Extension-Based Semantics for Incomplete Argumentation Frameworks

Jean-Guy Mailly(✉)

LIPADE, University of Paris, Paris, France
jean-guy.mailly@u-paris.fr

Abstract. Incomplete Argumentation Frameworks (IAFs) have been defined to incorporate some qualitative uncertainty in abstract argumentation: information such as "I am not sure whether this argument exists" or "I am not sure whether this argument attacks that one" can be expressed. Reasoning with IAFs is classically based on a set of completions, *i.e.* standard argumentation frameworks that represent the possible worlds encoded in the IAF. The number of these completions may be exponential with respect to the number of arguments in the IAF. This leads, in some cases, to an increase of the complexity of reasoning, compared to the complexity of standard AFs. In this paper, we follow an approach that was initiated for Partial AFs (a subclass of IAFs), which consists in defining new forms of conflict-freeness and defense, the properties that underly the definition of Dung's semantics for AFs. We generalize these semantics from PAFs to IAFs. We show that, among three possible types of admissibility, only two of them satisfy some desirable properties. We use them to define two new families of extension-based semantics. We study the properties of these semantics, and in particular we show that their complexity remains the same as in the case of Dung's AFs. Finally, we propose a logical encoding of these semantics, that paves the way to the development of SAT-based solvers for reasoning with our new semantics for IAFs.

Keywords: Abstract argumentation · Uncertainty

1 Introduction

Abstract argumentation has been a major subfield of Knowledge Representation and Reasoning since the seminal paper by Dung [14]. However, although it is very appealing, Dung's framework is limited in the kind of information that can be modeled: only (abstract) arguments and attacks between them. For this reason, many generalization of this framework have been proposed, introducing the notion of support between arguments [2], weighted attacks [15] or arguments [26], preferences between arguments [1], and so on.

Among these generalizations of Dung's framework, a very natural research direction is the introduction of uncertainty in the model. Indeed, uncertainty

P. Baroni et al. (Eds.): CLAR 2021, LNAI 13040, pp. 322–341, 2021.
https://doi.org/10.1007/978-3-030-89391-0_18

is omnipresent in real world, and must be taken into account in the modeling of agents that reason about their environment or about other agents. Moreover, when arguments are generated from natural language processing [21], the nuances that exist in natural language are likely to be sources of uncertainty [5] that should appear in the formal model. Two directions have been followed for integrating uncertainty in abstract argumentation: quantitative representation of uncertainties (*e.g.* probabilities [19,22]) and qualitative ones [6,11,12]. While quantitative representation of uncertainty is valuable when it is available, allowing fine grained reasoning about uncertainty, it may not be available in many realistic cases. The study of qualitative models of uncertainty is thus of utter importance for the design of AI systems.

In this paper, we follow this direction. Qualitative uncertainty in abstract argumentation was originally studied in a context of Argumentation Framework (AF) merging [12]: Partial Argumentation Frameworks (PAFs) are AFs with possible ignorance about the existence of some attacks. Semantics dedicated to these PAFs were then defined in [11]. However, most of the work in this field focuses on a generalization of PAFs, namely Incomplete AFs (IAFs), and reasoning is based on *completions*. A completion is an argumentation framework that represents one of the (uncertain) options encoded in the IAF. Classical reasoning tasks are then adapted in two versions: the possible view (is some property true for some completion?) and the necessary view (is some property true for each completion?). However, the number of completions is (in the worst case) exponential in the number of arguments. This means that various reasoning problems are harder for IAFs than their counterpart for standard AFs [6,8,18].

In this paper, we follow the approach initiated by [11]: we define new forms of conflict-freeness and defense based on the different types of information in an IAF. The combination of a notion of conflict-freeness and a notion of defense yields a notion of admissibility; we show that among the three possible variants of admissibility, only two of them satisfy some desirable property, namely Dung's Fundamental Lemma. This lemma states, in classical AFs, that an admissible set remains admissible if an argument defended by it is added to the set. From the two "fundamental" notions of admissibility for IAFs, we define variants of the classical complete, preferred and stable semantics. We study some properties of these semantics, and we show that their complexity remains the same as in the standard AF case. Finally, we propose logical encodings, in the same vein as [9], that pave the way to SAT-based implementations for reasoning with our new semantics.

The rest of the paper is organized as follows. Section 2 describes the background notions on abstract argumentation. In Sect. 3, we define our new semantics and study some of their properties, in particular the satisfaction of the Fundamental Lemma, and some inclusion relations between them. In Sect. 4, we show that the complexity remains the same as in the standard AF case,[1] and we provide a logical encoding for our semantics. Finally, Sect. 5 describes some related work, and Sect. 6 concludes the paper.

[1] At the exception of skeptical acceptability under the complete semantics, for which we do not have a tight complexity result yet.

2 Background

2.1 Abstract Argumentation Frameworks

Abstract argumentation is the study of relations between abstract pieces of information called *arguments*; the internal nature of arguments, as well as their origin, is considered as irrelevant. Only the interactions between arguments are considered in order to determine which arguments are acceptable or not. The most classical type of relationship is the so-called *attack* relation, that expresses a contradiction between arguments. An attack is generally directed from one argument to another one, meaning that the first one somehow *defeats* the second one. The seminal paper [14] has launched the strong interest for abstract argumentation in the last 25 years. In this section, we formally introduce this abstract framework and how it is used for reasoning.

We suppose the existence of a finite set of arguments $\mathbf{A}$.

Definition 1 (Argumentation Framework). *An argumentation framework (AF) is a pair* $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ *with* $\mathcal{A} \subseteq \mathbf{A}$ *the set of* arguments *and* $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ *the set of* attacks*.*

For $a, b \in \mathcal{A}$, we say that *a attacks b* if $(a, b) \in \mathcal{R}$. If b attacks some $c \in \mathcal{A}$, then *a defends c* against b. Similarly, a set $S \subseteq \mathcal{A}$ attacks (resp. defends) an argument b if there is some $a \in S$ that attacks (resp. defends) b.

Example 1. Figure 1 depicts an AF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$, with $\mathcal{A} = \{a, b, c\}$ (*i.e.* the nodes of the graph) and $\mathcal{R} = \{(b, a), (b, c), (c, b)\}$ (*i.e.* the edges of the graph).

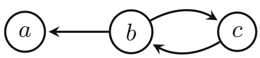

Fig. 1. An example of AF $\mathcal{F}$

The acceptability of arguments is classically evaluated through the concept of *extensions*, *i.e.* sets of arguments that are jointly acceptable. This form of joint acceptance can be interpreted as defining a coherent point of view about the argumentative scenario that is represented by the AF. Different semantics have been defined, that yield different sets of extensions. The usual semantics are based on two main principles: conflict-freeness and admissibility.

Definition 2 (Conflict-freeness and Admissibility). *Given* $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ *an AF, the set* $S \subseteq \mathcal{A}$ *is*

- conflict-free *iff* $\forall a, b \in S$, $(a, b) \notin \mathcal{R}$*;*
- admissible *iff it is conflict-free and* $\forall a \in S$, $\forall b \in \mathcal{A}$ *s.t.* $(b, a) \in \mathcal{R}$, $\exists c \in S$ *s.t.* $(c, b) \in \mathcal{R}$*.*

The meaning of conflict-freeness is quite easy to understand: we do not want to accept together arguments that are conflicting. Admissibility corresponds to a notion of "self-defense": a (conflict-free) set of arguments must be able to defend itself against external attacks in order to be considered as a valid point of view. We use $\mathsf{cf}(\mathcal{F})$ (resp. $\mathsf{ad}(\mathcal{F})$) to denote the set of conflict-free (resp. admissible) sets of an AF $\mathcal{F}$.

These principles are usually considered to be too weak to define semantics, but the classical semantics are based on them.[2] We recall now the definition of these semantics:

Definition 3 (Admissibility-based Semantics). *Given $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ an AF, the admissible set $S \subseteq \mathcal{A}$ is*

- *a* complete *extension iff S contains all the arguments that it defends;*
- *a* preferred *extension iff S is a $\subseteq$-maximal admissible set;*
- *a* grounded *extension iff S is a $\subseteq$-minimal complete extension.*

A fourth semantics is defined by Dung, that does not directly rely on the notion of admissibility:

Definition 4 (Stable Semantics). *Given $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ an AF, the conflict-free set $S \subseteq \mathcal{A}$ is a* stable extension *iff $\forall a \in \mathcal{A} \setminus S$, S attacks a.*

We use $\mathsf{co}(\mathcal{F})$, $\mathsf{pr}(\mathcal{F})$, $\mathsf{gr}(\mathcal{F})$ and $\mathsf{st}(\mathcal{F})$ for the sets of (respectively) complete, preferred, grounded and stable extensions. Among their basic properties:

- $\forall \mathcal{F}, |\sigma(\mathcal{F})| \geq 1$ for $\sigma \in \{\mathsf{co}, \mathsf{pr}, \mathsf{gr}\}$;
- $\forall \mathcal{F}, |\mathsf{gr}(\mathcal{F})| = 1$;
- $\forall \mathcal{F}, \mathsf{st}(\mathcal{F}) \subseteq \mathsf{pr}(\mathcal{F}) \subseteq \mathsf{co}(\mathcal{F})$.

The last point implies that stable extensions are admissible sets as well, even if they are not explicitly defined through admissibility.

Example 2. Considering again $\mathcal{F}$ from Example 1; its extensions for the four semantics defined previously are given in Table 1 (second column).

For further details about these semantics, as well as other semantics that have been defined subsequently, we refer the reader to [4,14].

Given an argumentation framework and a semantics, classical reasoning tasks include the verification that a given set of arguments is an extension, and that a given argument is credulously or skeptically acceptable, *i.e.* belongs to some or each extension. Formally:

σ-Ver Given an AF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ and $S \subseteq \mathcal{A}$, is S a σ-extension of $\mathcal{F}$?
σ-Cred Given an AF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ and $a \in \mathcal{A}$, does a belong to some σ-extension of $\mathcal{F}$?

[2] However, let us notice that we will sometimes include them in the family of studied semantics, for homogeneity of the presentation, *e.g.* in the complexity results (see Sect. 4.1).

σ-Skep Given an AF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ and $a \in \mathcal{A}$, does a belong to each σ-extension of $\mathcal{F}$?

We use $\mathsf{Cred}_\sigma(\mathcal{F})$ (resp. $\mathsf{Skep}_\sigma(\mathcal{F})$) to denote the set of credulously (resp. skeptically) accepted arguments of $\mathcal{F}$, *i.e.* those for which the answer to σ-Cred (resp. σ-Skep) is "YES".

Example 3. The credulously and skeptically accepted arguments in $\mathcal{F}$ from Example 1 are given in Table 1 (third and fourth columns).

Table 1. Extensions and acceptable arguments of $\mathcal{F}$, for $\sigma \in \{\mathsf{gr}, \mathsf{st}, \mathsf{co}, \mathsf{pr}\}$.

Semantics σ	$\sigma(\mathcal{F})$	$\mathsf{Cred}_\sigma(\mathcal{F})$	$\mathsf{Skep}_\sigma(\mathcal{F})$
gr	$\{\emptyset\}$	$\emptyset$	$\emptyset$
st	$\{\{b\}, \{a, c\}\}$	$\{a, b, c\}$	$\emptyset$
co	$\{\emptyset, \{b\}, \{a, c\}\}$	$\{a, b, c\}$	$\emptyset$
pr	$\{\{b\}, \{a, c\}\}$	$\{a, b, c\}$	$\emptyset$

The complexity of these problems for various semantics has been established, see *e.g.* [16] for an overview. The relevant results for this paper are summarized in Table 2. We assume that the reader is familiar with basic notions of complexity theory, otherwise see *e.g.* [3].

Table 2. Complexity of σ-Ver, σ-Cred and σ-Skep, for $\sigma \in \{\mathsf{cf}, \mathsf{ad}, \mathsf{gr}, \mathsf{st}, \mathsf{co}, \mathsf{pr}\}$. $\mathcal{C}$-c means $\mathcal{C}$-complete.

Semantics σ	σ-Ver	σ-Cred	σ-Skep
cf	in L	in L	Trivial
ad	in L	NP-c	Trivial
gr	P-c	P-c	P-c
st	in L	NP-c	coNP-c
co	in L	NP-c	P-c
pr	coNP-c	NP-c	Π_2^P-c

2.2 Qualitative Uncertainty in AFs

Now we present the existing models that incorporate qualitative uncertainty in abstract argumentation.

Incomplete Argumentation Frameworks

Definition 5 (Incomplete Argumentation Framework). *An incomplete argumentation framework (IAF) is a tuple* $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ *where*

- $\mathcal{A} \subseteq \mathbf{A}$ *is the set of* certain arguments*;*
- $\mathcal{A}^? \subseteq \mathbf{A}$ *is the set of* uncertain arguments*;*
- $\mathcal{R} \subseteq (\mathcal{A} \cup \mathcal{A}^?) \times (\mathcal{A} \cup \mathcal{A}^?)$ *the set of* certain attacks*;*
- $\mathcal{R}^? \subseteq (\mathcal{A} \cup \mathcal{A}^?) \times (\mathcal{A} \cup \mathcal{A}^?)$ *the set of* uncertain attacks.

$\mathcal{A}$ and $\mathcal{A}^?$ are disjoint sets of arguments, and $\mathcal{R}$, $\mathcal{R}^?$ are disjoint sets of attacks.

Intuitively, $\mathcal{A}$ and $\mathcal{R}$ correspond, respectively, to arguments and attacks that certainly exist, while $\mathcal{A}^?$ and $\mathcal{R}^?$ are those that may (or may not) actually exist.

Example 4. Figure 2 depicts an IAF $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ with $\mathcal{A} = \{a, b\}$ (plain nodes), $\mathcal{A}^? = \{c\}$ (square dashed node), $\mathcal{R} = \{(c, b)\}$ (plain edge) and $\mathcal{R}^? = \{(b, a)\}$ (dotted edge). It means that the arguments a and b certainly exist, and there is an uncertainty regarding the existence of the attack (b, a). Then, the argument c is uncertain, but if it exists then the attack (c, d) certainly exists as well.

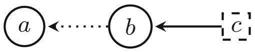

Fig. 2. An example of IAF $\mathcal{I}$

Reasoning with such IAFs is generally made through the notion of completion:

Definition 6 (Completion). *Let $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ be an IAF. A completion of $\mathcal{I}$ is a pair $\langle \mathcal{A}_c, \mathcal{R}_c \rangle$ such that*

- $\mathcal{A} \subseteq \mathcal{A}_c \subseteq \mathcal{A} \cup \mathcal{A}^?$*;*
- $\mathcal{R} \cap (\mathcal{A}_c \times \mathcal{A}_c) \subseteq \mathcal{R}_c \subseteq (\mathcal{R} \cup \mathcal{R}^?) \cap (\mathcal{A}_c \times \mathcal{A}_c)$.

Example 5. Figure 3 depicts the completions of $\mathcal{I}$ from Example 4. $\mathcal{F}_1$ shows the situation where none of the uncertain elements actually exists, while $\mathcal{F}_4$ shows the opposite situation (all the uncertain elements appear). $\mathcal{F}_2$ and $\mathcal{F}_3$ shows the intermediate situations, where only one uncertain element (either the argument c, or the attack (b, a)) exists.

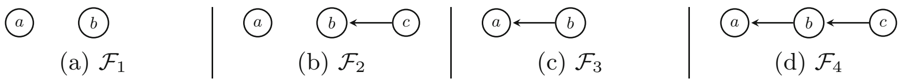

Fig. 3. The completions of $\mathcal{I}$

As seen with the previous example, the number of completions is generally exponential in the size of the IAF. More precisely, it is bounded by $2^{|\mathcal{A}^?|+|\mathcal{R}^?|}$.

Finally, reasoning tasks like credulous acceptance, skeptical acceptance or verification are defined with respect to some or each completion [6,8]: each classical reasoning task has two variants, following the possible view (the property holds in some completion) and the necessary view (the property holds in each completion). These reasoning tasks are, in most cases, computationally harder than their counterpart for standard AFs (under the usual assumption that the polynomial hierarchy does not collapse) [6,8]. This can be explained by the exponential number of completions.

Partial Argumentation Frameworks. Partial Argumentation Frameworks were initially defined as tool in a merging process [12]. They are tuples $\langle A, R, I, N\rangle$ with three binary relations over the set of arguments A: R is the (certain) attack relation, I the ignorance relation, and N the (certain) non-attack relation. Since $N = (A \times A)\setminus(R \cup I)$, a PAF can be identified with only $\langle A, R, I\rangle$. Since the meaning of I is exactly the same as the meaning of $\mathcal{R}^?$, PAFs actually form a subclass of IAFs:[3] any PAF $\langle A, R, I\rangle$ is an IAF $\langle \mathcal{A}, \emptyset, \mathcal{R}, \mathcal{R}^?\rangle$ with $\mathcal{A} = A$, $\mathcal{R} = R$, $\mathcal{R}^? = I$.

Extension-based semantics for PAFs have been defined in [11]. Intuitively, the idea consists in defining different forms of conflict-freeness and defense, and then combine them for defining three types of admissibility. From these new notions of admissibility, the authors define three variants of the preferred semantics, and study their properties. An interesting point is the fact that the complexity remains the same as in Dung's setting, contrary to the other reasoning methods for IAFs. These are the notions that are generalized from PAFs to IAFs in the next section.

3 Generalizing Extension-Based Semantics from PAFs to IAFs

In this section, we follow the same approach as [11] for defining semantics for IAFs. Instead of defining the extensions with respect to the set of completions of the IAF, we will generalize the basic concepts of conflict-freeness and defense to take into account the uncertainty in the IAF. Then, the usual admissibility-based semantics can be defined.

3.1 Conflict-Free and Admissible Sets of IAFs

We follow two approaches for defining conflict-freeness and defense for IAFs:

- Optimistic view: we consider that only certain arguments and attacks are harmful, so keep the definition of conflict-freeness and defense as in Dung's frameworks;
- Pessimistic view: we consider that all attacks are harmful, and must be defended by certain arguments and attacks only.

[3] This subclass was studied under the name *Attack-Incomplete AFs* [7].

By optimistic, we mean that the agent considers (*e.g.*) that $(a, b) \in \mathcal{R}^?$ does not make a a real "threat" against the acceptance of b. Roughly speaking, it means that the agent is tolerant to conflicts if they are uncertain. On the opposite, the pessimistic view means that the agent considers that all uncertain attacks against an argument are real threats against the acceptance of b, and that b must be defended by certain elements only in order to be accepted. Let us formally define the corresponding versions of conflict-freeness and defense.

Definition 7 (Weak and Strong Conflict-freeness). *Let* $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ *be an IAF. The set* $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ *is*

- weakly conflict-free *iff* $\forall a, b \in S \cap \mathcal{A}$, $(a, b) \notin \mathcal{R}$;
- strongly conflict-free *iff* $\forall a, b \in S$, $(a, b) \notin \mathcal{R} \cup \mathcal{R}^?$.

We use $\mathsf{cf}_w(\mathcal{I})$ and $\mathsf{cf}_s(\mathcal{I})$ to denote, respectively, the weakly and strongly conflict-free sets of an IAF $\mathcal{I}$.

Example 6. Figure 4 depicts an IAF $\mathcal{I}_2 = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$, with $\mathcal{A} = \{a, b, d, e\}$, $\mathcal{A}^? = \{c, f\}$, $\mathcal{R} = \{(c, b), (e, b), (e, f)\}$ and $\mathcal{R}^? = \{(b, a), (b, e), (d, e)\}$. The set $\{a, b, c\}$ is weakly conflict-free: the attack from b to a does not violate the weak conflict-freeness since it is uncertain, and the attack from c to b does not violate it either because the attacker (c) is uncertain. It is not strongly conflict-free because of the same two attacks.

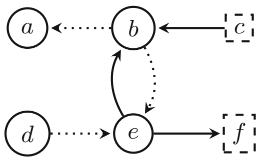

Fig. 4. An example of IAF $\mathcal{I}_2$

Strong conflict-freeness can be regarded as conflict-freeness applied on the "full" graph $\mathcal{F}_{full} = \langle \mathcal{A} \cup \mathcal{A}^?, \mathcal{R} \cup \mathcal{R}^? \rangle$, *i.e.* an AF made from the same arguments and attacks than the IAF, but without any uncertainty. However, weakly conflict-free sets do not correspond to the conflict-free sets of the "minimal" graph $\mathcal{F}_{min} = \langle \mathcal{A}, \mathcal{R} \cap (\mathcal{A} \times \mathcal{A}) \rangle$ (*i.e.* the AF obtained by simply ignoring the uncertain elements): see *e.g.* $\{a, b, c\}$ exhibited in Example 6, which is not a set of arguments in $\mathcal{F}_{min}$ (since $c \notin \mathcal{A}$).

Definition 8 (Weak and Strong Defense). *Let* $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ *be an IAF. Given a set of arguments* $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ *and an argument* $a \in \mathcal{A} \cup \mathcal{A}^?$,

- S weakly defends a *iff* $\forall b \in \mathcal{A}$ *such that* $(b, a) \in \mathcal{R}$, $\exists c \in S \cap \mathcal{A}$ *s.t.* $(c, b) \in \mathcal{R}$;
- S strongly defends a *iff* $\forall b \in \mathcal{A} \cup \mathcal{A}^?$ *such that* $(b, a) \in \mathcal{R} \cup \mathcal{R}^?$, $\exists c \in S \cap \mathcal{A}$ *s.t.* $(c, b) \in \mathcal{R}$.

Example 7. Considering again $\mathcal{I}_2$ from Example 6, we observe that $S = \{a\}$ weakly defends a, since there is no $x \in \mathcal{A}$ s.t. $(x, a) \in \mathcal{R}$. On the contrary, a is not strongly defended by S, because there is no argument in $S \cap \mathcal{A}$ that attacks b. But $S' = \{a, e\}$ strongly defends a: $e \in S' \cap \mathcal{A}$ (certainly) attacks b.

We observe that in the case where $\mathcal{A}^? = \emptyset$, then weak conflict-freeness and defense correspond to the notions of $\mathcal{R}$-conflict-freeness and $\mathcal{R}$-acceptability defined in [11], while the strong versions correspond to $\mathcal{RI}$-conflict-freeness and $\mathcal{RI}$-acceptability. Then, if $\mathcal{R}^? = \emptyset$ also holds, then both weak conflict-freeness and strong conflict-freeness coincide with the classical conflict-freeness [14], while both forms of defense defined here correspond with the classical defense.

For defining a notion of admissibility, we must combine conflict-freeness and defense. In theory, Definitions 7 and 8 induce four notions of admissibility. However, the following result shows that weak conflict-freeness and strong conflict-freeness induce the same notion of admissibility when combined with strong defense.

Proposition 1. *Let $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ be an IAF. Let $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ be a set of arguments such that S is weakly conflict-free and $\forall a \in S$, S strongly defends a. Then S is strongly conflict-free.*

The proof is similar to the proof of [11, Property 1].

Proof. Reasoning towards a contradiction, let us suppose that S is not strongly conflict-free, *i.e.* $\exists a, b \in S$ such that $(a, b) \in \mathcal{R}^?$ (we can exclude the option $(a, b) \in \mathcal{R}$ because S is assumed to be weakly conflict-free). Then, since S strongly defends all its elements, in particular it strongly defends b, so $\exists c \in S$ such that $(c, a) \in \mathcal{R}$. This is a contradiction with the weak conflict-freeness of S. So we can conclude that S is strongly conflict-free.

Now we define the three variants of admissibility.[4]

Definition 9 (Weak, Mixed and Strong Admissibility). *Given $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ an IAF, a set of arguments $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ is*

- weakly admissible *iff S is weakly conflict-free and weakly defends all its elements;*
- mixedly admissible *iff S is strongly conflict-free and weakly defends all its elements;*
- strongly admissible *iff S is strongly conflict-free and strongly defends all its elements.*

The weakly (resp. mixedly, strongly) admissible sets of an IAF $\mathcal{I}$ are denoted by $\mathsf{ad}_w(\mathcal{I})$ (resp. $\mathsf{ad}_m(\mathcal{I})$, $\mathsf{ad}_s(\mathcal{I})$).

The definitions imply that $\mathsf{ad}_s(\mathcal{I}) \subseteq \mathsf{ad}_m(\mathcal{I}) \subseteq \mathsf{ad}_w(\mathcal{I})$, for any IAF $\mathcal{I}$. Also, as in the standard Dung's framework, every IAF has at least one admissible set,

[4] The terminology "strong defense" and "strong admissibility" has been used with another meaning in [10], where it applies to classical AFs, not IAFs.

for all the variations of admissibility. Indeed, for any IAF $\mathcal{I}$, $\emptyset \in \mathsf{ad}_s(\mathcal{I})$. This fact will be useful later to guarantee the existence of extensions for the semantics based on admissibility.

Before going further with the definition of semantics based on these new notions of admissibility, we briefly discuss a property of classical semantics that we believe is important. It is called the *Fundamental Lemma* by Dung [14, Lemma 10]. This lemma states that if a set of arguments S is admissible, and defends an argument a, then $S \cup \{a\}$ is admissible. Besides its technical interest for proving some further results, this lemma describes an intuitive property of argumentation in general: if a point of view (*i.e.* a set of arguments) is seen as valid, then it should be jointly acceptable with any argument that it successfully defends. We thus consider this property as necessary for defining reasonable semantics. With the following lemma, we determine which of the notions of admissibility given in Definition 9 satisfy a notion of "fundamentality" similar to Dung's lemma. More precisely, we show that only weak and strong admissibility are suitable for defining semantics.

Lemma 1 (Fundamental Lemma). *Given $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ an IAF, and $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ a weakly (resp. strongly) admissible set, if S weakly (resp. strongly) defends some $a \in \mathcal{A} \cup \mathcal{A}^?$, then $S \cup \{a\}$ is weakly (resp. strongly) admissible.*

Proof. We first consider weak admissibility. Let us prove that $S \cup \{a\}$ is weakly conflict-free. First of all, notice that if $a \in \mathcal{A}^?$ then the set $S \cup \{a\}$ is weakly conflict-free iff S is weakly conflict-free, since only certain attacks between certain arguments violate weak conflict-freeness. So in the rest of the reasoning we suppose that $a \in \mathcal{A}$. Towards a contradiction, suppose that $S \cup \{a\}$ is not weakly conflict-free. Then, $\exists b \in S \cap \mathcal{A}$ such that, either $(b, a) \in \mathcal{R}$ or $(a, b) \in \mathcal{R}$. In the former case, since S weakly defends a, then there must be a $c \in S \cap \mathcal{A}$ with $(c, b) \in \mathcal{R}$, which is impossible since S is weakly conflict-free. Hence the contradiction. In the latter case ($(a, b) \in \mathcal{R}$), since S is weakly admissible, it must defend b against a, and the same reasoning applies for concluding the impossibility. Thus $S \cup \{a\}$ is weakly conflict-free.

The fact that $S \cup \{a\}$ weakly defends all its elements comes from the fact that S weakly defends all its elements, as well as a. So we conclude that $S \cup \{a\}$ is weakly admissible.

Now, consider S a strongly admissible set that strongly defends some $a \in \mathcal{A} \cup \mathcal{A}^?$. Suppose that $S \cup \{a\}$ is not strongly conflict-free. It means that some $b \in S$ is such that $(b, a) \in \mathcal{R} \cup \mathcal{R}^?$ or $(a, b) \in \mathcal{R} \cup \mathcal{R}^?$. In the first case, the fact that S strongly defends a (against b) means that some $c \in S \cap \mathcal{A}$ attacks b, which violates strong conflict-freeness of S. In the second case, since S strongly defends all its elements, there is a $c \in S \cap \mathcal{A}$ such that $(c, a) \in \mathcal{R}$, which is impossible for similar reasons to the first case. Hence $S \cup \{a\}$ is strongly conflict-free. Finally, the fact that $S \cup \{a\}$ strongly defends all its elements follows the fact that S strongly defends all its elements and a. So we conclude that $S \cup \{a\}$ is strongly admissible.

On the contrary, mixed admissibility does not satisfy a property of fundamentality.

Proposition 2. *There is an IAF* $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$*,* $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ *and an argument* $a \in \mathcal{A} \cup \mathcal{A}^?$ *such that* S *is mixedly admissible,* S *weakly defends* a*, and* $S \cup \{a\}$ *is not mixedly admissible.*

Proof. The IAF given at Fig. 5 provides an example. The set $S = \{b\}$ is mixedly admissible (it is strongly conflict-free, and it has no attacker). S weakly defends a (since there is no $x \in \mathcal{A}$ such that $(x, a) \in \mathcal{R}$, there is actually no need to weakly defend a). But $S \cup \{a\}$ is not strongly conflict-free, hence not mixedly admissible.

Fig. 5. A counter-example about fundamentality of mixed admissibility

Because of this reason, we do not consider mixed admissibility as suitable for defining semantics (*e.g.* mixed preferred or mixed complete semantics).

Example 8. Based on Example 6 and 7, we observe that, in $\mathcal{I}_2$ from Fig. 4, $\{a\}$ is weakly admissible but not strongly admissible. $\{a, e\}$ is not strongly admissible either, because it does not strongly defend e (against the uncertain attack (d, e)). The full sets of weakly and strongly admissible sets of $\mathcal{I}_2$ are given in Table 3.

Table 3. Weakly and Strongly Admissible Sets of $\mathcal{I}_2$.

$\mathsf{ad}_w(\mathcal{I}_2)$	$\emptyset$, $\{a\}$, $\{c\}$, $\{d\}$, $\{e\}$, $\{a,c\}$, $\{a,d\}$, $\{a,e\}$, $\{c,d\}$, $\{c,e\}$, $\{d,e\}$, $\{a,c,d\}$, $\{a,c,e\}$, $\{a,d,e\}$, $\{c,d,e\}$, $\{a,c,d,e\}$
$\mathsf{ad}_s(\mathcal{I}_2)$	$\emptyset$, $\{c\}$, $\{d\}$, $\{c,d\}$

3.2 Admissibility-Based Semantics for IAFs

The classical definitions of Dung's semantics can be adapted to IAFs, based on the two different notions of admissibility identified as suitable in Lemma 1.

Definition 10 (Admissibility-based Semantics). *Given* $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ *an IAF, a weakly (resp. strongly) admissible set of arguments* $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ *is*

- *a* weakly (resp. strongly) complete *extension iff S contains all the arguments that it weakly (resp. strongly) defends;*
- *a* weakly (resp. strongly) preferred *extension iff it is a $\subseteq$-maximal weakly (resp. strongly) admissible set.*

For $x \in \{w, s\}$ and $\sigma \in \{\mathsf{co}, \mathsf{pr}\}$, the set of x-σ extensions of an IAF $\mathcal{I}$ is denoted $\sigma_x(\mathcal{I})$. In the definition of the versions of complete semantics, the notion of defense used is the same as in the underlying notion of admissibility.

Example 9. We continue Example 8. From the weakly and strongly admissible sets described in Table 3, we deduce $\mathsf{co}_w(\mathcal{I}_2) = \mathsf{pr}_w(\mathcal{I}_2) = \{\{a, c, d, e\}\}$, and $\mathsf{co}_s(\mathcal{I}_2) = \mathsf{pr}_s(\mathcal{I}_2) = \{\{c, d\}\}$.

We observe some usual properties regarding these semantics.

Proposition 3. *Given $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ an IAF, and $x \in \{w, s\}$,*

- $\mathsf{pr}_x(\mathcal{I}) \neq \emptyset$*;*
- $\mathsf{pr}_x(\mathcal{I}) \subseteq \mathsf{co}_x(\mathcal{I})$*.*

Proof. The first item is a direct consequence of the fact that $\mathsf{ad}_x(\mathcal{I}) \neq \emptyset$, as seen previously. The existence of (finitely many) admissible sets implies the existence of $\subseteq$-maximal admissible sets.

Now, let S be a x-preferred extension of $\mathcal{I}$. Reasoning towards a contradiction, let us suppose that $S \notin \mathsf{co}_x(\mathcal{I})$. Since S is x-admissible, it means that S x-defends some argument a that it does not contain. According to Lemma 1, $S \cup \{a\}$ is x-admissible. This means that we have identified a proper superset of S which is x-admissible, thus S is not a $\subseteq$-maximal x-admissible set. This contradicts the fact that S is x-preferred. So we can conclude $S \in \mathsf{co}_x(\mathcal{I})$.

3.3 Stable Semantics for IAFs

Now we focus on a counterpart of stable semantics for IAFs.

Definition 11 (Stable Semantics). *Given $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ an IAF,*

- *a weakly conflict-free set of arguments $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ is a* weakly stable *extension iff $\forall a \in \mathcal{A} \setminus S$, there is some $b \in S \cap \mathcal{A}$ such that $(b, a) \in \mathcal{R}$;*
- *a strongly conflict-free set of arguments $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ is a* strongly stable *extension iff $\forall a \in (\mathcal{A} \cup \mathcal{A}^?) \setminus S$, there is some $b \in S \cap \mathcal{A}$ such that $(b, a) \in \mathcal{R}$.*

Weakly and strongly stable extensions of an IAF $\mathcal{I}$ are denoted by $\mathsf{st}_x(\mathcal{I})$, where $x \in \{w, s\}$.

Example 10. Continuing Example 9, we observe that the weakly preferred extension $S = \{a, c, d, e\}$ is weakly stable as well: the argument $e \in S \cap \mathcal{A}$ (certainly) attacks all the arguments in $\mathcal{A} \setminus S$. It is not strongly stable, since it is not strongly conflict-free.

On the contrary, the strongly preferred extension $S' = \{c, d\}$ is not strongly stable, since it does not attack all the arguments in $(\mathcal{A} \cup \mathcal{A}^?) \setminus S$ (*e.g.* a is not attacked by S').

In Dung's framework, although admissibility is not directly involved in the definition of the stable semantics, any stable extension is actually an admissible set. We show here that it is also the case for strong and weak stable semantics of IAFs.

Proposition 4 (Admissibility of Stable Extensions). *For any IAF* $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$, $\mathsf{st}_x \subseteq \mathsf{ad}_x(\mathcal{I})$, *with* $x \in \{w, s\}$.

Proof. Consider S a weakly stable extension of $\mathcal{I}$. By definition of weakly stable semantics, S is weakly conflict-free. Let us prove that it weakly defends all its elements. Consider any $a \in \mathcal{A} \setminus S$ such that $(a, b) \in \mathcal{R}$ for some $b \in S$. By definition of the weakly stable semantics, there is a $c \in S \cap \mathcal{A}$ such that $(c, a) \in \mathcal{R}$, so S weakly defends b. Thus, S weakly defends all its elements, hence it is weakly admissible.

Now we consider S a strongly stable extension of $\mathcal{I}$. Again, strong conflict-freeness is implied by the definition, so we just need to prove that S strongly defends all its elements. Consider any $a \in \mathcal{A} \cup \mathcal{A}^?$ such that $(a, b) \in \mathcal{R} \cup \mathcal{R}^?$, for some $b \in S$. By definition of strongly stable extensions, there is some $c \in S \cap \mathcal{A}$ such that $(c, a) \in \mathcal{R}$. Thus S strongly defends a, and then all its elements. We can conclude that it is strongly admissible.

Another classical results that still holds for our new semantics is the relationship between (weakly or strongly) stable and (weakly or strongly) preferred extensions.

Proposition 5 (Preferredness of Stable Extensions). *For any IAF* $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$, $\mathsf{st}_x \subseteq \mathsf{pr}_x(\mathcal{I})$, *with* $x \in \{w, s\}$.

Proof. Consider $S \in \mathsf{st}_w(\mathcal{I})$. From Proposition 4, we know that S is weakly admissible. Towards a contradiction, suppose that S is not weakly preferred, *i.e.* $\exists S' \in \mathsf{ad}_w(\mathcal{I})$ such that $S \subset S'$. This implies the existence of an argument $a \in S' \setminus S$. The weak stability of S implies the existence of some $b \in S \cap \mathcal{A}$ such that $(b, a) \in \mathcal{R}$, which violates the weak admissibility of S'. We reach a contradiction, and thus we conclude that $S \in \mathsf{pr}_w(af)$.

Now consider $S \in \mathsf{st}_s(\mathcal{I})$. Again, Proposition 4 implies the strong admissibility of S. Suppose the existence of $S' \in \mathsf{ad}_s(\mathcal{I})$ with $S \subset S'$. Take $a \in S' \setminus S$; the strong stability of S implies the existence of $b \in S \cap \mathcal{A}$ such that $(b, a) \in \mathcal{R}$, thus violating the strong admissibility of S'. We reach a contradiction, and conclude that S' does not exist, hence $S \in \mathsf{pr}_s(\mathcal{I})$.

Example 10 and Proposition 5 imply that $\mathsf{st}_s(\mathcal{I}_2) = \emptyset$. The non-existence of stable extensions in Dung's framework is one of the main differences between this semantics and the ones based on admissibility. We can simply show a similar example for the weakly stable semantics as well: add a certain argument g to $\mathcal{A}$ and $(g, g) \in \mathcal{R}$. The set $\{a, c, d, e\}$ remains the single weakly preferred extension, but it does not attack g, so it is not weakly stable in the new IAF.

4 Computational Issues

4.1 Computational Complexity

In this section, we study the complexity of the variants of verification, credulous acceptability and skeptical acceptability for IAFs. Formally, for $\sigma \in \{\mathsf{cf}, \mathsf{ad}, \mathsf{co}, \mathsf{pr}, \mathsf{st}\}$ and $x \in \{w, s\}$:

σ_x-Ver Given an IAF $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ and $S \subseteq \mathcal{A}$, is S a x-σ-extension of $\mathcal{F}$?
σ_x-Cred Given an IAF $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ and $a \in \mathcal{A} \cup \mathcal{A}^?$, does a belong to some x-σ-extension of $\mathcal{F}$?
σ_x-Skep Given an IAF $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ and $a \in \mathcal{A} \cup \mathcal{A}^?$, does a belong to each x-σ-extension of $\mathcal{F}$?

Lower Bounds. We can prove that reasoning with our semantics for IAFs is (at least) as hard as reasoning with the corresponding semantics for AFs. This can be done by showing that any AF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ can be transformed into an IAF $\mathcal{I}_\mathcal{F}$ that has the same extensions.

Definition 12 (IAF Associated with an AF). *Given $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ an AF, the IAF associated with $\mathcal{F}$ is $\mathcal{I}_\mathcal{F} = \langle \mathcal{A}, \emptyset, \mathcal{R}, \emptyset \rangle$.*

Now we prove the correspondance of extensions, *i.e.* $\sigma(\mathcal{F}) = \sigma_w(\mathcal{I}_\mathcal{F}) = \sigma_s(\mathcal{I}_\mathcal{F})$, for any $\sigma \in \{\mathsf{cf}, \mathsf{ad}, \mathsf{pr}, \mathsf{co}, \mathsf{st}\}$.

Proposition 6 (Dung Compatibility). *Given $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ an AF, $\sigma \in \{\mathsf{cf}, \mathsf{ad}, \mathsf{pr}, \mathsf{co}, \mathsf{st}\}$ and $x \in \{w, s\}$, $\sigma(\mathcal{F}) = \sigma_x(\mathcal{I}_\mathcal{F})$, where $\mathcal{I}_\mathcal{F}$ follows Definition 12.*

Proof. Observe that a set $S \subseteq \mathcal{A}$ is conflict-free (in $\mathcal{F}$) iff it is weakly and strongly conflict-free (in $\mathcal{I}_\mathcal{F}$). Then, a set $S \subseteq \mathcal{A}$ defends an argument $a \in \mathcal{A}$ against all it attackers (in $\mathcal{F}$) iff it weakly and strongly defends a against all its attackers (in $\mathcal{I}_\mathcal{F}$). These facts imply $\mathsf{ad}(\mathcal{F}) = \mathsf{ad}_w(\mathcal{I}_\mathcal{F}) = \mathsf{ad}_s(\mathcal{I}_\mathcal{F})$, which in turn imply the equivalence of complete and preferred extensions of $\mathcal{F}$ with the (weak and strong) complete and preferred extensions of $\mathcal{I}_\mathcal{F}$. Given $S \subseteq \mathcal{A}$, the equivalence between the conditions for S being stable in $\mathcal{F}$ and (weakly or strongly) stable in $\mathcal{I}_\mathcal{F}$ is straightforward.

This allows to prove that the complexity of reasoning with AFs provides a lower bound of the complexity of reasoning with IAFs.

Proposition 7. *Given $\sigma \in \{\mathsf{cf}, \mathsf{ad}, \mathsf{pr}, \mathsf{co}, \mathsf{st}\}$, $x \in \{w, s\}$, and $\mathcal{P} \in \{\mathsf{Ver}, \mathsf{Cred}, \mathsf{Skep}\}$, if σ-$\mathcal{P}$ is $\mathcal{C}$-hard, then σ_x-$\mathcal{P}$ is $\mathcal{C}$-hard.*

Proof. Proposition 6 provides a polynomial-time reduction from σ-$\mathcal{P}$ to σ_x-$\mathcal{P}$.

Upper Bounds for Extension Verification. Similarly to Dung's classical setting, most of the properties of extensions can be verified in polynomial time for our IAF semantics.

Lemma 2. *Given an IAF* $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ *and a set of arguments* $S \subseteq \mathcal{A} \cup \mathcal{A}^?$*, the following tasks are doable in polynomial time:*

1. *check whether* S *is weakly (resp. strongly) conflict-free,*
2. *check whether* S *weakly (resp. strongly) defends some argument* $a \in \mathcal{A}$ *(resp.* $a \in \mathcal{A} \cup \mathcal{A}^?$*),*
3. *check whether each argument in* $\mathcal{A} \backslash S$ *(resp.* $(\mathcal{A} \cup \mathcal{A}^?) \backslash S$*) is attacked by an argument in* $S \cap \mathcal{A}$*.*

Proof. For item 1., weak (resp. strong) conflict-freeness is checked by enumerating every $(a, b) \in S \times S$, and verifying whether $(a, b) \in \mathcal{R}$ (resp. $(a, b) \in \mathcal{R} \cup \mathcal{R}^?$). There are $|S|^2$ such pairs (a, b), and verifying the membership to $\mathcal{R}$ (resp. $\mathcal{R} \cup \mathcal{R}^?$) is bounded by $|\mathcal{A} \cup \mathcal{A}^?|^2$ (*i.e.* the maximal number of possible attacks in an IAF).

For item 2., identifying the arguments $b \in \mathcal{A}$ (resp. $b \in \mathcal{A} \cup \mathcal{A}^?$) such that $(b, a) \in \mathcal{R}$ (resp. $(b, a) \in \mathcal{R} \cup \mathcal{R}^?$) only requires to enumerate all the arguments in $\mathcal{A}$ (resp. $\mathcal{A} \cup \mathcal{A}^?$), and then polynomially check the membership to $\mathcal{R}$ (resp. $\mathcal{R} \cup \mathcal{R}^?$). Then, for each of these attackers b, enumerate all the arguments $c \in S \cap \mathcal{A}$ and check the membership of (c, b) to $\mathcal{R}$ (resp. $\mathcal{R} \cup \mathcal{R}^?$). All the enumerations are polynomially bounded.

Finally, for item 3., enumerate all the pairs (a, b) such that $a \in S \cap \mathcal{A}$ and $b \in \mathcal{A} \backslash S$ (resp. $b \in (\mathcal{A} \cup \mathcal{A}^?) \backslash S$), and then check whether $(a, b) \in \mathcal{R}$.

Combining these polynomial operations allows to check whether a set of arguments is an extension, for most of the semantics studied in this paper.

Proposition 8. *For* $\sigma \in \{\mathsf{cf}, \mathsf{ad}, \mathsf{co}, \mathsf{st}\}$ *and* $x \in \{w, s\}$*,* σ_x-Ver *is polynomial.*

Proof. The result straightforwardly follows Lemma 2.

Following Proposition 7, the verification of (weakly or strongly) preferred extensions is intractable (under the usual assumptions of complexity theory). The following results proves that it remains at the first level of the polynomial hierarchy, similarly to Dung's preferred semantics.

Proposition 9. *For* $x \in \{w, s\}$*,* pr_x-Ver *is in* coNP*.*

Proof. Given $S \subseteq \mathcal{A} \cup \mathcal{A}^?$, proving that S is not a weakly (resp. strongly) preferred extension is doable with the following non-deterministic polynomial algorithm:

1. Check whether S is weakly (resp. strongly) admissible. If not, then S is not weakly (resp. strongly) preferred.
2. Otherwise, guess a proper superset of S, *i.e.* $S \subset S' \subseteq \mathcal{A} \cup \mathcal{A}^?$. Verifying whether S' is a weakly (resp. strongly) admissible set is doable in polynomial time with a deterministic algorithm. If S' is weakly (resp. strongly) admissible, then S is not a weakly (resp. strong) preferred extension.

This algorithm proves that the complementary problem is in NP, thus we conclude that pr_x-$\mathsf{Ver} \in \mathsf{coNP}$ for $x \in \{w, s\}$.

Upper Bounds for Acceptability. First, consider the case of cf_x, for $x \in \{w, s\}$. An argument a is credulously accepted w.r.t. cf_x iff $\{a\} \in \mathsf{cf}_x(\mathcal{I})$. This can be easily checked, by verifying that $(a, a) \notin \mathcal{R}$ and $(a, a) \notin \mathcal{R}^?$. This is doable in polynomial time and logarithmic space. Thus cf_x-$\mathsf{Cred} \in \mathsf{L}$, for $x \in \{w, s\}$. Skeptical acceptability is even easier: since $\emptyset$ is weakly (resp. strongly) conflict-free, there is no skeptically acceptable argument w.r.t. cf_x for any IAF. The reasoning is the same for ad_x-Skep.

Proposition 10. *For* $\sigma \in \{\mathsf{ad}, \mathsf{co}, \mathsf{st}, \mathsf{pr}\}$ *and* $x \in \{w, s\}$*,* σ_x-Cred *is in* NP*.*

Proof. For $\sigma \in \{\mathsf{ad}, \mathsf{co}, \mathsf{st}\}$, guess a set of arguments that contains the queried argument a, and check (in polynomial time, see Proposition 8) whether it is a x-σ-extension. This is a NP algorithm for deciding σ_x-Cred.

For $\sigma = \mathsf{pr}$, notice that an argument belongs to some weakly (resp. strongly) preferred extension iff it belongs to some weakly (resp. strongly) admissible set, hence the result.

Proposition 11. *For* $\sigma \in \{\mathsf{co}, \mathsf{st}\}$ *and* $x \in \{w, s\}$*,* σ_x-Skep *is in* coNP*.*

Proof. Guess a set of arguments that does not contain the queried argument a and check (in polynomial time) whether it is a x-σ-extension, *i.e.* a is not skeptically accepted w.r.t. σ_x. This is a NP algorithm, thus σ_x-Skep is in coNP.

Proposition 12. *For* $x \in \{w, s\}$*,* pr_x-Skep *is in* Π_2^{P}*.*

Proof. Analogous to Proposition 11, except that the higher complexity of verification under the (weakly or strongly) preferred semantics yields a higher complexity upper bound for skeptical acceptability as well.

Discussion. We have proved that, in spite of the higher expressivity of IAFs compared to standard AFs, the complexity of most classical reasoning tasks remains the same. The only exception is skeptical acceptability under (weakly or strongly) complete semantics, for which we only have a coNP upper bound, while it is polynomial in standard Dung's AFs. We plan to study a counterpart of the grounded semantics for IAFs, which could bring new insights for the complete semantics. Finally, notice that using the weak or strong counterpart of our semantics does not have an impact on the complexity of reasoning.

4.2 SAT-Based Computational Approach

We follow the classical approach, initiated by [9], which consists in associating an AF with a propositional formula such that there is a bijection between the extensions of the AF and the models of the formula. Its has been applied with success for developing argumentation solvers [20,24].

In the following, we consider an IAF $\mathcal{I} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$, and we define a set of propositional variables $X_{\mathcal{A} \cup \mathcal{A}^?} = \{x_a \mid a \in \mathcal{A} \cup \mathcal{A}^?\}$. Intuitively, an interpretation ω corresponds to the set of arguments $S = \{a \in \mathcal{A} \cup \mathcal{A}^? \mid \omega(x_a) = \top\}$. We will provide in the rest of this section propositional formulas such that their models correspond to desirable sets of arguments (*e.g.* weakly or strongly conflict-free sets or extensions).

Conflict-Freeness. Recall that a set of arguments is weakly conflict-free if there is no certain attack between two certain arguments in it, while it is strongly conflict-free if there is no attack at all (neither certain nor uncertain) between any element of the set. This is encoded, respectively, by the following formulas ϕ^w_{cf} and ϕ^s_{cf}:

$$\phi^w_{\mathsf{cf}} = \bigwedge_{a,b \in \mathcal{A}, (a,b) \in \mathcal{R}} (\neg x_a \vee \neg x_b)$$

$$\phi^s_{\mathsf{cf}} = \bigwedge_{a,b \in \mathcal{A} \cup \mathcal{A}^?, (a,b) \in \mathcal{R} \cup \mathcal{R}^?} (\neg x_a \vee \neg x_b)$$

Admissibility. Weak (resp. strong) admissibility is based on weak (resp. strong) conflict-freeness, and weak (resp. strong) defense. We introduce a formula δ_w (resp. δ_s) which characterizes sets of arguments that weakly (resp. strongly) defend all their elements.

$$\delta_w = \bigwedge_{a \in \mathcal{A} \cup \mathcal{A}^?} x_a \rightarrow \bigwedge_{b \in \mathcal{A}, (b,a) \in \mathcal{R}} \bigvee_{c \in \mathcal{A}, (c,b) \in \mathcal{R}} x_c$$

$$\delta_s = \bigwedge_{a \in \mathcal{A} \cup \mathcal{A}^?} x_a \rightarrow \bigwedge_{b \in \mathcal{A} \cup \mathcal{A}^?, (b,a) \in \mathcal{R} \cup \mathcal{R}^?} \bigvee_{c \in \mathcal{A}, (c,b) \in \mathcal{R}} x_c$$

Then, weak and strong admissibility are encoded in

$$\phi^x_{\mathsf{ad}} = \phi^x_{\mathsf{cf}} \wedge \delta_x$$

where $x \in \{w, s\}$.

Complete Extensions. The formulas δ_w and δ_s characterize sets of arguments that (weakly or strongly) defend all their elements. To characterize complete extensions, we just need to replace the implication by an equivalence, which yields sets of arguments that defend all their elements and contain everything they defend. Formally,

$$\phi^x_{\mathsf{co}} = \phi^x_{\mathsf{cf}} \wedge \delta'_x$$

where $x \in \{w, s\}$, and

$$\delta'_w = \bigwedge_{a \in \mathcal{A} \cup \mathcal{A}^?} x_a \leftrightarrow \bigwedge_{b \in \mathcal{A}, (b,a) \in \mathcal{R}} \bigvee_{c \in \mathcal{A}, (c,b) \in \mathcal{R}} x_c$$

$$\delta'_s = \bigwedge_{a \in \mathcal{A} \cup \mathcal{A}^?} x_a \leftrightarrow \bigwedge_{b \in \mathcal{A} \cup \mathcal{A}^?, (b,a) \in \mathcal{R} \cup \mathcal{R}^?} \bigvee_{c \in \mathcal{A}, (c,b) \in \mathcal{R}} x_c$$

Stable Extensions. Weakly (resp. strongly) stable extensions are weakly (resp. strongly) conflict-free sets that attack all the certain arguments (resp. all the arguments) that they do not contain. Said otherwise, it means that an argument which is not attacked by (a certain argument in) the extension belongs to the extension. It can be characterized as follows:

$$\phi^w_{\mathsf{st}} = \phi^w_{\mathsf{cf}} \wedge \bigwedge_{a \in \mathcal{A}} ((\bigwedge_{b \in \mathcal{A}, (b,a) \in \mathcal{R}} \neg x_b) \rightarrow x_a)$$

$$\phi^s_{\mathsf{st}} = \phi^s_{\mathsf{cf}} \wedge \bigwedge_{a \in \mathcal{A} \cup \mathcal{A}^?} ((\bigwedge_{b \in \mathcal{A}, (b,a) \in \mathcal{R}} \neg x_b) \rightarrow x_a)$$

Preferred Extensions. Finally, weakly and strongly preferred semantics cannot (under the usual assumptions of complexity theory) be directly encoded as propositional formulas, since the complexity of reasoning with weak and strong preferred semantics is higher than the complexity of Boolean satisfiability (especially, skeptical acceptability is Π_2^{P}-complete). However, other techniques related to propositional logic have been used in the past for computing preferred extensions, *e.g.* quantified Boolean formulas [17], maximal satisfiable subsets [20] or CEGAR (CounterExample Guided Abstraction Refinement) [24]. These techniques could be adapted for computing weakly or strongly preferred extensions.

5 Related Work

Control Argumentation Frameworks (CAFs) [13,23,25] are highly related to IAFs. They add another kind of uncertainty (about the direction of an attack), and a "control part", which is a set of arguments and attacks that must be selected by the agent, the goal being to enforce the acceptability of a set of arguments in each (or some) completion, by means of the selected control arguments. Reasoning with CAFs is only based on completions, and generally the computational complexity is high (at least the same as reasoning with completions of IAFs, and sometimes higher).

Reasoning with weighted AFs (*i.e.* AFs with weights on the attacks) [15] consists, somehow, in relaxing conflict-freeness in order to jointly accept conflicting arguments, as soon as the total amount of conflict (*i.e.* the sum of the weights of the attacks) is lower than a given inconsistency budget. We could adapt this principle for IAFs, by accepting only a given amount of uncertain attacks in extensions.

6 Conclusion

In this paper, we have continued an effort started by [11], and defined extension-based semantics for Incomplete Argumentation Frameworks that do not rely on the completions of the IAF. We have studied the properties of our new semantics, and provided complexity results and logical encoding that pave the way to SAT-based computation.

Future work include, naturally, missing complexity results (*i.e.* tight results for the skeptical acceptability under weakly and strongly complete semantics) and implementation of our encodings. In particular, the comparison of the various methods that reach the second level of the polynomial hierarchy for computing the preferred extensions is an enthralling question. The study of the grounded semantics will fulfill our study of Dung-style semantics for IAFs. We also plan to apply this kind of semantics to Control Argumentation Frameworks, which would decrease the complexity of controllability. This requires to take into account the additional type of information, namely the uncertainty about the direction of attacks. The link with weighted AFs, *i.e.* integrating an inconsistency budget in the weak variants of our semantics, is also a promising line for future research.

References

1. Amgoud, L., Cayrol, C.: A reasoning model based on the production of acceptable arguments. Ann. Math. Artif. Intell. **34**(1–3), 197–215 (2002)
2. Amgoud, L., Cayrol, C., Lagasquie-Schiex, M., Livet, P.: On bipolarity in argumentation frameworks. Int. J. Intell. Syst. **23**(10), 1062–1093 (2008)
3. Arora, S., Barak, B.: Computational Complexity - A Modern Approach. Cambridge University Press, Cambridge (2009)
4. Baroni, P., Caminada, M., Giacomin, M.: Abstract argumentation frameworks and their semantics. In: Baroni, P., Gabbay, D., Giacomin, M., van der Torre, L. (eds.) Handbook of Formal Argumentation, pp. 159–236. College Publications (2018)
5. Baroni, P., Giacomin, M., Liao, B., van der Torre, L.: Encompassing uncertainty in argumentation schemes. In: Proceedings of the Workshop on Frontiers and Connections between Argumentation Theory and Natural Language Processing, Forlì-Cesena, Italy, 21–25 July 2014 (2014)
6. Baumeister, D., Järvisalo, M., Neugebauer, D., Niskanen, A., Rothe, J.: Acceptance in incomplete argumentation frameworks. Artif. Intell. **295**, 103470 (2021)
7. Baumeister, D., Neugebauer, D., Rothe, J.: Verification in attack-incomplete argumentation frameworks. In: Walsh, T. (ed.) ADT 2015. LNCS (LNAI), vol. 9346, pp. 341–358. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-23114-3_21
8. Baumeister, D., Neugebauer, D., Rothe, J., Schadrack, H.: Verification in incomplete argumentation frameworks. Artif. Intell. **264**, 1–26 (2018)
9. Besnard, P., Doutre, S.: Checking the acceptability of a set of arguments. In: 10th International Workshop on Non-Monotonic Reasoning (NMR 2004), pp. 59–64 (2004)
10. Caminada, M., Dunne, P.E.: Strong admissibility revisited: theory and applications. Argument Comput. **10**(3), 277–300 (2019)
11. Cayrol, C., Devred, C., Lagasquie-Schiex, M.C.: Handling ignorance in argumentation: semantics of partial argumentation frameworks. In: Mellouli, K. (ed.) ECSQARU 2007. LNCS (LNAI), vol. 4724, pp. 259–270. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-75256-1_25
12. Coste-Marquis, S., Devred, C., Konieczny, S., Lagasquie-Schiex, M., Marquis, P.: On the merging of Dung's argumentation systems. Artif. Intell. **171**(10–15), 730–753 (2007)

13. Dimopoulos, Y., Mailly, J.G., Moraitis, P.: Control argumentation frameworks. In: 32nd AAAI Conference on Artificial Intelligence (AAAI 2018), pp. 4678–4685 (2018)
14. Dung, P.M.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. Artif. Intell. **77**(2), 321–358 (1995)
15. Dunne, P.E., Hunter, A., McBurney, P., Parsons, S., Wooldridge, M.J.: Weighted argument systems: basic definitions, algorithms, and complexity results. Artif. Intell. **175**(2), 457–486 (2011)
16. Dvorák, W., Dunne, P.E.: Computational problems in formal argumentation and their complexity. In: Baroni, P., Gabbay, D., Giacomin, M., van der Torre, L. (eds.) Handbook of Formal Argumentation, pp. 631–688. College Publications (2018)
17. Egly, U., Woltran, S.: Reasoning in argumentation frameworks using quantified boolean formulas. In: Computational Models of Argument: Proceedings of COMMA 2006, 11–12 September 2006, Liverpool, UK, vol. 144, pp. 133–144 (2006)
18. Fazzinga, B., Flesca, S., Furfaro, F.: Revisiting the notion of extension over incomplete abstract argumentation frameworks. In: 29th International Joint Conference on Artificial Intelligence (IJCAI 2020), pp. 1712–1718 (2020)
19. Hunter, A.: Some foundations for probabilistic abstract argumentation. In: Verheij, B., Szeider, S., Woltran, S. (eds.) Computational Models of Argument - Proceedings of COMMA 2012, Vienna, Austria, 10–12 September 2012. Frontiers in Artificial Intelligence and Applications, vol. 245, pp. 117–128. IOS Press (2012)
20. Lagniez, J.M., Lonca, E., Mailly, J.G.: CoQuiAAS: a constraint-based quick abstract argumentation solver. In: 27th IEEE International Conference on Tools with Artificial Intelligence, ICTAI 2015, Vietri sul Mare, Italy, 9–11 November 2015, pp. 928–935. IEEE Computer Society (2015)
21. Lawrence, J., Reed, C.: Argument mining: a survey. Comput. Linguist. **45**(4), 765–818 (2019)
22. Li, H., Oren, N., Norman, T.J.: Probabilistic argumentation frameworks. In: Proceedings of the First International Workshop on Theory and Applications of Formal Argumentation (TAFA 2011), pp. 1–16 (2011)
23. Mailly, J.G.: Possible controllability of control argumentation frameworks. In: 8th International Conference on Computational Models of Argument (COMMA 2020), pp. 283–294 (2020)
24. Niskanen, A., Järvisalo, M.: μ-toksia: an efficient abstract argumentation reasoner. In: Calvanese, D., Erdem, E., Thielscher, M. (eds.) Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR 2020, Rhodes, Greece, 12–18 September 2020, pp. 800–804 (2020)
25. Niskanen, A., Neugebauer, D., Järvisalo, M.: Controllability of control argumentation frameworks. In: Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020, pp. 1855–1861 (2020)
26. Rossit, J., Mailly, J.G., Dimopoulos, Y., Moraitis, P.: United we stand: accruals in strength-based argumentation. Argument Comput. **12**(1), 87–113 (2021)