neg\psi$ is unsatisfiable and $p \notin \mathcal{G}''_D(\emptyset, S)$ for all $p \in P$. Obviously φ_z is satisfiable whence $z \in \mathcal{G}''_D(\emptyset, S)$. Thus $\mathcal{G}''_D(\emptyset, S) = \{z\}$. Furthermore $\mathcal{G}'_D(\emptyset, S) = \emptyset$ since all statements have undecided parents. Thus $\mathcal{G}_D(\emptyset, S) = (\emptyset, \{z\})$. Now since z does not occur in the acceptance formula of z, it is clear that $\mathcal{G}_D(\emptyset, \{z\}) = (\{z\}, \{z\}) = \mathcal{G}_D(\{z\}, \{z\})$. Thus z is true in the grounded semantics of $\mathcal{G}_D$.*

"only if": *Let $lfp(\mathcal{G}_D) = (X, Y)$ and $z \in X$. By the acceptance condition of z and the fact that (X, Y) is a fixpoint of $\mathcal{G}_D$ we get $P \cap Y = \emptyset$. Since $X \subseteq Y$ we have $(X, Y) = (\{z\}, \{z\})$. Assume to the contrary that ψ is not a tautology. Then $\neg\psi$ is satisfiable and $P \subseteq Y = \mathcal{G}''_D(\emptyset, S)$. Contradiction.* □

Regarding skeptical reasoning for the remaining semantics we need only show the results for complete and preferred semantics, in all other cases the complexity coincides with credulous reasoning or is trivial. For complete semantics it is easy to see that skeptical reasoning coincides with skeptical reasoning under grounded semantics, since the grounded pair is the $\leq_i$-least complete pair.

Corollary 19. *Let Ξ be an ADF, $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$ and $s \in S$. $\mathsf{Skept}^{\mathcal{O}}_{\mathsf{com}}(s)$ is coNP-complete.*

Similar to reasoning on AFs, we step up one level of the polynomial hierarchy by changing from credulous to skeptical reasoning with respect to preferred semantics, which makes skeptical reasoning under preferred semantics particularly hard. We apply proof ideas by Dunne and Bench-Capon [2002] to prove Π^P_3-hardness.

Theorem 20. *Let Ξ be an ADF, $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$ and $s \in S$. $\mathsf{Skept}^{\mathcal{O}}_{\mathsf{pre}}(s)$ is Π^P_3-complete.*

Proof. *Membership is given by Theorem 8. Hardness is shown by a reduction from the Π^P_3-hard problem $QBF_{3,\forall}$-TRUTH. Let $\forall P \exists Q \forall R \psi$ be a QBF. We define an ADF Ξ as follows:*

- $S = P \cup Q \cup \text{-}Q \cup \{f\}$, *with* $\text{-}Q = \{\text{-}q \mid q \in Q\}$,
- $\varphi_p = p$ *for* $p \in P$,
- $\varphi_q = \neg f \wedge \neg\text{-}q$ *for* $q \in Q$,
- $\varphi_{\text{-}q} = \neg f \wedge \neg q$ *for* $\text{-}q \in \text{-}Q$,
- $\varphi_r = \neg r$ *for* $r \in R$,
- $\varphi_f = \neg f \wedge \neg \psi$.

We now show that all preferred pairs (X, Y) in Ξ w.r.t. $\mathcal{O}$ have $f \in S \setminus Y$ iff $\forall P \exists Q \forall R \psi$ is true. First observe some helpful facts. For any preferred pair (X, Y) in Ξ w.r.t. $\mathcal{O}$ we have $P \cap (Y \setminus X) = \emptyset$, since otherwise it would not be $\leq_i$-maximal. Suppose the contrary, i.e. (X, Y) is preferred and $p \in P$ and $p \in Y \setminus X$. We know that $(X, Y) \leq_i \mathcal{O}(X, Y)$ due to admissibility of the pair and trivially it holds that $(X, Y) \leq_i (X \cup \{p\}, Y)$. Due to monotonicity of $\mathcal{O}$ w.r.t. $\leq_i$ we know that $\mathcal{O}(X, Y) \leq_i \mathcal{O}(X \cup \{p\}, Y)$. This implies that $(X, Y) \leq_i \mathcal{O}(X \cup \{p\}, Y)$. This means that $X \subseteq \mathcal{O}'(X \cup \{p\}, Y)$ and $\mathcal{O}''(X \cup \{p\}, Y) \subseteq Y$. Further we can deduce that $p \in \mathcal{G}'_\Xi(X \cup \{p\}, Y)$ as well as $p \in \mathcal{U}'_\Xi(X \cup \{p\}, Y)$, hence $X \cup \{p\} \subseteq \mathcal{O}'(X \cup \{p\}, Y)$ and thus $(X \cup \{p\}, Y)$ is admissible w.r.t. $\mathcal{O}$ and (X, Y) is not preferred, since clearly $(X, Y) <_i (X \cup \{p\}, Y)$. The case where p is set to false is symmetric.

Suppose that f is set to true in a preferred pair (X, Y), i.e. $f \in X$, then for any $Q' \subseteq Q$ and $R' \subseteq R$ we have $X \cup Q' \cup R' \not\models \varphi_f$, hence f cannot be in X. If $f \in Y \setminus X$, then it holds that $q, \text{-}q \in Y \setminus X$ for any $q \in Q$ and $\text{-}q \in \text{-}Q$, or in other words, if f is undefined in the preferred pair then also all statements in $Q \cup \text{-}Q$ are undefined. Assume (X, Y) is an admissible pair w.r.t. $\mathcal{O}$, then if f is undefined in this pair, for no $q \in Q$ and $\text{-}q \in \text{-}Q$ it holds that q or $\text{-}q$ is true in the admissible pair. For $\mathcal{O} = \mathcal{G}_\Xi$ this is immediate, since this operator requires that all parents of $q, \text{-}q$ ($\{f, \text{-}q\}$ and $\{f, q\}$ respectively) are not undefined. For $\mathcal{O} = \mathcal{U}_\Xi$ the formula $\varphi_q^{(X,Y)}$ is not tautological if f is undefined (similarly for $\text{-}q$). Further for q ($\text{-}q$) it holds that its truth value is false in an admissible pair only if $\text{-}q$ is true (q is true). This is easy to see, since f is never true in an admissible pair, hence $\text{-}q$ (q) must be true. This implies that if f is undefined, both q and $\text{-}q$ are undefined.

"if": *Assume that the formula is valid. We now have to show that in all preferred pairs w.r.t.* $\mathcal{O}$ *we have that* f *is not in the upper bound. For any* $P' \subseteq P$ *there is a preferred pair* (X, Y) *with* $P' \subseteq X$ *(because* (P', S) *is admissible w.r.t.* $\mathcal{O}$*) and since* $P \cap (Y \setminus X) = \emptyset$ *it follows that in the preferred pair* P' *are set to true and* $P \setminus P'$ *to false. For such a* P' *we know that there is a* $Q' \subseteq Q$ *such that for any* $R' \subseteq R$ *we have* $P' \cup Q' \cup R' \models \psi$*. We now show there is a preferred pair* (X', Y') *with* $X' = P' \cup Q' \cup \text{-}(Q \setminus Q')$ *and* $Y' = X' \cup R$ *in* Ξ *w.r.t.* $\mathcal{O}$*. We set* $\text{-}(Q \setminus Q') = \{\text{-}q \mid q \in (Q \setminus Q')\}$*. It is easy to see that* $P' \subseteq \mathcal{G}'_\Xi(X', Y')$ *and also for the ultimate operator. Likewise since* $f \in S \setminus Y$ *also* $Q' \subseteq \mathcal{G}'_\Xi(X', Y')$ *and also for* $\mathcal{U}'_\Xi$ *as well as for the negated statements. Similarly we can show that* $P \setminus P'$ *is not in the upper bound, as well as for* $Q \setminus Q'$ *and the negated statements. The statements* $r \in R$ *are always undecided for both operators. To see that* $f \notin \mathcal{G}''_\Xi(X', Y')$*, we know that* $P' \cup Q' \cup R' \models \psi$ *by assumption. Thus* (X', Y') *is indeed a preferred pair w.r.t.* $\mathcal{O}$*. Now we know that there exists a preferred pair which sets* P' *to true and* f *to false. To see that there is no preferred pair* (X'', Y'') *which sets* P' *to true and* $P \setminus P'$ *to false such that* $f \in Y''$*, it suffices to show that for this case then also* $Q, \text{-}Q \subseteq Y'' \setminus X''$*, which holds since* f *is undecided and then* (X'', Y'') *is smaller w.r.t.* $\leq_i$ *to the preferred pair* (X', Y') *and hence cannot be a preferred pair in this ADF. Summarizing, all preferred pairs set a* $P' \subseteq P$ *to true and* $P \setminus P'$ *to false and for each such choice there exists a preferred pair setting* f *to false. Further if for such a choice a preferred pair exists with* f *set to false, we know that there is no preferred pair with the same assignment to the statements in* P *and setting* f *not to false. Thus any preferred pair sets* f *to false and hence* f *is skeptically rejected w.r.t.* $\mathcal{O}$*.*

"only if": *Assume that in any preferred pair w.r.t.* $\mathcal{O}$ *we have that* $f \in S \setminus Y$*. As in the "if" direction we know that for any* $P' \subseteq P$ *there exists a preferred pair* (X, Y) *with* $P \cap (Y \setminus X) = \emptyset$ *and* $P' \subseteq X$*. Since* $f \in S \setminus Y$ *we have that in any such preferred pair also* $(Q \cup \text{-}Q) \cap (Y \setminus X) = \emptyset$*, since otherwise it would not be maximal w.r.t.* $\leq_i$ *by a similar argument as above. Further we know that* $R \subseteq Y \setminus X$*. This means that there exists a* $Q' \subseteq Q$ *such that* $(P' \cup Q') \subseteq X$*. Thus for any* $R' \subseteq R$ *we have* $P' \cup Q' \cup R' \not\models \varphi_f$*, since* $f \in S \setminus Y$ *and hence* $P' \cup Q' \cup R' \models \psi$*. Since for any* $P' \subseteq P$ *such a preferred pair exists, the QBF is valid.* ☐

4.1 Two-valued semantics

The complexity results we have obtained so far might lead the reader to ask why we bother with the approximate operator $\mathcal{G}_\Xi$ at all: the ultimate operator $\mathcal{U}_\Xi$ is at least as precise and for all the semantics considered up to now, it has the same computational costs. We now show that for the verification of two-valued stable models, the operator for the upper bound plays no role and therefore the complexity difference between the lower bound operators for approximate (in P) and ultimate (coNP-hard) semantics comes to bear.

For the ultimate two-valued stable semantics, Brewka et al. [2013a] already have some complexity results: model verification is in D_2^P (see Proposition 8), and model existence is Σ_2^P-complete (see Theorem 9). We will show next that we can do better for the approximate version, with this deterministic polytime decision procedure for verifying that a given set $X \subseteq Y$ of statements is the least fixpoint of $\mathcal{G}'_\Xi(\cdot, Y)$.

Proposition 21. *Let* Ξ *be an ADF and* $X \subseteq Y \subseteq S$*. Verifying that* X *is the least fixpoint of* $\mathcal{G}'_\Xi(\cdot, Y)$ *is in* P*.*

Proof. We provide the following polynomial-time decision procedure with input Ξ, X, Y*.*

1. *Set* $i = 0$ *and* $X_0 = \emptyset$*.*

2. *For each statement* $s \in S$*, do the following:*
   (a) *If* $par(s) \cap (Y \setminus X_i) = \emptyset$ *and* $C_s(par(s) \cap X_i) = \mathbf{t}$*, then set* $s \in X_{i+1}$*.*
3. *If* $X_{i+1} = X_i = X$ *then return "Yes".*
4. *If* $X_{i+1} = X_i \subsetneq X$ *then return "No".*
5. *If* $X_{i+1} \not\subseteq X$ *then return "No".*
6. *Increment* i *and go to step 2.*

Overall, the loop between steps 2 and 6 is executed at most $|S|$ *times, since* $X_i \subseteq X_{i+1}$ *for all* $i \in \mathbb{N}$ *and we can add at most all statements one by one. In each execution of the loop, step 2a is executed* $|S|$ *times. The conditions of step 2a, in particular* $par(s) \cap X_i \models \varphi_s$*, can be verified in polynomial time.*

It remains to show that X *is the least fixpoint of* $\mathcal{G}'_\Xi(\cdot, Y)$ *iff the procedure returns "Yes".*

"if": *Assume the procedure returned "Yes" on input* Ξ, X, Y*.*

- X *is a fixpoint of* $\mathcal{G}'_\Xi(\cdot, Y)$*, that is,* $\mathcal{G}'_\Xi(X, Y) = X$*:*

  "$\subseteq$": *Let* $s \in \mathcal{G}'_\Xi(X, Y)$*. Then there is a* $B \subseteq X \cap par(s)$ *such that* $C_s(B) = \mathbf{t}$ *and* $par(s) \setminus B \subseteq S \setminus Y$*. As in the proof of Proposition 4, we get that* $B = X \cap par(s)$*,* $C_s(par(s) \cap X) = \mathbf{t}$ *and* $par(s) \cap (Y \setminus X) = \emptyset$*. Since the procedure answered "Yes", there was an* $i \in \mathbb{N}$ *with* $X_{i+1} = X_i = X$*. From step 2a of the procedure, we know that* $par(s) \cap (Y \setminus X_i) = \emptyset$ *and* $C_s(par(s) \cap X_i) = \mathbf{t}$ *means that* $s \in X_{i+1} = X$*.*

  "$\supseteq$": *Let* $s \in X$*. Since the procedure answered "Yes", there was an* $i \in \mathbb{N}$ *with* $X_{i+1} = X_i = X$*. Now* $s \in X_{i+1}$ *by step 2a of the procedure means that* $par(s) \cap (Y \setminus X_i) = \emptyset$ *and* $C_s(par(s) \cap X_i) = \mathbf{t}$*. Thus there exists a* $B = par(s) \cap X$ *with* $C_s(B) = \mathbf{t}$ *and* $par(s) \setminus B \subseteq S \setminus Y$*, and* $s \in \mathcal{G}'_\Xi(X, Y)$*.*
- X *is the least fixpoint: Assume to the contrary that there is some* $X' \subsetneq X$ *that is a fixpoint of* $\mathcal{G}'_\Xi(\cdot, Y)$*. But then step 4 of the procedure would have detected* $X_{i+1} = X_i = X' \subsetneq X$ *and returned "No", contradiction.*

"only if": *Let* X *be the least fixpoint of* $\mathcal{G}'_\Xi(\cdot, Y)$ *and assume to the contrary that the procedure answered "No".*

- *The procedure answered "No" in step 4. By the argument above, we can show that there is a fixpoint* $X' \subsetneq X$*, contradiction.*
- *The procedure answered "No" in step 5. We have* $X_{i+1} \not\subseteq X$ *for some* $i \in \mathbb{N}$*, that is, there is some* $s \in X_{i+1}$ *with* $s \notin X$*. Since* $s \in X_{i+1}$*, we have* $par(s) \cap (Y \setminus X_i) = \emptyset$ *and* $C_s(par(s) \cap X_i) = \mathbf{t}$*. Since the procedure did not terminate with* X_i *already, we know that* $X_i \subseteq X$*. Therefore,* $par(s) \cap (Y \setminus X) = \emptyset$ *and* $C_s(par(s) \cap X) = \mathbf{t}$*. This means* $s \in \mathcal{G}'_\Xi(X, Y) = X$*. Contradiction.* □

In particular, the procedure can decide whether Y is the least fixpoint of $\mathcal{G}'_\Xi(\cdot, Y)$, that is, whether (Y, Y) is a two-valued stable model of $\mathcal{G}_\Xi$. This yields the next result.

Theorem 22. *Let* Ξ *be an ADF and* $X \subseteq S$*. 1.* $\mathsf{Ver}^{\mathcal{G}_\Xi}_{\mathsf{2st}}(X, X)$ *is in* P*. 2.* $\mathsf{Exists}^{\mathcal{G}_\Xi}_{\mathsf{2st}}$ *is* NP*-complete.*

Proof. 1. *We have to verify that* X *is the least fixpoint of the operator* $\mathcal{G}'_\Xi(\cdot, X)$*, which can be done in polynomial time by Proposition 21.*

2. *Deciding whether Ξ has a two-valued stable model is* NP*-complete:*

   in NP: *To decide whether there is a two-valued stable model, we guess a set $X \subseteq S$ and verify as above that (X, X) is indeed a two-valued stable model.*

   NP-hard: *Carries over from AFs.* □

The hardness direction of the second part is clear since the respective result from stable semantics of abstract argumentation frameworks carries over.

Brewka et al. [2013a] showed that $\mathsf{Ver}_{2\mathsf{st}}^{\mathcal{U}_\Xi}$ is in D_2^P (Proposition 8). We can improve that upper bound to coNP: the proof is not that trivial, but basically the operator for the upper bound (contributing the NP part) is not really needed. Using the complexity of the lower revision operator $\mathcal{U}'_\Xi$, we can even show completeness for coNP.

Proposition 23. *Let Ξ be an ADF and $X \subseteq S$. $\mathsf{Ver}_{2\mathsf{st}}^{\mathcal{U}_\Xi}(X, X)$ is* coNP*-complete.*

Proof. in coNP: *Given an ADF $\Xi = (S, L, C)$ and a set $M \subseteq S$ we first construct the reduct Ξ^M in polynomial time. Now M is an ultimate two-valued stable model of Ξ iff all statements in M are true in the grounded semantics of Ξ^M. We will show that the co-problem (there is an $s \in M$ that is false or undecided in the grounded semantics of Ξ^M) is in* NP*. To this end consider the claim shown in the proof of Theorem 12 for the ultimate semantics. We guess two sets $X \subseteq Y \subseteq M$ and a statement $s \in M \setminus X$. Furthermore we guess witnesses to verify that (X, Y) satisfies the presumptions in the claim, which shows that $lfp(\mathcal{U}_\Xi) \leq_i (X, Y)$ and thus s is not true in the grounded pair of Ξ^M. Note that for each statement we need at most two two-valued interpretations over S as witnesses, which can easily be constructed in polynomial time.*

coNP-hard: *Let ψ be a propositional formula over a vocabulary P. We define an ADF D over statements P with $\varphi_p = \psi$ for all $p \in P$. Now P is an ultimate two-valued stable model of D iff P is the least fixpoint of $\mathcal{U}'_D(\cdot, P)$ iff $\mathcal{U}'_D(\emptyset, P) = P = \mathcal{U}'_D(P, P)$ iff $p \in \mathcal{U}'_D(\emptyset, P)$ for all $p \in P$ iff $\varphi_p^{(\emptyset, P)}$ is a tautology iff $\psi^{(\emptyset, P)}$ is a tautology iff ψ is a tautology.* □

We now turn to the credulous and skeptical reasoning problems for the two-valued semantics. We first recall that a two-valued pair (X, X) is a supported model (or model for short) of an ADF Ξ iff $\mathcal{G}_\Xi(X, X) = (X, X)$. Thus it could equally well be characterized by the two-valued operator by saying that X is a model iff $G_\Xi(X) = X$. Now since $\mathcal{U}_\Xi$ is the ultimate approximation of G_Ξ, also $\mathcal{U}_\Xi(X, X) = (X, X)$ in this case. Rounding up, this recalls that approximate and ultimate two-valued supported models coincide. Hence we get the following results for reasoning with this semantics.

Corollary 24. *Let Ξ be an ADF, $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$ be an operator and $s \in S$. The problem $\mathsf{Cred}_{2\mathsf{su}}^{\mathcal{O}}(s)$ is* NP*-complete; $\mathsf{Skept}_{2\mathsf{su}}^{\mathcal{O}}(s)$ is* coNP*-complete.*

Proof. The membership parts are clear since $\mathsf{Ver}_{2\mathsf{su}}^{\mathcal{O}}$ is in P*. Hardness carries over from AFs [Dimopoulos and Torres, 1996].*

For the approximate two-valued stable semantics, the fact that model verification can be decided in polynomial time leads to the next result.

Corollary 25. *Let Ξ be an ADF and $s \in S$. $\mathsf{Cred}_{2\mathsf{st}}^{\mathcal{G}_\Xi}(s)$ is* NP*-complete; $\mathsf{Skept}_{2\mathsf{st}}^{\mathcal{G}_\Xi}(s)$ is* coNP*-complete.*

Proof. *The membership parts are clear since* $\mathsf{Ver}^{\mathcal{G}_\Xi}_{2\mathsf{st}}$ *is in* P*. Hardness carries over from AFs [Dimopoulos and Torres, 1996].*

For the ultimate two-valued stable semantics, things are bit more complex. First of all, we recapitulate a result of Brewka et al. [2013a] because we will need the proof later on.

Theorem ([Brewka et al., 2013a, Theorem 9]). *Let* $\Xi = (S, L, C)$ *be an ADF. Deciding whether* Ξ *has an ultimate two-valued stable model is* Σ_2^{P}*-complete.*

Proof. *For membership, we first guess a set* $M \subseteq S$*. We can verify in polynomial time that* M *is a two-valued supported model of* Ξ*, and compute the reduct* Ξ_M*. Using the* NP *oracle, we can compute the grounded semantics* (K', K'') *of the reduct in polynomial time. It then only remains to check* $K' = M$*.*

For hardness, we provide a reduction from the Σ_2^{P}*-complete problem of deciding whether a* $QBF_{2,\exists}$*-formula is valid. Let* $\exists P \forall Q \psi$ *be an instance of* $QBF_{2,\exists}$*-TRUTH where* ψ *is in DNF and* $P, Q \neq \emptyset$*. We have to construct an ADF* D *such that* D *has a stable model iff* $\exists P \forall Q \psi$ *is true.*

First of all, define $\text{-}P = \{\text{-}p \mid p \in P\}$ *for abbreviating the negations of* $p \in P$*. For guessing an interpretation for* P*, define the acceptance formulas* $\varphi_p = \neg\text{-}p$ *and* $\varphi_{\text{-}p} = \neg p$ *for* $p \in P$*. Define* ψ' *as the formula* $\psi[\neg p/\text{-}p]$ *where all occurrences of* $\neg p$ *have been replaced by* -p*. (Note that* ψ *is in DNF and thus* ψ' *is a DNF without negation.) Further add a statement* z *with* $\varphi_z = \neg z \wedge \neg\psi'$*, an integrity constraint that ensures truth of* ψ' *in any model. For* $q \in Q$ *we set* $\varphi_q = \psi'$*. Thus we get the statements* $S = P \cup \text{-}P \cup Q \cup \{z\}$*. We have to show that* D *has a stable model iff* $\exists P \forall Q \psi$ *is true.*

"if": *Let* $M_P \subseteq P$ *be such that the following formula over vocabulary* Q *is a tautology:*

$$\phi = \psi^{(M_P, M_P \cup Q)}$$

We now construct a stable model $M = M_P \cup Q \cup \{\text{-}p \in \text{-}P \mid p \notin M_P\}$*. We first show that* M *is a model of* D*: For each* $p \in M_P$*, we have* -$p \notin M$ *by definition and hence* $M \models \varphi_p = \neg\text{-}p$*. Conversely, if* $p \notin M_P$ *then* -$p \in M$ *and* $M \models \varphi_{\text{-}p} = \neg p$*. For* $q \in Q$*, we have that* $\varphi_q = \psi'$ *and so we have to show* $M \models \psi'$*. This is however immediate since* ϕ *(the partial evaluation of* ψ *with* M *as interpretation for* P*) is a tautology. Finally, by definition* $z \notin M$*, and since* $M \models \psi'$ *we get* $M \not\models \varphi_z = \neg z \wedge \neg\psi'$ *as required.*

To show that M *is a stable model, we have to show that all statements in* M *are true in the ultimate Kripke-Kleene semantics of the reduct* D_M*. The reduct is given by*

- $D_M = (M, L_M, C_M)$ *with*
- $\varphi_p = \neg\bot$ *for* $p \in M$*,*
- $\varphi_{\text{-}p} = \neg\bot$ *for* -$p \in M$*,*
- $\varphi_q = \psi'^{(\emptyset, M)}$*.*

The computation of the Kripke-Kleene semantics starts with $(\emptyset, M)$ *and leads to the first revision* $(K_0', K_0'') = \mathcal{U}_\Xi(\emptyset, M)$*. Since the acceptance condition of any* $p,\text{-}p \in M$ *is tautological, we have* $p, \text{-}p \in K_0'$*, that is, the statements* $p, \text{-}p \in M$ *are considered true. For the next step, the acceptance formula of any* $q \in Q$ *can thus be simplified to*

$$\begin{aligned}\varphi_q^{(M \setminus Q, M)} &= \left(\psi'^{(\emptyset, M)}\right)^{(M \setminus Q, M)} \\ &= \psi'^{(M \setminus Q, M)} \\ &= \psi'[p/\bot : p \notin M, \text{-}p/\bot : \text{-}p \notin M, p/\top : p \in M, \text{-}p/\top : \text{-}p \in M],\end{aligned}$$

a formula over Q that is equivalent to $\phi = \psi^{(M_P, M_P \cup Q)}$. By presumption, ϕ is a tautology. Hence at this point all acceptance formulas partially evaluated by (K'_0, K''_0) are tautologies and thus $\mathcal{U}_{\not\in}(K'_0, K''_0) = (M, M)$, which has already been shown to be a fixpoint of $\mathcal{U}_{\not\in}$.

"only if": *Let $M \subseteq S$ be an ultimate two-valued stable model of D. We have to show that $\exists P \forall Q \psi$ is true. Define $M_P = M \cap P$ and $\phi = \psi^{(M_P, M_P \cup Q)}$. We show that ϕ is a tautology.*

First of all, since M is a model of D_M we have $z \notin M$: assume to the contrary that $z \in M$, then M is a model for $\varphi_z = \neg z \wedge \neg \psi' \equiv \bot \wedge \neg \psi'$, contradiction. Hence $M \not\models \neg z \wedge \neg \psi'$, that is, $M \not\models \neg \psi'$. This shows that $M \models \psi'$, that is, $M \models \varphi_q$ for all $q \in Q$, whence $Q \subseteq M$. Thus the evaluation of $p \in P$ and -$p \in$ -P defined by M shows the truth of the formula

$$\psi'^{(M,M)} = \psi'[p/\top : p \in M, \text{-}p/\top : \text{-}p \in M, p/\bot : p \notin M, \text{-}p/\bot : \text{-}p \notin M][q/\top : q \in Q]$$

Now since M is a stable model of D, the pair (M, M) is the ultimate grounded semantics of the reduct D_M again given by

- $D_M = (M, L_M, C_M)$ *with*
- $\varphi_p = \neg\bot$ *for* $p \in M$,
- $\varphi_{\text{-}p} = \neg\bot$ *for* -$p \in M$,
- $\varphi_q = \psi'^{(\emptyset, M)}$.

To show that ϕ is a tautology, assume to the contrary that ϕ is refutable. As observed in the "if" part, ϕ is equivalent to the formula $\varphi_q^{(M \setminus Q, M)}$. Thus also φ_q is refutable, whence $q \notin \mathcal{U}'_{D_M}(\emptyset, M)$ for all $q \in Q$ and $\mathcal{U}'_{D_M}(\emptyset, M) = M \setminus Q$. Furthermore we know that $\mathcal{U}''_{D_M}(\emptyset, M) = M$. Now $\varphi_q^{(M \setminus Q, M)}$ is refutable and thus $\mathcal{U}_{D_M}(M \setminus Q, M) = (M \setminus Q, M)$. Since $Q \neq \emptyset$, we find that (M, M) is not the least fixpoint of $\mathcal{U}_{D_M}$. Contradiction. □

The hardness reduction in this proof makes use of a statement z that is false in any ultimate two-valued stable model. This can be used to show the same hardness for the credulous reasoning problem for this semantics: we introduce a new statement x that behaves just like $\neg z$, then x is true in some model iff there exists a model.

Proposition 26. *Let Ξ be an ADF and $s \in S$. The problem $\mathsf{Cred}^{\mathcal{U}_{\not\in}}_{\mathsf{2st}}(s)$ is Σ^P_2-complete.*

Proof. in Σ^P_2: *We can guess a set $X \subseteq S$ with $s \in X$ and verify in* coNP *that it is an ultimate two-valued stable model.*

Σ^P_2-hard: *Let $\exists P \forall Q \psi$ be a QBF. We use the same ADF construction as in the hardness proof of $\mathsf{Exists}^{\mathcal{U}_{\not\in}}_{\mathsf{2st}}$ and augment D by an additional statement x with $\varphi_x = \neg z$. It is clear that in any model of D, z must be false and so x must be true. So x is true in some two-valued stable model of D iff D has a two-valued stable model iff $\exists P \forall Q \psi$ is true.* □

A similar argument works for the skeptical reasoning problem: Given a QBF $\forall P \exists Q \psi$, we construct its negation $\exists P \forall Q \neg\psi$, whose associated ADF D has an ultimate two-valued stable model (where z is false) iff $\exists P \forall Q \neg\psi$ is true iff the original QBF $\forall P \exists Q \psi$ is false. Hence $\forall P \exists Q \psi$ is true iff z is true in all ultimate two-valued stable models of D.

Proposition 27. *Let Ξ be an ADF and $s \in S$. The problem $\mathsf{Skept}^{\mathcal{U}_{\not\in}}_{\mathsf{2st}}(s)$ is Π^P_2-complete.*

Proof. in Π^P_2: *To decide the co-problem, we guess a set $X \subseteq S$ with $s \notin X$ and verify in* coNP *that it is an ultimate two-valued stable model.*

approximate ($\mathcal{G}_\sqsubseteq$), σ	admissible	complete	preferred	grounded	model	stable model
$\mathsf{Ver}_\sigma^{\mathcal{G}_\sqsubseteq}$	coNP-c (Proposition 10)	D_2^P-c (Proposition 11)	Π_2^P-c (Proposition 16)	D_2^P-c (Theorem 12)	in P ([Brewka et al., 2013a, Prop. 5])	in P (Theorem 22)
$\mathsf{Exists}_\sigma^{\mathcal{G}_\sqsubseteq}$	Σ_2^P-c (Theorem 13)	Σ_2^P-c (Corollary 14)	Σ_2^P-c (Corollary 14)	coNP-c (Proposition 15)	NP-c ([Brewka et al., 2013a, Prop. 5])	NP-c (Theorem 22)
$\mathsf{Cred}_\sigma^{\mathcal{G}_\sqsubseteq}$	Σ_2^P-c (Proposition 17)	Σ_2^P-c (Proposition 17, Lemma 7)	Σ_2^P-c (Proposition 17, Lemma 7)	coNP-c (Proposition 18)	NP-c (Corollary 24)	NP-c (Corollary 25)
$\mathsf{Skept}_\sigma^{\mathcal{G}_\sqsubseteq}$	trivial	coNP-c (Corollary 19)	Π_3^P-c (Theorem 20)	coNP-c (Proposition 18)	coNP-c (Corollary 24)	coNP-c (Corollary 25)

ultimate ($\mathcal{U}_\sqsubseteq$), σ	admissible	complete	preferred	grounded	model	stable model
$\mathsf{Ver}_\sigma^{\mathcal{U}_\sqsubseteq}$	coNP-c ([Brewka et al., 2013a, Prop. 10])	D_2^P-c ([Brewka et al., 2013a, Cor. 7])	Π_2^P-c (Proposition 16)	D_2^P-c ([Brewka et al., 2013a, Thm. 6])	in P ([Brewka et al., 2013a, Prop. 5])	coNP-c (Proposition 23)
$\mathsf{Exists}_\sigma^{\mathcal{U}_\sqsubseteq}$	Σ_2^P-c (Theorem 13)	Σ_2^P-c (Corollary 14)	Σ_2^P-c (Corollary 14)	coNP-c (Proposition 15)	NP-c ([Brewka et al., 2013a, Prop. 5])	Σ_2^P-c ([Brewka et al., 2013a, Thm. 9])
$\mathsf{Cred}_\sigma^{\mathcal{U}_\sqsubseteq}$	Σ_2^P-c (Proposition 17)	Σ_2^P-c (Proposition 17, Lemma 7)	Σ_2^P-c (Proposition 17, Lemma 7)	coNP-c (Proposition 18)	NP-c (Corollary 24)	Σ_2^P-c (Proposition 26)
$\mathsf{Skept}_\sigma^{\mathcal{U}_\sqsubseteq}$	trivial	coNP-c (Corollary 19)	Π_3^P-c (Theorem 20)	coNP-c (Proposition 18)	coNP-c (Corollary 24)	Π_2^P-c (Proposition 27)

Table 2: Complexity results for semantics of Abstract Dialectical Frameworks.

Π_2^P-hard: *Let $\forall P \exists Q \psi$ be a QBF with ψ in CNF. Define the QBF $\exists P \forall Q \neg\psi$ and observe that $\neg\psi$ can be transformed into DNF in linear time. We use this new QBF to construct an ADF D as we did in the hardness proof of $\mathsf{Exists}_{2\mathsf{st}}^{\mathcal{U}_\sqsubseteq}$. As observed in the proof of Proposition 26, the special statement z is false in all ultimate two-valued stable models of D. To show that $\forall P \exists Q \psi$ is true iff z is true in all ultimate two-valued stable models of D, we show that $\forall P \exists Q \psi$ is false iff D has an ultimate two-valued stable model where z is false: $\forall P \exists Q \psi$ is false iff $\neg\forall P \exists Q \psi$ is true iff $\exists P \forall Q \neg\psi$ is true iff D has an ultimate two-valued stable model iff D has an ultimate two-valued stable model where z is false.* □

5 Complexity of Bipolar ADFs

We first note that since BADFs are a subclass of ADFs, all membership results from the previous section immediately carry over. However, we can show that many problems will in fact become easier. Intuitively, computing the revision operators is now P-easy because the associated satisfiability/tautology problems only have to treat restricted acceptance formulas. In bipolar ADFs, by definition, if in some three-valued pair (X, Y) a statement s is accepted by a revision operator ($s \in \mathcal{O}'(X, Y)$), it will stay so if we set its undecided supporters to true and its undecided attackers to false. Symmetrically, if a statement is rejected by an operator ($s \notin \mathcal{O}''(X, Y)$), it will stay so if we set its undecided supporters to false and its undecided attackers to true. This is the key idea underlying the next result.

Proposition 28. *Let Ξ be a BADF with $L = L^+ \cup L^-$, $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$, $s \in S$ and $X \subseteq Y \subseteq S$.*

1. *Deciding $s \in \mathcal{O}'(X, Y)$ is in* P.
2. *Deciding $s \in \mathcal{O}''(X, Y)$ is in* P.

Proof. It suffices to show the claims for $\mathcal{O} = \mathcal{U}_\Xi$, since the result that $s \in \mathcal{U}''_\Xi(X, Y)$ is computable in P implies that deciding $s \in \mathcal{G}''_\Xi(X, Y)$ is in P, due to coincidence of the two operators. Further due to Proposition 4 we know that deciding $s \in \mathcal{G}'_\Xi(X, Y)$ is a problem in P.

Recall that for $M \subseteq S$, if a link (z, s) is attacking, then it cannot be the case that $M \not\models \varphi_s$ and $M \cup \{z\} \models \varphi_s$. Similarly if (z, s) is supporting, then it cannot be the case that $M \models \varphi_s$ and $M \cup \{z\} \not\models \varphi_s$. If (x, s) is attacking and supporting then for any $M \subseteq S$ we have $M \models \varphi_s$ iff $M \cup \{z\} \models \varphi_s$, i.e. a change of the truth value of z does not change the evaluation of φ_s.

Given a consistent pair (X, Y) and $s \in S$ we can use a "canonical" interpretation representing all $X \subseteq Z \subseteq Y$ as follows. Note that the aforementioned "redundant" links, i.e. links in the intersection $L^+ \cap L^-$ can be disregarded completely and for ease of notation we will assume in the proof that no such link is present (formally if (x, s) is a redundant link, then we can replace each x in φ_s uniformly with $\top$ or $\bot$). Let $Z \subseteq S$, $Z' \subseteq att_\Xi(s)$ and $Z'' \subseteq supp_\Xi(s)$. Then

$$\begin{aligned} & s \in \mathcal{U}'_\Xi(Z, Z) \\ \text{iff } & s \in \mathcal{U}'_\Xi(Z \setminus Z', Z) \\ \text{iff } & s \in \mathcal{U}'_\Xi(Z \setminus Z', Z \cup Z''). \end{aligned}$$

The "if" direction is both times trivially satisfied. This can be seen by the easy fact that if $\varphi_s^{L,U}$ is tautological, then so is $\varphi_s^{L',U'}$ with $(L, U) \leq_i (L', U')$. Suppose the first "only if" does not hold, i.e. the first line holds, but the second is not true. Then there exists a set H with $(Z \setminus Z') \subseteq H \subseteq Z$ such that $H \not\models \varphi_s$. By assumption $Z \models \varphi_s$ and since $H \cup (Z' \cap Z) = Z$ also $H \cup (Z' \cap Z) \models \varphi_s$, which is a contradiction, since $Z' \subseteq att_\Xi(s)$ and thus $(Z' \cap Z) \subseteq att_\Xi(s)$, which implies that there exists a statement in $att_\Xi(s)$ which is not attacking.

Suppose the second only if does not hold, then there exists a H with $(Z \setminus Z') \subseteq H \subseteq (Z \cup Z'')$ such that $H \not\models \varphi_s$. Since we have that $(Z \setminus Z') \subseteq (H \setminus (Z'' \setminus Z)) \subseteq Z$ it follows that $H \setminus (Z'' \setminus Z) \models \varphi_s$, which is a contradiction since Z'' consists only of supporters of s.

Now we set the canonical interpretation as $Z = X \cup (Y \setminus supp_\Xi(s))$. Observe that there exists $Z' \subseteq att_\Xi(s)$ and $Z'' \subseteq supp_\Xi(s)$ s.t. $X = Z \setminus Z'$ and $Y = Z \cup Z''$, thus $s \in \mathcal{U}'_\Xi(Z, Z)$ iff $s \in \mathcal{U}'_\Xi(X, Y)$. Since we can construct Z in polynomial time if L^+ and L^- are known and deciding $s \in \mathcal{U}'_\Xi(Z, Z)$ simply amounts to evaluating a formula under a valuation, the first claim follows.

Showing the second claim is similar. Let $Z \subseteq S$, $Z' \subseteq supp_\Xi(s)$ and $Z'' \subseteq att_\Xi(s)$. Then

$$\begin{aligned} & s \in \mathcal{U}''_\Xi(Z, Z) \\ \text{iff } & s \in \mathcal{U}''_\Xi(Z \setminus Z', Z) \\ \text{iff } & s \in \mathcal{U}''_\Xi(Z \setminus Z', Z \cup Z''). \end{aligned}$$

□

Using the generic upper bounds of Theorem 8, it is now straightforward to show membership results for BADFs with known link types.

Corollary 29. *Let Ξ be a BADF with $L = L^+ \cup L^-$, consider any operator $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$ and semantics $\sigma \in \{adm, com\}$. For $X \subseteq Y \subseteq S$ and $s \in S$, we find that*

- $\mathsf{Ver}_\sigma^{\mathcal{O}}(X, Y)$ *and* $\mathsf{Ver}_{\mathsf{grd}}^{\mathcal{O}}(X, Y)$ *are in* P*;*
- $\mathsf{Ver}_{\mathsf{pre}}^{\mathcal{O}}(X, Y)$ *is in* coNP*;*
- $\mathsf{Exists}_\sigma^{\mathcal{O}}$, $\mathsf{Exists}_{\mathsf{pre}}^{\mathcal{O}}$, $\mathsf{Cred}_\sigma^{\mathcal{O}}(s)$ *and* $\mathsf{Cred}_{\mathsf{pre}}^{\mathcal{O}}(s)$ *are in* NP*;*
- $\mathsf{Exists}_{\mathsf{grd}}^{\mathcal{O}}$, $\mathsf{Cred}_{\mathsf{grd}}^{\mathcal{O}}(s)$, $\mathsf{Skept}_{\mathsf{grd}}^{\mathcal{O}}(s)$, $\mathsf{Skept}_{\mathsf{com}}^{\mathcal{O}}(s)$ *are in* P*;*
- $\mathsf{Skept}_{\mathsf{pre}}^{\mathcal{O}}(s)$ *is in* Π_2^P*.*

Proof. *Membership is due to Theorem 8 and the complexity bounds of the operators in BADFs in Proposition 28, just note that* $\Sigma_0^P = \Pi_0^P = \mathsf{P}$*. Further, due to Corollary 9, we can compute the grounded pair in* $\mathsf{P}^{\mathsf{P}} = \mathsf{P}$*. For the existence of non-trivial pairs we can simply guess and check in polynomial time for admissible pairs and thus also for complete and preferred semantics.* □

Hardness results straightforwardly carry over from AFs.

Proposition 30. *Let Ξ be a BADF with $L = L^+ \cup L^-$, consider any operator $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$ and semantics $\sigma \in \{adm, com, pre\}$. For $X \subseteq Y \subseteq S$ and $s \in S$:*

- $\mathsf{Ver}_{\mathsf{pre}}^{\mathcal{O}}(X, Y)$ *is* coNP *hard;*
- $\mathsf{Exists}_\sigma^{\mathcal{O}}$ *and* $\mathsf{Cred}_\sigma^{\mathcal{O}}(s)$ *are* NP *hard;*
- $\mathsf{Skept}_{\mathsf{pre}}^{\mathcal{O}}(s)$ *is* Π_2^P *hard.*

Proof. *Hardness results from AFs for these problems carry over to BADFs as for all semantics AFs are a special case of BADFs [Brewka et al., 2013a; Strass, 2013a]. The complexities of the problems on AFs for admissible and preferred semantics are shown by Dimopoulos and Torres [1996], except for the* Π_2^P*-completeness result of skeptical preferred semantics, which is shown by Dunne and Bench-Capon [2002]. The complete semantics is studied by Coste-Marquis et al. [2005].* □

We next show that there is no hope that the existence problems for approximate and ultimate two-valued stable models coincide as there are cases when the semantics differ.

Example 3. Consider the BADF $F = (S, L, C)$ with statements $S = \{a, b, c\}$ and acceptance formulas $\varphi_a = \mathbf{t}$, $\varphi_b = a \vee c$ and $\varphi_c = a \vee b$. The only two-valued supported model is (S, S) where all statements are true. This pair is also an ultimate two-valued stable model, since $\mathcal{U}_F'(\emptyset, S) = \{a\}$, and both $\varphi_b^{(\{a\},S)} = \mathbf{t} \vee c$ and $\varphi_c^{(\{a\},S)} = \mathbf{t} \vee b$ are tautologies, whence we have $\mathcal{U}_F'(\{a\}, S) = S$. However, (S, S) is not an approximate two-valued stable model: although $\mathcal{G}_F'(\emptyset, S) = \{a\}$, then $\mathcal{G}_F'(\{a\}, S) = \{a\}$ and we thus cannot reconstruct the upper bound S. Thus F has no approximate two-valued stable models.

So approximate and ultimate two-valued (stable) model semantics are indeed different. However, we can show that the respective existence problems have the same complexity.

Proposition 31. *Let Ξ be a BADF with $L = L^+ \cup L^-$, $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$ an operator and semantics $\sigma \in \{2su, 2st\}$. For $X \subseteq S$,* $\mathsf{Ver}_\sigma^{\mathcal{O}}(X, X)$ *is in* P*;* $\mathsf{Exists}_\sigma^{\mathcal{O}}$ *is* NP*-complete.*

$\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$, σ	admissible	complete	preferred	grounded	model	stable model
$\mathsf{Ver}_\sigma^\mathcal{O}$	in P (Corollary 29, Proposition 30)	in P (Corollary 29, Proposition 30)	coNP-c (Corollary 29, Proposition 30)	in P (Corollary 29)	in P (Proposition 31)	in P (Proposition 31)
$\mathsf{Exists}_\sigma^\mathcal{O}$	NP-c (Corollary 29, Proposition 30)	NP-c (Corollary 29, Proposition 30)	NP-c (Corollary 29, Proposition 30)	in P (Corollary 29)	NP-c (Proposition 31)	NP-c (Proposition 31)
$\mathsf{Cred}_\sigma^\mathcal{O}$	NP-c (Corollary 29, Proposition 30)	NP-c (Corollary 29, Proposition 30)	NP-c (Corollary 29, Proposition 30)	in P (Corollary 29)	NP-c (Corollary 32)	NP-c (Corollary 32)
$\mathsf{Skept}_\sigma^\mathcal{O}$	trivial	in P (Corollary 29)	Π_2^P-c (Corollary 29, Proposition 30)	in P (Corollary 29)	coNP-c (Corollary 32)	coNP-c (Corollary 32)

Table 3: Complexity results for semantics of bipolar Abstract Dialectical Frameworks.

Proof. *Membership carries over – for supported models from [Brewka et al., 2013a, Proposition 5], for approximate stable models from Theorem 22. For membership for ultimate stable models, we can use Proposition 28 to adapt the decision procedure of Proposition 21. In any case, hardness carries over from AFs [Dimopoulos and Torres, 1996].* □

For credulous and skeptical reasoning over the two-valued semantics, membership is straightforward and hardness again carries over from argumentation frameworks.

Corollary 32. *Let Ξ be a BADF with $L = L^+ \cup L^-$; consider any operator $\mathcal{O} \in \{\mathcal{G}_\Xi, \mathcal{U}_\Xi\}$ and semantics $\sigma \in \{2su, 2st\}$. For $s \in S$, $\mathsf{Cred}_\sigma^\mathcal{O}(s)$ is* NP*-complete;* $\mathsf{Skept}_\sigma^\mathcal{O}(s)$ *is* coNP*-complete.*

6 Discussion

In this paper we studied the computational complexity of abstract dialectical frameworks using approximation fixpoint theory. We showed several novel results for two families of ADF semantics, the approximate and ultimate semantics, which are themselves inspired by argumentation and AFT. We showed that in most cases the complexity increases by one level of the polynomial hierarchy compared to the corresponding reasoning tasks on AFs. A notable difference between the two families of semantics lies in the stable semantics, where the approximate version is easier than its ultimate counterpart. For the restricted, yet powerful class of bipolar ADFs we proved that for the corresponding reasoning tasks AFs and BADFs have the same complexity, which suggests that many types of relations between arguments can be introduced without increasing the worst-time complexity. On the other hand, our results again emphasize that arbitrary (non-bipolar) ADFs cannot be compiled into equivalent Dung AFs in deterministic polynomial time unless the polynomial hierarchy collapses to the first level.

Although AFT may not have been developed with the intention of studying newly conceived formalisms and defining semantics for them, we show that indeed AFT is a fruitful basis for investigating such new formal models, in particular it is well-suited for analyzing their complexity.

Our results lay the foundation for future algorithms and their implementation, for example augmenting the ADF system DIAMOND [Ellmauthaler and Strass, 2013] to support also the approximate semantics family, as well as devising efficient methods for the interesting class of BADFs.

For further future work several promising directions are possible. Studying easier fragments of ADFs as well as parameterized complexity analysis can lead to efficient decision procedures, as is witnessed for AFs [Dvořák et al., 2014; Dvořák et al., 2012]. We also deem it auspicious to use alternative representations of acceptance conditions, for instance by employing techniques from knowledge compilation [Darwiche and Marquis, 2002].

AFs also feature several other useful semantics for which a detailed analysis would reveal further insights, for example semi-stable semantics [Caminada et al., 2012] and naive-based semantics, such as cf2 [Baroni et al., 2005]. Furthermore in [Polberg et al., 2013] an extension-based semantics for ADFs is proposed and a complexity analysis would be interesting.

Regarding semantical analysis, it would be useful to consider principle-based evaluations of ADF semantics as was done for AFs [Baroni and Giacomin, 2007]. Furthermore it appears natural to compare (ultimate) ADF semantics and ultimate logic programming semantics [Denecker et al., 2004] using approximation fixpoint theory, in particular with respect to computational complexity.

Acknowledgements

This research has been supported by DFG (project BR 1817/7-1) and FWF (project I1102).

References

Pietro Baroni and Massimiliano Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence*, 171(10-15):675–700, 2007.

Pietro Baroni, Massimiliano Giacomin, and Giovanni Guida. SCC-recursiveness: a general schema for argumentation semantics. *Artificial Intelligence*, 168(1-2):162–210, 2005.

Trevor J. M. Bench-Capon and Paul E. Dunne. Argumentation in Artificial Intelligence. *Artificial Intelligence*, 171(10-15):619–641, July 2007. ISSN 0004-3702. doi: 10.1016/j.artint.2007.05.001. URL `http://dx.doi.org/10.1016/j.artint.2007.05.001`.

Gerhard Brewka and Thomas F. Gordon. Carneades and Abstract Dialectical Frameworks: A Reconstruction. In *Computational Models of Argument: Proceedings of COMMA 2010*, volume 216 of *FAIA*, pages 3–12. IOS Press, September 2010.

Gerhard Brewka and Stefan Woltran. Abstract Dialectical Frameworks. In *Proceedings of the Twelfth International Conference on the Principles of Knowledge Representation and Reasoning (KR)*, pages 102–111, 2010.

Gerhard Brewka, Paul E. Dunne, and Stefan Woltran. Relating the Semantics of Abstract Dialectical Frameworks and Standard AFs. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence, IJCAI 2011*, pages 780–785. IJCAI/AAAI, 2011. ISBN 978-1-57735-516-8. URL `http://ijcai.org/papers11/Papers/IJCAI11-137.pdf`.

Gerhard Brewka, Stefan Ellmauthaler, Hannes Strass, Johannes Peter Wallner, and Stefan Woltran. Abstract Dialectical Frameworks Revisited. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, pages 803–809. AAAI Press, August 2013a.

Gerhard Brewka, Sylwia Polberg, and Stefan Woltran. Generalizations of Dung frameworks and their role in formal argumentation. *Intelligent Systems, IEEE*, PP(99), 2013b. ISSN 1541-1672. doi: 10.1109/MIS.2013.122. In press.

Martin W.A. Caminada and Leila Amgoud. On the evaluation of argumentation formalisms. *Artificial Intelligence*, 171(5–6):286–310, 2007.

Martin W.A. Caminada, Walter A. Carnielli, and Paul E. Dunne. Semi-stable Semantics. *Journal of Logic and Computation*, 22(5):1207–1254, 2012.

Sylvie Coste-Marquis, Caroline Devred, and Pierre Marquis. Symmetric Argumentation Frameworks. In *Proceedings of the 8th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, ECSQARU 2005*, volume 3571 of *LNCS*, pages 317–328. Springer, 2005. ISBN 3-540-27326-3.

Adnan Darwiche and Pierre Marquis. A Knowledge Compilation Map. *Journal of Artificial Intelligence Research (JAIR)*, 17:229–264, 2002.

B.A. Davey and H.A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, second edition, 2002.

Marc Denecker, Victor W. Marek, and Miroslaw Truszczyński. Approximations, Stable Operators, Well-Founded Fixpoints and Applications in Nonmonotonic Reasoning. In *Logic-Based Artificial Intelligence*, pages 127–144. Kluwer Academic Publishers, 2000.

Marc Denecker, Victor W. Marek, and Miroslaw Truszczyński. Uniform Semantic Treatment of Default and Autoepistemic Logics. *Artificial Intelligence*, 143(1):79–122, 2003.

Marc Denecker, Victor W. Marek, and Miroslaw Truszczyński. Ultimate approximation and its application in nonmonotonic knowledge representation systems. *Information and Computation*, 192(1):84–121, 2004.

Yannis Dimopoulos and Alberto Torres. Graph Theoretical Structures in Logic Programs and Default Theories. *Theoretical Computer Science*, 170(1-2):209–244, 1996.

Yannis Dimopoulos, Bernhard Nebel, and Francesca Toni. On the computational complexity of assumption-based argumentation for default reasoning. *Artificial Intelligence*, 141(1/2):57–78, 2002.

Phan Minh Dung. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. *Artificial Intelligence*, 77(2):321–358, 1995.

Paul E. Dunne and Trevor J. M. Bench-Capon. Coherence in Finite Argument Systems. *Artificial Intelligence*, 141(1/2):187–203, 2002.

Wolfgang Dvořák, Sebastian Ordyniak, and Stefan Szeider. Augmenting tractable fragments of abstract argumentation. *Artificial Intelligence*, 186:157–173, 2012. ISSN 0004-3702. doi: 10.1016/j.artint.2012.03.002.

Wolfgang Dvořák, Matti Järvisalo, Johannes P. Wallner, and Stefan Woltran. Complexity-Sensitive Decision Procedures for Abstract Argumentation. *Artificial Intelligence*, 206:53–78, 2014. doi: 10.1016/j.artint.2013.10.001.

Stefan Ellmauthaler and Hannes Strass. The DIAMOND System for Argumentation: Preliminary Report. In *Proceedings of ASPOCP*, pages 97–107, 2013.

Thomas F. Gordon, Henry Prakken, and Douglas Walton. The Carneades model of argument and burden of proof. *Artificial Intelligence*, 171(10–15):875–896, 2007.

Sylwia Polberg, Johannes P. Wallner, and Stefan Woltran. Admissibility in the Abstract Dialectical Framework. In *Proceedings of CLIMA XIV*, volume 8143 of *LNAI*, pages 102–118. Springer, 2013. doi: 10.1007/978-3-642-40624-9_7.

Hannes Strass. Approximating operators and semantics for abstract dialectical frameworks. *Artificial Intelligence*, 205:39–70, December 2013a.

Hannes Strass. Instantiating knowledge bases in abstract dialectical frameworks. In *Proceedings of CLIMA XIV*, volume 8143 of *LNAI*, pages 86–101. Springer, September 2013b.

Bas Van Gijzel and Henry Prakken. Relating Carneades with abstract argumentation. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence – Volume Two*, pages 1113–1119. AAAI Press, 2011.

Adam Wyner, Trevor J. M. Bench-Capon, and Paul E. Dunne. On the instantiation of knowledge bases in abstract argumentation frameworks. In *Proceedings of CLIMA XIV*, volume 8143 of *LNAI*, pages 34–50. Springer, September 2013.