Recursion in Abstract Argumentation is Hard — On the Complexity of Semantics Based on Weak Admissibility

Wolfgang Dvořák,[1] Markus Ulbricht,[2] Stefan Woltran[1]
[1] TU Wien, Institute of Logic and Computation, Austria
[2] Leipzig University, Department of Computer Science, Germany
{dvorak,woltran}@dbai.tuwien.ac.at, mulbricht@informatik.uni-leipzig.de

Abstract

We study the computational complexity of abstract argumentation semantics based on weak admissibility, a recently introduced concept to deal with arguments of self-defeating nature. Our results reveal that semantics based on weak admissibility are of much higher complexity (under typical assumptions) compared to all argumentation semantics which have been analysed in terms of complexity so far. In fact, we show PSPACE-completeness of all non-trivial standard decision problems for weak-admissible based semantics. We then investigate potential tractable fragments and show that restricting the frameworks under consideration to certain graph-classes significantly reduces the complexity. As a strategy for implementation we also provide a polynomial-time reduction to DATALOG with stratified negation.

Introduction

Abstract argumentation frameworks (AFs) as introduced by Dung (1995) are nowadays identified as key concept to understand the fundamental mechanisms behind formal argumentation and non-monotonic reasoning. In these frameworks, it is solely the attack relation between (abstract) arguments that is used to determine the semantics of a given AF, i.e. jointly acceptable sets of arguments called extensions.

Most of the existing argumentation semantics were either based on the concept of naivety or admissibility (van der Torre and Vesic 2017). The former is satisfied if the selected sets are maximal conflict-free. For the latter, it is required that the sets defend themselves (each attacker of an argument in the set is counter-attacked by the set).

There is a wide consensus that the absence of defense in naive extensions potentially leads to undesired results. However, already Dung noticed that also the concept of defense can be seen problematic; in particular, when self-defeating arguments are involved, that is, arguments which attack themselves directly or indirectly through an odd loop of arguments. Such "dummy" arguments may block the acceptance state of other reasonable ones, while never standing a chance of being accepted themselves. This issue has been known for a long time, and inspired several approaches to mitigate the effect of self-defeating arguments, see e.g. (Bodanza and Tohmé 2009; Amendola and Ricca 2019; Fazzinga, Flesca, and Furfaro 2020). However, no semantics for abstract argumentation among the numerous invented so far (see e.g. (Baroni, Caminada, and Giacomin 2011)) has addressed this problem in a commonly agreed way.

In a recent paper, Baumann, Brewka, and Ulbricht (2020b) propose a mediating position between naivity and admissibility and introduced the concept of weak admissibility. This new concept aims at limiting the effect of self-defeating arguments by verifying the credibility of arguments in a recursive fashion: any conflict-free set of arguments is considered acceptable unless attacked by some serious rival. On top of handling self-defeating arguments in a more reasonable way, the introduced semantics possess several promising theoretical properties which were already pointed out in (Baumann, Brewka, and Ulbricht 2020b) by showing that weak admissibility inherits many of the desirable properties of its classical Dung-style counterpart. These observations triggered further investigations of these semantics w.r.t. well-known postulates discussed in the literature (see (Baroni, Caminada, and Giacomin 2018; van der Torre and Vesic 2017)): in particular, Dauphin, Rienstra, and van der Torre (2020) have studied the aforementioned postulates in a comprehensive fashion while Baumann, Brewka, and Ulbricht (2020a) address concepts like strong equivalence for semantics based on weak admissibility.

In light of these solid theoretical results an investigation from a computational point of view stands to reason as well. In this paper, we take several steps towards this direction by thoroughly analyzing the computational complexity of weak admissibility as well as providing a DATALOG encoding.

More specifically, our main contributions are as follows:

- We show that all standard decision problems for weak-admissible based semantics (with the exception of skeptical w-admissible acceptance) are PSPACE-complete.
- We analyze the effect of restricting the AFs under consideration to certain graph-classes, which in most cases renders the "weak" semantics computationally comparable to their Dung-style counterparts, which is a significant drop in their complexity.
- Towards implementation we provide a polynomial-time reduction to non-recursive DATALOG with stratified negation which is known to be PSPACE-complete in terms of program-complexity (cf. (Dantsin et al. 2001)).

The complexity analysis we provide is of particular interest, since all known complexity results of standard tasks for argumentation semantics are located within the first two layers of the polynomial hierarchy (see, e.g. (Dvořák and Dunne 2018)). This holds even for semantics which have a certain recursive nature like cf2- or stage2-semantics; see (Gaggl and Woltran 2013; Dvořák and Gaggl 2016) for the respective complexity analyses. We recall that under the assumption that the polynomial hierarchy does not collapse, problems complete for PSPACE are rated as significantly harder than problems located at lower levels of the polynomial hierarchy. Our results are mirrored in the complexity landscape of nonmonotonic reasoning in the broad sense, where decision problems for many prominent formalisms (like default logic or circumscription) are located on the second level of the polynomial hierarchy (see, e.g. (Cadoli and Schaerf 1993; Thomas and Vollmer 2010) for survey articles), and only a few formalisms reach PSPACE-hardness. Examples for the latter are nested circumscription (Cadoli, Eiter, and Gottlob 2005), nested counterfactuals (Eiter and Gottlob 1996), model-preference default logic (Papadimitriou 1991), and theory curbing (Eiter and Gottlob 2006).

Background

Let us start by giving the necessary preliminaries.

Standard Concepts and Classical Semantics

We fix a non-finite background set $\mathcal{U}$. An argumentation framework (AF) (Dung 1995) is a directed graph $F = (A, R)$ where $A \subseteq \mathcal{U}$ represents a set of arguments and $R \subseteq A \times A$ models *attacks* between them. $\mathcal{F}$ denotes the set of all finite AFs over $\mathcal{U}$; we shall consider finite AFs only.

Now assume $F = (A, R)$. For $S \subseteq A$ we let $F{\downarrow_S} = (A \cap S, R \cap (S \times S))$. For $a, b \in A$, if $(a, b) \in R$ we say that a *attacks* b as well as a *attacks* (the set) E given that $b \in E \subseteq A$. Moreover, E is conflict-free in F (for short, $E \in cf(F)$) iff for no $a, b \in E$, $(a, b) \in R$. We say a set E *classically defends* (c-defends) an argument a (in F) if any attacker of a is attacked by some argument of E, i.e. for each $(b, a) \in R$, there is $c \in E$ such that $(c, b) \in R$.

A *semantics* σ is a mapping $\sigma : \mathcal{F} \to 2^{2^{\mathcal{U}}}$ where $F \mapsto \sigma(F) \subseteq 2^A$, i.e. given an AF $F = (A, R)$ a semantics returns a subset of 2^A. We consider here *admissible complete*, *grounded*, and *preferred* semantics (abbr. ad, co, gr, pr).

Definition 0.1. Let $F = (A, R)$ be an AF and $E \in cf(A)$.

1. $E \in ad(F)$ iff E c-defends all its elements,
2. $E \in co(F)$ iff $E \in ad(F)$ and, for any x c-defended by E, we have $x \in E$,
3. $E \in gr(F)$ iff E is $\subseteq$-minimal in $co(F)$,
4. $E \in pr(F)$ iff E is $\subseteq$-maximal in $ad(F)$.

Weak Admissible-Based Semantics

The reduct is the central notion in the definition of weak admissible semantics (Baumann, Brewka, and Ulbricht 2020b). For a compact definition, we use, given an AF $F = (A, R)$, $E_F^+ = \{a \in A \mid E \text{ attacks } a \text{ in } F\}$ as well as $E_F^\oplus = E \cup E_F^+$. The latter set is known as the *range* of E in F. When clear from the context, we omit the subscript F.

Definition 0.2. Let $F = (A, R)$ be an AF and let $E \subseteq A$. The E-reduct of F is the AF $F^E = (E^*, R \cap (E^* \times E^*))$ where $E^* = A \setminus E_F^\oplus$.

By definition, F^E is the subframework of F obtained by removing the range of E as well as corresponding attacks, i.e. $F^E = F{\downarrow_{A \setminus E^\oplus}}$. Intuitively, the E-reduct contains those arguments whose status still needs to be decided, assuming the arguments in E are accepted. This intuition is captured in the forthcoming central definition.

Definition 0.3. For an AF $F = (A, R)$, $E \subseteq A$ is called *weakly admissible* (or *w-admissible*) in F ($E \in ad^w(F)$) iff

1. $E \in cf(F)$ and
2. for any attacker y of E we have $y \notin \bigcup ad^w\left(F^E\right)$.

The major difference between the standard definition of admissibility and the "weak" one is that extensions do not have to defend themselves against *all* attackers: attackers which do not appear in any w-admissible set of the reduct can be neglected.

Example 0.4. Consider the following simple example:

F: $\qquad\qquad$ $F^{\{a\}} = F^{\{b\}}$:

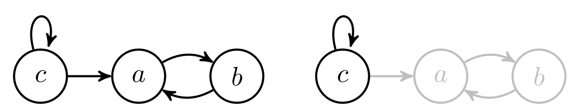

While we observe $\{a\} \notin ad(F)$, we can verify weak admissibility of $\{a\}$ in F. Obviously, $\{a\}$ is conflict-free in F (condition 1). Since c is the only attacker of $\{a\}$ in $F^{\{a\}}$ we have to check $c \notin \bigcup ad^w\left(F^{\{a\}}\right)$ (condition 2). Since $\{c\}$ is not conflict-free in the reduct $F^{\{a\}} = (\{c\}, \{(c, c)\})$ we find $\{c\} \notin ad^w\left(F^{\{a\}}\right)$ yielding $\bigcup ad^w\left(F^{\{a\}}\right) = \emptyset$. Hence, $c \notin \bigcup ad^w\left(F^{\{a\}}\right)$, and thus $\{a\} \in ad^w(F)$. $\diamond$

Following the classical Dung-style semantics, *weakly preferred* extensions are defined as $\subseteq$-maximal w-admissible extensions.

Definition 0.5. For an AF $F = (A, R)$, $E \subseteq A$ is called *weakly preferred* (or *w-preferred*) in F ($E \in pr^w(F)$) iff E is $\subseteq$-maximal in $ad^w(F)$.

In order to define the "weak" counterparts to Dung's grounded and complete extensions, the following notion of "weak defense" has been proposed in (Baumann, Brewka, and Ulbricht 2020b):

Definition 0.6. Let $F = (A, R)$ be an AF. Given two sets $E, X \subseteq A$. We say E *weakly defends* (or *w-defends*) X iff for any attacker y of X we have,

1. E attacks y, or (c-defense)
2. $y \notin \bigcup ad^w\left(F^E\right)$, $y \notin E$ and $X \subseteq X' \in ad^w(F)$.

Now weakly complete and weakly grounded extensions can be defined analogously to complete and grounded ones:

Definition 0.7. For an AF $F = (A, R)$, $E \subseteq A$ is called *weakly complete* (or just *w-complete*) in F ($E \in co^w(F)$) iff

$E \in ad^w(F)$ and for any X, s.t. $E \subseteq X$ and X w-defended by E, we have $X \subseteq E$.

A set $E \subseteq A$ is called *weakly grounded* (or *w-grounded*) in F ($E \in gr^w(F)$) iff E is $\subseteq$-minimal in $co^w(F)$.

The following relations between weak admissibility semantics and their Dung-style counterparts will be useful throughout the present paper: We have $ad(F) \subseteq ad^w(F) \subseteq cf(F)$ as well as $pr^w(F) \subseteq co^w(F) \subseteq ad^w(F)$, and $gr^w(F) \subseteq co^w(F)$ (see (Baumann, Brewka, and Ulbricht 2020b, Proposition 5.6)). Moreover, as it is the case for the classical semantics, a set $E \subseteq A$ is w-preferred in F iff it is $\subseteq$-maximal in $co^w(F)$ (Baumann, Brewka, and Ulbricht 2020b, Theorem 5.3).

Towards a more convenient notion of weak defense, the following characterization has been developed in (Baumann, Brewka, and Ulbricht 2020a); it is suitable in all cases that "matter", i.e. cases where w-completeness of a given set is to be verified:

Proposition 0.8. *Let F be an AF and let $E \in ad^w(F)$. Then, for any $X = E \dot\cup D$ we have that E w-defends X iff*

1. *for any attacker y of D, $y \notin \bigcup ad^w(F^E)$, and*
2. *there is a set $D \subseteq D'$ with $D' \in ad^w\left(F^E\right)$.*

Example 0.9. Consider the AF F:

F:
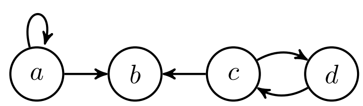

Let us verify that $E = \{d\}$ w-defends $X = \{b, d\}$. Since $\{d\}$ itself is w-admissible, the conditions of the above proposition are met. We thus consider the reduct F^E:

F^E:
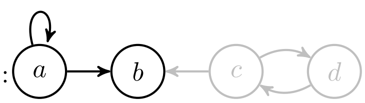

Now $D = \{b\}$ is not attacked by a w-admissible argument (since a is a self-attacker) and is itself w-admissible in F^E. Hence $X = E \cup D$ is w-defended by E. Thus $\{b\}$ is not w-complete (but of course $\{b, d\}$ is). It is thus easy to verify that $co^w(F) = \{\emptyset, \{c\}, \{b, d\}\}$. $\diamond$

For more details regarding the definition and basic properties of weak admissibility we refer the reader to (Baumann, Brewka, and Ulbricht 2020b).

Decision Problems and Complexity Classes

For an AF $F = (A, R)$ and a semantics σ, we say an argument $a \in A$ is *credulously accepted* (*skeptically accepted*) in F w.r.t. σ if $a \in \bigcup \sigma(F)$ ($a \in \bigcap \sigma(F)$). The corresponding decision problems for a semantics σ, given an AF F and argument a, are as follows: Credulous Acceptance $Cred_\sigma$: Deciding whether a is credulously accepted in F w.r.t. σ; Skeptical Acceptance $Skept_\sigma$: Deciding whether a is skeptically accepted in F w.r.t. σ. We also consider the following decision problems, given an AF F: Verification of an extension Ver_σ: Deciding whether a set of arguments is in $\sigma(F)$; and Existence of a non-empty extension $NEmpty_\sigma$: Deciding whether $\sigma(F)$ contains a non-empty set.

σ	$Cred_\sigma$	$Skept_\sigma$	Ver_σ	$NEmpty_\sigma$
ad	NP-c	trivial	in P	NP-c
co	NP-c	P-c	in P	NP-c
gr	P-c	P-c	P-c	in P
pr	NP-c	Π_2^P-c	coNP-c	NP-c
ad^w	PSPACE-c	trivial	PSPACE-c	PSPACE-c
co^w	PSPACE-c	PSPACE-c	PSPACE-c	PSPACE-c
gr^w	PSPACE-c	PSPACE-c	PSPACE-c	PSPACE-c
pr^w	PSPACE-c	PSPACE-c	PSPACE-c	PSPACE-c

Table 1: Complexity of classical and weak-admissible based semantics ("c" is used as shorthand for "complete").

Finally, we assume the reader to be familiar with the basic concepts of computational complexity theory (see, e.g. (Dvořák and Dunne 2018)) as well as the standard classes P, NP as well as coNP. In addition we consider the class $\Pi_2^P = \text{coNP}^{\text{NP}}$ of problems that can be solved in non-deterministic polynomial time when the algorithm has access to an NP oracle. Finally, PSPACE contains the problems that can be solved using only polynomial space of memory. We have $\text{P} \subseteq \text{NP}/\text{coNP} \subseteq \Pi_2^P \subseteq \text{PSPACE}$.

Complexity Analysis

In this section, we investigate the complexity of the standard decision problems in argumentation for the four semantics based on weak admissibility. Our results are summarized and compared with the classical ones (Dvořák and Dunne 2018) in Table **??**.

Membership Results

We provide an algorithm that runs in PSPACE and closely follows the definition of w-admissibility.

Lemma 0.10. *Ver_{ad^w} is in* PSPACE.

Proof. An algorithm for verifying that $E \in ad^w(F)$ proceeds as follows: (1) test whether $E \in cf(F)$; if not return false, (2) compute the reduct F^E, (3) iterate over all subsets S of F^E that contain at least one attacker of E and test whether S is w-admissible; if so return false; else return true. Notice that the last step involves recursive calls. However, the size of the considered AF is decreasing in each step and thus the recursion depth is in $O(n)$. Moreover, we only need to store the current AF as well as the set S to verify. Finally, iterating over all subsets of an AF can be done in PSPACE as well. Hence, the above algorithm is in PSPACE. $\square$

Given that verification is in PSPACE we can adapt standard algorithms to obtain the PSPACE membership of the other problems. Notice that $Skept_{ad^w}$ is always false as the empty-set is always w-admissible.

Proposition 0.11. *For $\sigma \in \{gr^w, ad^w, co^w, pr^w\}$, $Cred_\sigma$, $Skept_\sigma$, Ver_σ, and $NEmpty_\sigma$ can be solved in* PSPACE.

Proof. $Ver_{ad^w} \in$ PSPACE is by Lemma 0.10. The other memberships are by the following algorithms that can be easily implemented in PSPACE with calls to other PSPACE problems , e.g. Ver_{ad^w}, and thus are themselves in PSPACE.

Ver_{pr^w} can be solved by first verifying that the set is w-admissible and then iterating over all super-sets and verifying that they are not w-admissible.

$Ver_{co^w} \in$ PSPACE: To test whether a set E is w-complete, first test whether it is w-admissible, then compute $Cred = \bigcup ad^w(F^E)$ (which is in PSPACE as we show below), and finally test for each set $D \subseteq A \setminus E$ whether it is w-defended by E. The latter can be done by first testing whether $Cred$ attacks D and then iterating over all $D' \supsetneq D$ and test $D' \in ad^w\left(F^E\right)$ (which by the above is in PSPACE). If none of the sets D is w-defended by E then E is w-complete and thus we obtain $Ver_{co^w} \in$ PSPACE. $Ver_{gr^w} \in$ PSPACE: To test whether a set E is w-grounded, we first test whether it is w-complete, and then that each E' with $E' \subsetneq E$ is not w-complete.

For $Cred_\sigma$ with $Ver_\sigma \in$ PSPACE: we iterate over all sets of the arguments that contain the query argument and test whether the set is a σ-extension. As soon as we find a subset that is a σ-extension we can stop and return that the argument is credulously accepted. Otherwise if none of the sets is a σ-extension the argument is not credulously accepted.

For $Skept_\sigma$ we iterate over all subsets of the arguments that do not contain the query argument and test whether the set is a σ-extension. As soon as we find a subset that is a σ-extension we can stop and return that the argument is not skeptically accepted. Otherwise if none of the sets is a σ-extension the argument is skeptically accepted.

For $NEmpty_\sigma$ we iterate over all non-empty subsets of the arguments and test whether the set is a σ-extension. If one of them is a σ-extension we terminate and return true otherwise we return false. □

Hardness Results

We show hardness by a reduction from the PSPACE-complete problem of deciding whether a QBF is valid. To this end we consider QBFs of the form

$$\Phi = \forall x_n \exists x_{n-1} \ldots \forall x_2 \exists x_1 : \phi(x_1, x_2, \ldots, x_{n-1}, x_n).$$

Notice that Φ might start with a universal or existential quantifier and then alternates between universal and existential quantifiers after each variable and ends with an existential quantifier. ϕ is a propositional formula in CNF given by a set of clauses C, i.e, $\phi = \bigwedge_{c \in C} \bigvee_{l \in c} l$. We call a QBF starting with a universal quantifier a $\forall$-QBF and a QBF starting with an existential quantifier an $\exists$-QBF. Finally, observe that we named variables in reverse order to avoid renaming variables in our proofs by induction.

We start with a reduction that maps QBFs to AFs such that the validity of the QBF can be read off by inspecting the w-admissible sets of the AF. We will later extend this reduction to encode the specific decision problems under our considerations.

Reduction 0.12. Given a QBF Φ with propositional formula $\phi(x_1, \ldots, x_n)$ we define the AF $G_\Phi = (A, R)$ with

$$A = \{x_i, \bar{x}_i, p_i \mid 1 \leq i \leq n\} \cup \{c \mid c \in C\} \text{ and}$$

$$\begin{aligned} R = &\{(x_i, \bar{x}_i), (\bar{x}_i, x_i) \mid 1 \leq i \leq n\} \cup \\ &\{(x_i, x_{i+1}), (x_i, \bar{x}_{i+1}) \mid 1 \leq i < n\} \cup \\ &\{(\bar{x}_i, x_{i+1}), (\bar{x}_i, \bar{x}_{i+1}) \mid 1 \leq i < n\} \cup \\ &\{(x_i, c) \mid x_i \in c \in C\} \cup \{(\bar{x}_i, c) \mid \neg x_i \in c \in C\} \cup \\ &\{(c, x_1), (c, \bar{x}_1) \mid c \in C\} \cup \{(p_i, p_{i+1}) \mid 1 \leq i < n\} \cup \\ &\{(x_i, p_i), (\bar{x}_i, p_i) \mid 1 \leq i \leq n\} \cup \\ &\{(p_i, x_{i-1}), (p_i, \bar{x}_{i-1}) \mid 2 \leq i \leq n\} \cup \{(p_1, c) \mid c \in C\}. \end{aligned}$$

Example 0.13. Let us consider the valid QBF $\forall x_2 \exists x_1 : \phi$ with $\phi = c_1 \wedge c_2 = (\neg x_2 \vee x_1) \wedge (x_2 \vee \neg x_1)$ and apply Reduction 0.12 to obtain an AF F. It will be convenient to think of several layers, each one induced by a variable occurring in the QBF at hand. We thus have two layers here, with x_i and $\bar{x}_i$ attacking each other in the expected way and each layer attacked by its predecessor. The x-arguments attack the c-arguments in the natural way. The c-arguments attack the X_1 layer only. The arguments p_1 and p_2 induce odd cycles to forbid certain possible extensions. Schematically, this looks as follows.

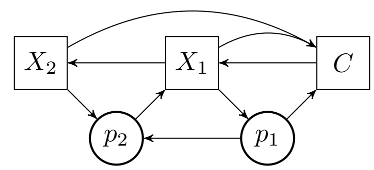

In detail, Reduction 0.12 applied to our QBF yields:

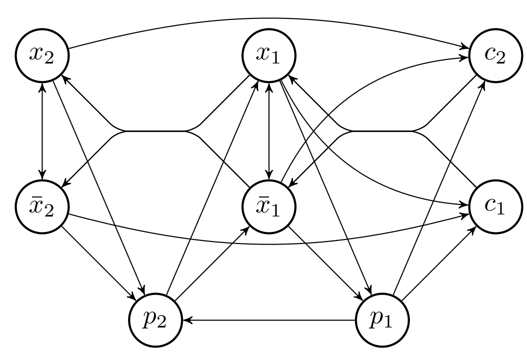

Now regarding our QBF note that setting x_2 to true requires x_1 is to be true as well and setting x_2 to false requires x_1 to be false. This translates to F as follows: Take $E = \{\bar{x}_2\}$, corresponding to setting x_2 to false. The set E is not w-admissible in F. To see this, consider the reduct F^E: Here $\{\bar{x}_1\}$ (corresponding to $\neg x_1$ in the QBF) is w-admissible in F^E (even admissible) and attacks $\bar{x}_2$ witnessing that $E \notin ad^w(F)$. Similarly, $\{x_2\}$ is not w-admissible since it is attacked by x_1 in the corresponding reduct. ◇

Let us collect some properties we require:

Proposition 0.14. *For a QBF Φ,*

1. *if Φ is of the form $\exists x_n \forall x_{n-1} \ldots \exists x_1 : \phi(x_1, x_2, \ldots, x_n)$ we have that $ad^w(G_\Phi) \cap \{\{x_n\}, \{\bar{x}_n\}\} \neq \emptyset$ if Φ is valid and $ad^w(G_\Phi) = \{\emptyset\}$ otherwise; and*
2. *if Φ is of the form $\forall x_n \exists x_{n-1} \ldots \exists x_1 : \phi(x_1, x_2, \ldots, x_n)$ we have that $ad^w(G_\Phi) = \{\emptyset\}$ if Φ is valid and $ad^w(G_\Phi) \cap \{\{x_n\}, \{\bar{x}_n\}\} \neq \emptyset$ otherwise.*

Moreover, in both cases $ad^w(G_\Phi) \subseteq \{\{x_n\}, \{\bar{x}_n\}, \emptyset\}$.

We briefly sketch the main ideas of the proof. First, we have that all conflict-free sets E that contain an argument a different from x_n and $\bar{x}_n$ yield a reduct G_Φ^E with unattacked argument b that attacks a in G_Φ and thus E is not w-admissible. That is, $\{x_n\}$, $\{\bar{x}_n\}$, and $\emptyset$ are the only candidates for being w-admissible. The remainder of the proof is then by induction on the number of variables n, starting with $n = 1$. In the induction step we exploit that when considering one of the sets $E = \{x_n\}$, $E = \{\bar{x}_n\}$ respectively, we have that the reduct G_Φ^E corresponds to the AF G'_Φ where Φ' is the QBF we obtain from Φ when eliminating the variable x_n by replacing it by $\top$, $\bot$ respectively. That is, we have that $\{x_n\}$ is weakly admissible in G_Φ iff neither $\{x_{n-1}\}$ nor $\{\bar{x}_{n-1}\}$ is weakly admissible in G'_Φ and as Φ' has only $n-1$ variables one can exploit the induction hypothesis.

We next extend our reduction by two further arguments ϕ and p_{n+1} in order to show our hardness results.

Reduction 0.15. Given a $\forall$-QBF $\Phi = \forall x_n \exists x_{n-1} \dots \exists x_1 : \phi(x_1, \dots, x_n)$ we define $F_\Phi = G_\Phi \cup (\{\phi, p_{n+1}\}, \{(\phi, p_{n+1}), (p_n, p_{n+1}), (p_{n+1}, x_n), (p_{n+1}, \bar{x}_n), (x_n, \phi), (\bar{x}_n, \phi)\})$.

We now formally characterize the potential w-admissible sets in Reduction 0.15.

Lemma 0.16. *For a QBF Φ, $ad^w(F_\Phi) \subseteq \{\emptyset, \{\phi\}\}$.*

Proof. In comparison to Proposition 0.14, we are only left to consider p_{n+1}. The assumption $p_{n+1} \in E \in ad^w(F)$ yields an analogous contradiction: $E \in cf(F)$ implies $\phi, x_n \notin E$ and hence p_{n+1} is attacked by $\{\phi\} \in ad^w(F^E)$. $\square$

Proposition 0.17. *Given a $\forall$-QBF $\Phi = \forall x_n \exists x_{n-1} \dots \exists x_1 : \phi(x_1, x_2, \dots, x_n)$ we have that Φ is valid if and only if $ad^w(F_\Phi) = \{\emptyset, \{\phi\}\}$, and $ad^w(F_\Phi) = \{\emptyset\}$ otherwise.*

Proof. We have that the empty-set is always w-admissible and by Lemma 0.16 that $\{\phi\}$ is the only candidate for being a w-admissible set. Now consider $\{\phi\}$ and the reduct $F_\Phi^{\{\phi\}}$. We have that $F_\Phi^{\{\phi\}} = G_\Phi$ and x_n and $\bar{x}_n$ being the attackers of ϕ. By Proposition 0.14 we have that $\{x_n\}$ or $\{\bar{x}_n\}$ is w-admissible in the reduct iff Φ is not valid. Thus $\{\phi\}$ is w-admissible iff Φ is valid. $\square$

Theorem 0.18. *All of the following problems are* PSPACE*-complete: Cred$_{ad^w}$, Ver$_{ad^w}$, NEmpty$_{ad^w}$, and Cred$_\sigma$, Skept$_\sigma$, Ver$_\sigma$, NEmpty$_\sigma$ for any $\sigma \in \{co^w, gr^w, pr^w\}$.*

Proof. The membership results are by Proposition 0.11. The hardness results are all by Reduction 0.15 and Proposition 0.17. It only remains to state the precise problem instances that are equivalent to testing the validity of the $\forall$-QBF Φ. First, consider *Cred*$_{ad^w}$ = *Cred*$_{co^w}$ = *Cred*$_{pr^w}$. In the AF F_Φ we have that ϕ is credulously accepted w.r.t. w-admissible semantics iff $\{\phi\} \in ad^w(F_\Phi)$ iff Φ is valid. Now, consider *Ver*$_{ad^w}$ and *Ver*$_{pr^w}$. We have that $\{\phi\} \in ad^w(F_\Phi)$ iff $\{\phi\} \in pr^w(F_\Phi)$ iff Φ is valid. Next, consider *Skept*$_{pr^w}$. We have that ϕ is skeptically accepted iff $pr^w(F_\Phi) = \{\{\phi\}\}$ iff Φ is valid. Moreover, for *NEmpty*$_{ad^w}$ = *NEmpty*$_{pr^w}$, the only w-preferred/w-admissible extension is the empty-set iff Φ is not valid.

For the remaining problems it suffices to show $gr^w(F_\Phi) = co^w(F_\Phi) = \{\{\phi\}\}$ whenever Φ is valid and otherwise, $gr^w(F_\Phi) = co^w(F_\Phi) = \{\emptyset\}$. Regarding the former, if Φ is valid, then $ad^w(F_\Phi) = \{\emptyset, \{\phi\}\}$. We show that in this case, $\emptyset \notin co^w(F_\Phi)$. To this end we show that $\emptyset$ w-defends $\{\phi\}$. We may apply Proposition 0.8 to the w-admissible set $E = \emptyset$ and $X = E \dot\cup D = \emptyset \cup \{\phi\} = \{\phi\}$ and see that $\emptyset$ w-defends $\{\phi\}$ since (1) no attacker of ϕ can be w-admissible in $F_\Phi^E = F_\Phi^\emptyset = F_\Phi$, and (2) $\{\phi\}$ itself is w-admissible in $F_\Phi^E = F_\Phi^\emptyset = F_\Phi$. If, on the other hand, Φ is not valid, then $ad^w(F_\Phi) = \{\emptyset\}$ so the only candidate for a w-complete, w-grounded resp., extension is $\emptyset$. Since there is no other w-admissible set in the reduct $F_\Phi^\emptyset = F_\Phi$, $\emptyset$ does not w-defend any set and is thus itself w-complete and hence w-grounded. Hence, we obtain for $\sigma \in \{gr^w, co^w\}$ that Φ is valid iff ϕ is credulously, skeptically respectively, accepted in F_Φ, iff $\{\phi\} \in \sigma(F_\Phi)$ iff F_Φ has a non-empty σ-extensions. Thus, Reduction 0.15 provides a reduction from $\forall$-QBF to all of the considered problems, and as it can be clearly performed in polynomial time, the PSPACE-hardness follows. $\square$

Complexity for Specific Graph-classes

In the previous section we have shown the standard reasoning problems to be computationally hard. A common approach towards tractability is to consider AFs that have a special graph structure (Dunne 2007). To this end, we consider graph classes that have been shown to be tractable fragments for the traditional argumentation semantics and we focus on the problems *Cred*$_\sigma$ and *Skept*$_\sigma$. As we will see, some results follow from the fact that weak-admissible semantics coincide with the standard semantics on certain classes of AFs. However, certain cases require a dedicated analysis.

First, we consider the class of symmetric AFs (A,R) (Coste-Marquis, Devred, and Marquis 2005) which require that if $(a, b) \in R$ then also $(b, a) \in R$. In symmetric AFs we have that conflict-free and admissible sets coincide and thus also w-admissible and conflict-free sets coincide (recall that we always have $cf(F) \supseteq ad^w(F) \supseteq ad(F)$). If we additionally assume that there is no self-attack then all arguments are credulously accepted and an argument is only w-defended if it is not attacked at all. That is, c-defence and w-defence coincide and thus also complete and w-complete as well as grounded and w-grounded semantics coincide.

Lemma 0.19. *For symmetric AFs F we have $ad^w(F) = ad(F)$ and $pr^w(F) = pr(F)$. Moreover, if F has no self-attacks then also $co^w(F) = co(F)$ and $gr^w(F) = gr(F)$.*

By (Baumann, Brewka, and Ulbricht 2020b), we can remove self-attacking arguments without changing the extensions of weakly-admissible based semantics and thus by the known complexity results for the standard semantics (cf. (Dvořák and Dunne 2018)) we obtain that reasoning on symmetric AFs is in polynomial time.

Proposition 0.20. *For symmetric AFs and $\sigma \in \{gr^w, ad^w, co^w, pr^w\}$, Cred$_\sigma$ and Skept$_\sigma$ can be solved in* P.

Next we consider graph classes that are based on the absence of (certain types) of cycles. To this end, we first re-

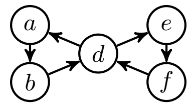

Figure 1: Noeven AF with two pr^w extensions $\{a\}$ and $\{e\}$.

call recent characterizations from (Baumann, Brewka, and Ulbricht 2020a) that are crucial for the following investigations: (a) for odd-cycle free AFs the w-preferred and preferred extensions coincide; (b) for odd-cycle free AFs there is a unique w-grounded extension that consists of the skeptically preferred accepted arguments; and (c) for acyclic AFs gr^w, co^w, pr^w, gr, co and pr coincide.

First, let us consider odd-cycle free AFs. For the standard semantics odd-cycle free AFs are not a tractable fragment but Ver_{pr} becomes tractable and the complexity of $Skept_{pr}$ drops to coNP-c. For w-admissible based semantics we have a similar effect with a more drastic drop in complexity. Given the results of (Baumann, Brewka, and Ulbricht 2020a) and the results for preferred semantics (Dvořák and Dunne 2018) we obtain the following result.

Proposition 0.21. *For odd-cycle free AFs, $Cred_\sigma$ is* NP*-complete for $\sigma \in \{ad^w, co^w, pr^w\}$, and $Skept_{co^w} = Skept_{gr^w} = Cred_{gr^w}$ and $Skept_{pr}$ are* coNP*-complete.*

Next we consider the class of acyclic AFs and exploit the fact that the grounded extension is the only w-preferred extension and can be computed in polynomial time.

Proposition 0.22. *For acyclic AFs and $\sigma \in \{gr^w, ad^w, co^w, pr^w\}$, $Cred_\sigma$ and $Skept_\sigma$ are in* P.

An interesting observation is that even in acyclic AFs w-admissible semantics differs from admissible and strongly admissible (Caminada 2014) sets, while the latter two coincide for AFs from this class.

Next we investigate the class of even-cycle free AFs (noeven AFs) (Dvořák and Dunne 2018) which allow to decide admissible based semantics in polynomial-time. We have that noeven AFs have a unique preferred extension, which is however not true for w-preferred extensions. Consider the AF F in Figure 1 which has only odd cycles. We have $ad^w = \{\emptyset, \{a\}, \{e\}\}$ and thus $pr^w = \{\{a\}, \{e\}\}$ (moreover $co^w = gr^w = \{\{a\}, \{e\}\}$). We next show that noeven AFs are not a tractable fragment for weak admissibility-based semantics. To this end we use the above AF F to adapt the standard reduction from propositional logic in order to obtain a noeven AF. That is, we replace the symmetric attacks, which model setting a variable to either true or false, by sub-AFs that are isomorphic to F (cf. Figure 2).

Reduction 0.23. Given a propositional formula $\phi(x_1, \ldots, x_n) = \bigwedge_{c \in C} \bigvee_{l \in c} l$ we define the AF $H_\phi = (A, R)$ with A and R as follows.

$$\begin{aligned} A =& \{x_i, \bar{x}_i, b_i, d_i, f_i \mid 1 \le i \le n\} \cup \{c \mid c \in C\} \cup \{t, \bar{t}\} \\ R =& \{(x_i, b_i), (b_i, d_i), (d_i, x_i), (d_i, \bar{x}_i), (\bar{x}_i, f_i), (f_i, d_i) \\ & \mid 1 \le i \le n\} \cup \{(c, t) \mid c \in C\} \cup \{(t, \bar{t})\} \cup \\ & \{(x_i, c) \mid x_i \in c \in C\} \cup \{(\bar{x}_i, c) \mid \neg x_i \in c \in C\} \end{aligned}$$

Lemma 0.24. *For every propositional formula ϕ we have that (1) ϕ is satisfiable iff t is credulously accepted in H_ϕ w.r.t. σ, for $\sigma \in \{gr^w, ad^w, co^w, pr^w\}$; and (2) ϕ is unsatisfiable iff $\bar{t}$ is skeptically accepted in H_ϕ w.r.t. σ for $\sigma \in \{gr^w, co^w, pr^w\}$.*

By the above, the NP-hard SAT problem can be reduced to credulous acceptance and the coNP-hard UNSAT problem can be reduced to skeptical acceptance.

Proposition 0.25. *For noeven AFs, $Cred_\sigma$ is* NP*-hard for $\sigma \in \{gr^w, ad^w, co^w, pr^w\}$ and $Skept_\tau$ is* coNP*-hard for $\tau \in \{gr^w, co^w, pr^w\}$.*

Finally, we consider the class of bipartite AFs. We have that bipartite AFs are a sub-class of odd-cycle free AFs and thus we can again use the correspondence of w-preferred and preferred semantics. Moreover, preferred semantics has been shown to be tractable on bipartite AFs (Dunne 2007).

Proposition 0.26. *For bipartite AFs and $\sigma \in \{gr^w, ad^w, co^w, pr^w\}$, $Cred_\sigma$ and $Skept_\sigma$ are in* P.

Proof. The results for pr^w are directly from the corresponding results for pr in (Dunne 2007). Further, by (Baumann, Brewka, and Ulbricht 2020b), we have $Cred_{gr^w} = Skept_{gr^w} = Skept_{pr^w}$ and thus reasoning with gr^w is in P. The results for ad^w, co^w semantics are by its correspondences with tasks for gr^w or pr^w. □

DATALOG Encoding

In this section we provide a DATALOG encoding for w-admissible semantics. Our reduction will generate a polynomial size logic program that falls into the class of non-recursive DATALOG with stratified negation which is known to be PSPACE-complete in terms of program-complexity (Dantsin et al. 2001). For our encoding we consider an AF $F = (A, R)$ with arguments $A = \{a_1, \ldots, a_n\}$. The weakly-admissible sets will be encoded as an n-ary predicate $wadm(e_1, \ldots e_n)$ where variable e_i indicates whether argument a_i is in the extension or not. That is, our fixed database will be over the Boolean domain $\{0, 1\}$.

The encoding closely follows the definition of w-admissible sets, which of course is recursive. To avoid recursion in the DATALOG program we will exploit that the recursion depth is bounded by n. We will introduce n-copies of certain predicates, each of which can only be used on a certain recursion depth of the w-admissible definition.

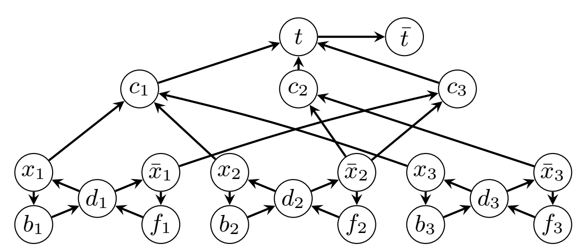

Figure 2: Illustration of the AF H_ϕ, for ϕ with clauses $\{\{x_1, x_2, x_3\}, \{\bar{x}_2, \bar{x}_3\}, \{\bar{x}_1, \bar{x}_2\}\}$.

The input database contains a unary predicate $dom = \{0,1\}$ defining the Boolean domain of our variables and standard predicates that allow to encode the arithmetic operation we are using in our rules, e.g. the binary predicates $equal = \{(0,0),(1,1)\}$ and $leq = \{(0,0),(0,1),(1,1)\}$ (below we denote them via "=" and "≤" symbols).

We will first introduce certain auxiliary predicates which we require in order to define $wadm(e_1, \dots e_n)$. In our encoding we will use variables $\{x_i, y_i, d_i, e_i \mid 1 \leq i \leq n\}$ to represent whether arguments are in certain sets or not. We will use the following shorthands to group variables that together represent a set of arguments: $X = x_1, \dots, x_n$, $Y = y_1, \dots, y_n$, $D = d_1, \dots, d_n$, and $E = e_1, \dots, e_n$. We will use each set of variables to represent a set of arguments such that the i-th variable is set to 1 iff the i-th argument is in the set and 0 otherwise. We start with encoding the subset relation between two sets of arguments X, Y and define a predicate $cf(\cdot)$ encoding conflict-free sets by a rule which for each attack checks that not both incident arguments are in the set.

$$X \subseteq Y \leftarrow \bigwedge_{i=1}^{n} x_i \leq y_i.$$

$$cf(E) \leftarrow \bigwedge_{i=1}^{n} dom(e_i), \bigwedge_{(i,j)\in R} e_i + e_j \leq 1.$$

Notice that the $dom(e_i)$ predicates in the body of the second rule are used to meet the safety condition of DATALOG.

Next we define the predicate $Att(\cdot,\cdot)$ which encodes that the first set of arguments attacks the second set. To this end, for each attack $(a_j, a_k) \in R$, we add the following rule to the DATALOG program (again using $dom(\cdot)$ for safety):

$$Att(D,E) \leftarrow \bigwedge_{i=1}^{n} dom(d_i), \bigwedge_{i=1}^{n} dom(e_i), d_j = 1, e_k = 1.$$

In the following we will use $F{\downarrow_X}$ to refer to the sub-AF of F that is given by the arguments in the set represented by X, i.e. $F{\downarrow_X} = (A', R \cap (A' \times A'))$ with $A' = \{a_i \mid x_i = 1\}$. We define a predicate $Range(X, E, D)$ that defines the range D of an extension E in the AF $F{\downarrow_X}$.

$$Range(X,E,D) \leftarrow E \subseteq D, \bigwedge_{(i,j)\in R}^{n} \min(e_i, x_j) \leq d_j,$$
$$\bigwedge_{i=1}^{n} (d_i \leq e_i + \max_{(j,i)\in R} e_j),\ D \subseteq X.$$

The first constraint ensures that each argument in E is also in the range. The second constraint ensures that each argument in $F{\downarrow_X}$ that is attacked by E is in the range (but makes no statement about arguments not in $F{\downarrow_X}$). The third constraint encodes that an argument is only in the range if it is in E or attacked by E and the final constraint ensures that only arguments in $F{\downarrow_X}$ can be in the range.

We are now ready to encode w-admissible semantics. In a first step we treat the reduct operation and use a predicate $Red(X, E, Y)$ encoding that when we are in the sub-AF $F{\downarrow_X}$ and build the reduct for the argument set E we obtain the sub-AF $F{\downarrow_Y}$.

$$Red(X,E,Y) \leftarrow Range(X,E,D), \bigwedge_{i=1}^{n} y_i = x_i - d_i$$

$Range(X, E, D)$ defines the range of E within the subframework $F{\downarrow_X}$ and the second constraint makes sure that exactly those argument which are in X but not in the range of E are included in the reduct $F{\downarrow_Y}$ (notice that by the definition of $Range$ we have $d_i \leq x_i$).

In order to define the predicate $wadm(E)$ we introduce predicates $P_i(X, E), 1 \leq i \leq n$ that encode the w-admissible sets of the reducts on the i-th recursion level. Recall that the recursion depth of the computation is bounded by n. The variables X in $P_i(X, E)$ encode the arguments of the reduct and the variables E encode the extension, i.e. E represents a w-admissible set of $F{\downarrow_X}$. The initial AF F corresponds to the reduct containing all arguments.

$$wadm(E) \leftarrow P_1(1, \dots, 1, E).$$

Next we define the w-admissible sets of each reduct. We first state the rules for $1 \leq i \leq n-1$ and then consider the special case P_n with at most one argument in the reduct.

$$P_i(X,E) \leftarrow E \subseteq X, cf(E), \texttt{not}\ Q_i(X,E).$$
$$Q_i(X,E) \leftarrow Red(X,E,Y), D \subseteq Y, Att(D,E), P_{i+1}(Y,D).$$
$$P_n(X,E) \leftarrow \sum_{i=1}^{n} x_i \leq 1, E \subseteq X, cf(E).$$

The first two rules are a direct encoding of the definition of w-admissible sets. That is, a set E is weakly admissible in F if it is conflict-free in F and there is no weakly-admissible set D in the reduct F^E that attacks that E. The last rule covers the special case at recursion depth n we have that at most one argument is left in the reduct.

Notice that our DATALOG encoding is indeed nonrecursive and thus can be solved in PSPACE.

Conclusion

In this paper, we investigated the computational complexity of the standard reasoning problems for weakly admissibility-based semantics and showed that all of them, except the trivial skeptical acceptance for ad^w, are PSPACE-complete in general. In the light of this high computational complexity we investigated graph classes as tractable fragments and as a first step towards suitable algorithms we provided a DATALOG encoding for w-admissible semantics. Directions for future include: (a) Studying the potential for fixed-parameter tractable algorithms in particular backdoor approaches on top of the tractable fragments (Dvořák, Ordyniak, and Szeider 2012; Dvořák, Pichler, and Woltran 2012). (b) Settling the exact complexity of weak-admissible based semantics for noeven AFs. (c) Extend the DATALOG approach to the remaining semantics, and provide and evaluate a corresponding implementation. We envision a decomposition approach where the components belonging to easier fragments are handed over to dedicated solvers and an extended version of our DATALOG encoding takes care of combining the results with the computation of the remaining parts.

Acknowledgements

This work was supported by the Vienna Science and Technology Fund (WWTF) through project ICT19-065, the Austrian Science Fund (FWF) through project P30168, and the German Federal Ministry of Education and Research (BMBF, 01/S18026A-F) by funding the competence center for Big Data and AI "ScaDS.AI Dresden/Leipzig". We thank the reviewers of this paper and of a earlier version (Dvořák, Ulbricht, and Woltran 2020) for their valuable comments.

References

Amendola, G.; and Ricca, F. 2019. Paracoherent Answer Set Semantics meets Argumentation Frameworks. *Theory Pract. Log. Program.* 19(5-6): 688–704. doi:10.1017/S1471068419000139.

Baroni, P.; Caminada, M.; and Giacomin, M. 2011. An introduction to argumentation semantics. *The Knowledge Engineering Review* 26: 365–410. doi:10.1017/S0269888911000166.

Baroni, P.; Caminada, M.; and Giacomin, M. 2018. Abstract Argumentation Frameworks and Their Semantics. In *Handbook of Formal Argumentation*. College Publications.

Baumann, R.; Brewka, G.; and Ulbricht, M. 2020a. Comparing Weak Admissibility Semantics to their Dung-style Counterparts - Reduct, Modularization, and Strong Equivalence in Abstract Argumentation. In *Proc. KR*, 79–88. doi:10.24963/kr.2020/9.

Baumann, R.; Brewka, G.; and Ulbricht, M. 2020b. Revisiting the Foundations of Abstract Argumentation: Semantics Based on Weak Admissibility and Weak Defense. In *Proc. AAAI*, 2742–2749. AAAI Press. doi:10.1609/aaai.v34i03.5661.

Bodanza, G. A.; and Tohmé, F. A. 2009. Two approaches to the problems of self-attacking arguments and general odd-length cycles of attack. *Journal of Applied Logic* 7(4): 403–420. doi:10.1016/j.jal.2007.06.012.

Cadoli, M.; Eiter, T.; and Gottlob, G. 2005. Complexity of propositional nested circumscription and nested abnormality theories. *ACM Trans. Comput. Log.* 6(2): 232–272. doi:10.1145/1055686.1055688.

Cadoli, M.; and Schaerf, M. 1993. A Survey of Complexity Results for Nonmonotonic Logics. *J. Log. Program.* 17(2/3&4): 127–160. doi:10.1016/0743-1066(93)90029-G.

Caminada, M. 2014. Strong Admissibility Revisited. In *Proc. COMMA*, 197–208. IOS press. doi:10.3233/978-1-61499-436-7-197.

Coste-Marquis, S.; Devred, C.; and Marquis, P. 2005. Symmetric Argumentation Frameworks. In *Proc. ECSQARU*, 317–328. Springer. doi:10.1007/11518655_28.

Dantsin, E.; Eiter, T.; Gottlob, G.; and Voronkov, A. 2001. Complexity and expressive power of logic programming. *ACM Comput. Surv.* 33(3): 374–425. doi:10.1145/502807.502810.

Dauphin, J.; Rienstra, T.; and van der Torre, L. 2020. A Principle-Based Analysis of Weakly Admissible Semantics. In *Proc. COMMA*. IOS press. doi:10.3233/FAIA200502.

Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77(2): 321–357. doi:10.1016/0004-3702(94)00041-X.

Dunne, P. E. 2007. Computational properties of argument systems satisfying graph-theoretic constraints. *Artificial Intelligence* 171(10-15): 701–729. doi:10.1016/j.artint.2007.03.006.

Dvořák, W.; and Dunne, P. E. 2018. Computational Problems in Formal Argumentation and their Complexity. In *Handbook of Formal Argumentation*. College Publications.

Dvořák, W.; and Gaggl, S. A. 2016. Stage semantics and the SCC-recursive schema for argumentation semantics. *J. Log. Comput.* 26(4): 1149–1202. doi:10.1093/logcom/exu006.

Dvořák, W.; Ordyniak, S.; and Szeider, S. 2012. Augmenting tractable fragments of abstract argumentation. *Artificial Intelligence* 186: 157–173. doi:10.1016/j.artint.2012.03.002.

Dvořák, W.; Pichler, R.; and Woltran, S. 2012. Towards fixed-parameter tractable algorithms for abstract argumentation. *Artificial Intelligence* 186: 1–37. doi:10.1016/j.artint.2012.03.005.

Dvořák, W.; Ulbricht, M.; and Woltran, S. 2020. Weak Admissibility is PSPACE-complete. In *Proc. NMR*, 100–109. URL https://nmr2020.dc.uba.ar/WorkshopNotes.pdf.

Eiter, T.; and Gottlob, G. 1996. The Complexity of Nested Counterfactuals and Iterated Knowledge Base Revisions. *J. Comput. Syst. Sci.* 53(3): 497–512. doi:10.1006/jcss.1996.0083.

Eiter, T.; and Gottlob, G. 2006. Reasoning under minimal upper bounds in propositional logic. *Theor. Comput. Sci.* 369(1-3): 82–115. doi:10.1016/j.tcs.2006.07.054.

Fazzinga, B.; Flesca, S.; and Furfaro, F. 2020. Embedding The Trust Degrees of Agents in Abstract Argumentation. In *Proc. ECAI*, 737–744. IOS Press. doi:10.3233/FAIA200161.

Gaggl, S. A.; and Woltran, S. 2013. The cf2 argumentation semantics revisited. *J. Log. Comput.* 23(5): 925–949. doi:10.1093/logcom/exs011.

Papadimitriou, C. H. 1991. On Selecting a Satisfying Truth Assignment (Extended Abstract). In *Proc. FOCS*, 163–169. IEEE Computer Society. doi:10.1109/SFCS.1991.185365.

Thomas, M.; and Vollmer, H. 2010. Complexity of non-monotonic logics. *Bulletin of the EATCS* 102: 53–82. URL http://eatcs.org/beatcs/index.php/beatcs/article/view/154.

van der Torre, L.; and Vesic, S. 2017. The Principle-Based Approach to Abstract Argumentation Semantics. *FLAP* 4(8). URL http://www.collegepublications.co.uk/downloads/ifcolog00017.pdf.