Revisiting initial sets in abstract argumentation

Matthias Thimm

Artificial Intelligence Group, University of Hagen, Germany

Abstract. We revisit the notion of *initial sets* by Xu and Cayrol (In *Proceedings of the 1st Chinese Conference on Logic and Argumentation (CLAR'16)* 2016), i. e., non-empty minimal admissible sets in abstract argumentation frameworks. Initial sets are a simple concept for analysing conflicts in an abstract argumentation framework and to explain why certain arguments can be accepted. We contribute with new insights on the structure of initial sets and devise a simple non-deterministic construction principle for any admissible set, based on iterative selection of initial sets of the original framework and its induced reducts. In particular, we characterise many existing admissibility-based semantics via this construction principle, thus providing a constructive explanation on the structure of extensions. We also investigate certain problems related to initial sets with respect to their computational complexity.

Keywords: Abstract argumentation, initial sets

1. Introduction

Formal argumentation [3,6] encompasses approaches for non-monotonic reasoning that focus on the role of arguments and their interactions. The most well-known approach is that of *abstract argumentation frameworks* [19] that model arguments as vertices in a directed graph, where a directed edge from an argument *a* to an argument *b* denotes an *attack* from *a* to *b*. Although conceptually simple, abstract argumentation frameworks can be used in a variety of argumentative scenarios such as persuasion dialogues [16], explanations in recommendation systems [31], mathematical modelling [11], or as a target formalism of *structured argumentation formalisms* [28,35]. Formal semantics are given to abstract argumentation frameworks by *extensions*, i. e., sets of arguments that can jointly be accepted and represent a coherent standpoint on the conflicts between the arguments. Different variants of such semantics have been proposed [5], but there are also other (non set-based) approaches such as ranking-based semantics [1] and probabilistic approaches [26].

A particular advantage of approaches to formal argumentation is the capability to *explain* the reasoning behind certain conclusions using human-accessible concepts such as arguments and counterarguments. Works such as [2,27,29,31–33,36] have already explored certain aspects of the explanatory power of approaches to formal argumentation. Amgoud and Prakken [2] and Rago et. al [31] develop argumentation formalisms for decision-making that augment recommendations with arguments. In [33] an extension to abstract argumentation frameworks is developed that explicitly includes a relationship for an "explanation", while Liao and van der Torre [27] define "explanation semantics" for ordinary abstract argumentation frameworks. Saribatur, Wallner, and Woltran [32] as well as Niskanen and Järvisalo [29] address computational problems and develop a notion for explaining non-acceptability of arguments to,

e. g., verify results of an argumentation solver. Finally, [36] presents *strong explanations* as a mechanism to explain acceptability of (sets of) arguments.

In this paper, we revisit one of the fundamental concepts underlying approaches to formal argumentation (and abstract argumentation in particular) for the purpose of explaining, namely *admissibility*. Informally speaking, a set of arguments is *admissible* if each of its members is defended against any attack from the outside (we will provide formal details in Section 2). Many popular semantics for abstract argumentation rely on the notion of admissibility. In particular, a *preferred extension* is a maximal (wrt. set inclusion) admissible set and preferred semantics satisfies many desirable properties [5]. However, since a preferred extension is a maximal admissible set, it can hardly be used for explaining why a certain argument is acceptable: such an extension may contain many irrelevant arguments and its size alone distracts from the particular reasons why a certain member is acceptable. Our aim is to investigate why certain arguments are contained in, e. g., a preferred extension and how we can decompose such large extensions into smaller sets that allow us to justify the reasoning process behind such complex semantics. As a tool for our investigation, we consider *initial sets*, i. e., non-empty admissible sets that are minimal wrt. set inclusion. Initial sets have been introduced in [38] and further analysed in [39,40]. We contribute to this analysis with new insights on the structure of initial sets and, in particular, to the use of initial sets for the task of explanation. In fact, initial sets can exactly be used for the purpose mentioned before [38]: they allow us to decompose large extensions into smaller fragments, each of them representing a single resolved issue in the argumentation framework and thus showcases the reasoning behind certain semantics. This has been done already in some form in [38–40] but we present a new, and arguably more elegant, formalisation of that idea that allows us to derive new results as well. Using the notion of a *reduct* [8], we can concisely represent any admissible set as a sequence of initial sets of the original framework and derived reducts. Moreover, we characterise many admissibility-based semantics through a step-wise construction process using certain selections of initial sets (this has been hinted at using the original formalisation for complete and preferred semantics in [40]). We round up our analysis with a characterisation of the computational complexity of certain tasks related to initial sets, which is also missing so far from the literature.

In summary, the contributions of this paper are as follows:

(1) We revisit initial sets and investigate further properties (Section 3)
(2) We provide a characterisation result of admissible sets and many admissibility-based semantics (Section 4)
(3) We analyse certain computational problems wrt. their complexity (Section 5)

Section 2 introduces preliminaries on abstract argumentation and Section 6 concludes this paper.

Complete proofs can be found in the appendix. A short paper presenting the main ideas of this work has been published before in [34].

2. Abstract argumentation

Let $\mathfrak{A}$ denote a universal set of arguments. An *abstract argumentation framework* AF is a tuple AF = (A, R) where A $\subseteq$ $\mathfrak{A}$ is a finite set of arguments and R is a relation R $\subseteq$ A $\times$ A [19]. Let $\mathfrak{AF}$ denote the set of all abstract argumentation frameworks. For two arguments $a, b \in$ A the relation aRb means that argument a attacks argument b. For AF = (A, R) and AF' = (A', R') we write AF' $\sqsubseteq$ AF iff A' $\subseteq$ A and

$R' = R \cap (A' \times A')$. For a set $X \subseteq A$, we denote by $AF|_X = (X, R \cap (X \times X))$ the projection of AF on X. For a set $S \subseteq A$ we define

$$S^+ = \{a \in A \mid \exists b \in S : bRa\}$$

$$S^- = \{a \in A \mid \exists b \in S : aRb\}$$

If S is a singleton set, we omit brackets for readability, i. e., we write a^- (a^+) instead of $\{a\}^-$ ($\{a\}^+$). For two sets S and S' we write SRS' iff $S^+ \cap S' \neq \emptyset$. We say that a set $S \subseteq A$ is *conflict-free* if for all $a, b \in S$ it is not the case that aRb. A set S *defends* an argument $b \in A$ if for all a with aRb there is $c \in S$ with cRa. A conflict-free set S is called *admissible* if S defends all $a \in S$. Let $adm(AF)$ denote the set of admissible sets of AF.

Different semantics can be phrased by imposing constraints on admissible sets [5]. In particular, an admissible set E

- is a *complete* (co) extension iff for all $a \in A$, if E defends a then $a \in E$,
- is a *grounded* (gr) extension iff E is complete and minimal,
- is a *stable* (st) extension iff $E \cup E^+ = A$,
- is a *preferred* (pr) extension iff E is maximal.
- is a *semi-stable* (sst) extension iff $E \cup E^+$ is maximal,
- is an *ideal* (id) extension iff E is the maximal admissible set with $E \subseteq E'$ for each preferred extension E'.
- is a *strongly admissible* (sa) extension iff $E = \emptyset$ or each $a \in E$ is defended by some strongly admissible $E' \subseteq E \setminus \{a\}$.

All statements on minimality/maximality are meant to be with respect to set inclusion. For $\sigma \in \{co, gr, st, pr, sst, id, sa\}$ let $\sigma(AF)$ denote the set of σ-extensions of AF. We say a semantics σ is *admissibility-based* if $\sigma(AF) \subseteq adm(AF)$ for all AF. Note that all semantics above are admissibility-based but there are also others such as CF2 semantics [7] and weak admissibility-based semantics [8].

3. Revisiting initial sets

Admissibility captures the basic intuition for an explanation *why* a certain argument can be regarded as acceptable. More concretely, if S is an admissible set then $a \in S$ is accepted *because* all arguments in S are accepted, every attacker of a is attacked back by some argument in S. However, admissibility alone is not sufficient to model explainability as it does not take *relevance* into account.

Example 1. Consider the argumentation framework AF_0 depicted in Fig. 1. There are eight admissible sets containing the argument e:

$$S_1 = \{b, e, f, h, i\} \qquad S_2 = \{b, e, f, i\} \qquad S_3 = \{b, e, h, i\}$$

$$S_4 = \{e, f, h, i\} \qquad S_5 = \{b, e, i\} \qquad S_6 = \{f, e, i\}$$

$$S_7 = \{h, e, i\} \qquad S_8 = \{e, i\}$$

S_1 is also a preferred extension. However, it is also clear that arguments b, f, and h are not integral for defending e and the set S_8 presents a concise description of what is needed in order to deem e as acceptable (wrt. admissibility), namely only e and i.
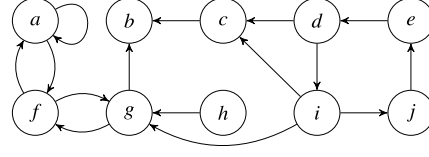
Fig. 1. AF_0 from example 1.

In the following, we take relevance into account by considering *minimal* (wrt. set inclusion) admissible sets. Of course, a notion of minimal admissible set without further constraints is not a useful concept as the empty set is always admissible and constitutes the unique minimal admissible set. Non-empty minimal admissible sets have been coined *initial sets* by Xu and Cayrol in [38].

Definition 1 ([38])**.** For $AF = (A, R)$, a set $S \subseteq A$ with $S \neq \emptyset$ is called an *initial set* if S is admissible and there is no admissible $S' \subsetneq S$ with $S' \neq \emptyset$. Let $IS(AF)$ denote the set of initial sets of AF.

Example 2. We continue Example 1. There are four initial sets of AF_0: $\{f\}$, $\{h\}$, $\{d, j\}$, and $\{e, i\}$.

As the previous example shows, an initial set is not supposed to provide a "solution" to the whole argumentation represented in an abstract argumentation framework, but "solves" a single atomic conflict (or in the case of $\{h\}$ points to an obvious deterministic inference step). In fact, we can identify three different types of initial sets.

Definition 2. For $AF = (A, R)$ and $S \in IS(AF)$, we say that

(1) S is *unattacked* iff $S^- = \emptyset$,
(2) S is *unchallenged* iff $S^- \neq \emptyset$ and there is no $S' \in IS(AF)$ with $S'RS$,
(3) S is *challenged* iff there is $S' \in IS(AF)$ with $S'RS$.

Note that only unattacked initial sets have been considered explicitly in [40]; in particular, note that every unattacked initial set S is a singleton $S = \{a\}$. Observe that the notions of unattacked, unchallenged, and challenged initial sets are mutually exclusive and exhaustive. Let $IS^{\leftarrow}(AF)$, $IS^{\leftrightarrow}(AF)$, and $IS^{\leftrightarrow}(AF)$ denote the set of unattacked, unchallenged, and challenged initial sets, respectively. So we have $IS(AF) = IS^{\leftarrow}(AF) \cup IS^{\leftrightarrow}(AF) \cup IS^{\leftrightarrow}(AF)$. Moreover, for $S \in IS(AF)$ let

$$\text{conflicts}(S, AF) = \left\{ S' \in IS(AF) \mid S'RS \right\}$$

denote the set of conflicting initial sets of S, which is always empty in the case of unattacked and unchallenged initial sets. Note that SRS' implies $S'RS$ for any $S, S' \in IS(AF)$ as S' is admissible and therefore defends itself.

Example 3. We continue Example 2. Here we have

$$IS^{\leftarrow}(AF) = \left\{\{h\}\right\}$$
$$IS^{\leftrightarrow}(AF) = \left\{\{f\}\right\}$$
$$IS^{\leftrightarrow}(AF) = \left\{\{d, j\}, \{e, i\}\right\}$$

and $\{d, j\}$ and $\{e, i\}$ are in conflict with each other, i. e., $\text{conflicts}(\{d, j\}, AF) = \{\{e, i\}\}$ and $\text{conflicts}(\{e, i\}, AF) = \{\{d, j\}\}$.

Before we continue with characterising arbitrary admissible sets using initial sets in Section 4, we first contribute some new results on the structure of initial sets, therefore extending the analysis from [38–40].

Initial sets have an interesting property with respect to strongly connected components as follows. Recall that we can decompose an abstract argumentation framework AF into its strongly connected components. More precisely, an abstract argumentation framework $\mathsf{AF}' = (\mathsf{A}', \mathsf{R}')$ is a strongly connected component (SCC) of AF, if $\mathsf{AF}' \sqsubseteq \mathsf{AF}$ s.t. there is a directed path between any pair $a, b \in \mathsf{A}'$ in AF' and there is no larger AF'' with that property. Let $\mathsf{SCC}(\mathsf{AF})$ be the set of SCCs of AF.

Example 4. Consider again the argumentation framework AF_0 from Fig. 1. AF_0 decomposes as follows into SCCs: $\mathsf{SCC}(\mathsf{AF}_0) = \{\mathsf{AF}_0|_{\{d,e,i,j\}}, \mathsf{AF}_0|_{\{c\}}, \mathsf{AF}_0|_{\{h\}}, \mathsf{AF}_0|_{\{g,f,a\}}, \mathsf{AF}_0|_{\{b\}}\}$

The following result shows that initial sets are always completely contained in a single SCC.

Proposition 1. *If S is an initial set of AF then there is $\mathsf{AF}' = (\mathsf{A}', \mathsf{R}') \in SCC(\mathsf{AF})$ s.t. $S \subseteq \mathsf{A}'$.*

If S is an initial set let $\mathsf{SCC}(S)$ denote its SCC as in the above proposition. Initial sets can actually be characterised by their SCC as follows.

Proposition 2. *S is an initial set of AF if and only if S is an initial set of $SCC(S) = (\mathsf{A}', \mathsf{R}')$ and $S^- \subseteq \mathsf{A}'$.*

In other words, a set S is an initial set iff it is an initial set of a single SCC and it is not attacked by arguments outside of the SCC.

We close this investigation on the relationship of initial sets with SCCs by making some straightforward observations regarding the types of initial sets.

**Proposition 3.** *Let $S \in IS(\mathsf{AF})$ and $SCC(S) = (\mathsf{A}', \mathsf{R}')$.*

(1) *If S is unattacked then $|\mathsf{A}'| = 1$.*
(2) *If S is challenged or unchallenged then $|\mathsf{A}'| > 1$.*
(3) *If S is challenged and $S' \in conflicts(S, \mathsf{AF})$ then $SCC(S) = SCC(S')$.*

In particular, the final observation in the previous proposition states that conflicts between initial sets are always within a single SCC.

4. Characterising admissibility-based semantics through initial sets

In [38,40] it has been shown that any admissible set (and in particular every complete and preferred extension) can be constructed by (1) selecting a set of non-conflicting initial sets, (2) adding further defended arguments, and (3) iterating this procedure taking so-called "J-acceptable" sets into account. In particular, the described mechanism involves iterative application of the characteristic function [19], computation of the grounded extension, and said notion of J-acceptability to provide those characterisations (and some further concepts). In this section, we provide a (arguably) more elegant formalisation of these ideas that allows us to derive characterisations of further semantics as well as some impossibility results. Results that are (partly) due to works [38–40] are annotated as such, all remaining results are new.

Our characterisations rely on the notion of the *reduct* [8].
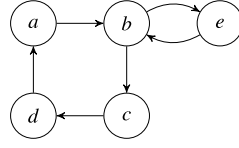
Fig. 2. AF_1 from Example 5.

Definition 3. For $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ and $S \subseteq \mathsf{A}$, the S-reduct AF^S of AF is defined via $\mathsf{AF}^S = \mathsf{AF}|_{\mathsf{A} \setminus (S \cup S^+)}$.

As a single initial set S solves an atomic conflict in an abstract argumentation framework AF, "committing" to it by moving to AF^S may reveal further conflicts and thus new initial sets. We can make the following observations on this aspect.

Proposition 4. *Let* $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ *be an abstract argumentation framework and* $S, S' \in IS(\mathsf{AF})$ *with* $S \neq S'$.

(1) *If* $S' \in IS^{\leftarrow}(\mathsf{AF})$ *then* $S' \in IS^{\leftarrow}(\mathsf{AF}^S)$
(2) *If* $S' \in conflicts(S, \mathsf{AF})$ *then* $S' \notin IS(\mathsf{AF}^S)$
(3) *If* $S' \notin conflicts(S, \mathsf{AF})$ *then* $S' \cap \bigcup IS(\mathsf{AF}^S) \neq \emptyset$

The above observations give an impression on how initial sets behave under reducts. So unattacked initial sets are always retained (item 1), conflicting initial sets are always removed (item 2), and non-conflicting initial sets are "essentially" retained (item 3). More precisely, while it may not be guaranteed that non-conflicting initial sets are still initial sets in the reduct, their arguments are still potentially acceptable, as the following example shows.

Example 5. Consider the argumentation framework AF_1 depicted in Fig. 2. We have

$$IS(\mathsf{AF}_1) = \big\{\{a, c\}, \{b, d\}, \{e\}\big\}$$

and $\{b, d\}$ and $\{e\}$ are conflicting (there are no further conflicts). Now we have

$$IS\big(\mathsf{AF}_1^{\{e\}}\big) = \big\{\{c\}\big\}$$

So the initial set $\{a, c\}$ of AF_1 is not retained in $\mathsf{AF}_1^{\{e\}}$, despite $\{a, c\}$ and $\{e\}$ not being in conflict. However, we have that $\{c\} \subseteq \{a, c\}$ is an initial set of $\mathsf{AF}_1^{\{e\}}$. Furthermore, $\{a\}$ (the "remaining" argument of the initial set $\{a, c\}$) is actually the unique initial set of $(\mathsf{AF}_1^{\{e\}})^{\{c\}}$.

The following results show that by an iterative selection of initial sets on the corresponding reducts, we can re-construct every admissible set. These observations have been hinted at in [38–40], but no formal proof had been provided. We make up for that now.

Theorem 1. *Let* $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ *be an abstract argumentation framework and* $S \subseteq \mathsf{A}$. S *is admissible if and only if either*

- $S = \emptyset$ *or*
- $S = S_1 \cup S_2$, $S_1 \in IS(\mathsf{AF})$ *and* S_2 *is admissible in* AF^{S_1}.

Proof. Let S be an admissible set and $S \neq \emptyset$. By definition of initial sets, there is $S_1 \in \mathsf{IS}(\mathsf{AF})$ with $S_1 \subseteq S$. It remains to show that $S_2 = S \setminus S_1$ is admissible in AF^{S_1}. Let $a \in S_2$ and let $b_1, \ldots, b_n \in \mathsf{A}$ be the attackers of a in AF. Since S is admissible, there are arguments $c_1, \ldots, c_n \in S$ so that c_i attacks b_i, $i = 1, \ldots, n$ (possibly some of the c_i are identical). Without loss of generality, assume $c_1, \ldots, c_k \in S_1$ for some $k \leqslant n$. Then $b_1, \ldots, b_k$ are not present in AF^{S_1}, thus a must only be defended against $b_{k+1}, \ldots, b_n$ in AF^{S_1}. However, since $S_2 = S \setminus S_1$ we have that $c_{k+1}, \ldots, c_n \in S_2$, showing that a is defended by S_2 in AF^{S_1} and, thus, S_2 is admissible in AF^{S_1}.

For the other direction,[1] if $S = \emptyset$ then S is also admissible. Assume $S = S_1 \cup S_2$, $S_1 \in \mathsf{IS}(\mathsf{AF})$ and S_2 is an admissible set of AF^{S_1}. We have to show that S is admissible. Let $a \in S$ and let $b_1, \ldots, b_n$ be the attackers of a in AF. If $a \in S_1$ then there are $c_1, \ldots, c_n \in S_1 \subseteq S$ such that c_i attacks b_i since S_1 is admissible. If $a \in S_2$, assume for the sake of contradiction that there is an attacker b of a such that there is no $b \in S$ that attacks c in AF. It follows that b is also in AF^{S_1} and a is undefended by S_2 in AF^{S_1}. This contradicts the assumption that S_2 is admissible in AF^{S_1}. $\square$

By recursively applying the above theorem, we obtain the following corollary.

Corollary 1. *Every non-empty admissible set S can be written as $S = S_1 \cup \cdots \cup S_n$ with pairwise disjoint S_i, $i = 1, \ldots, n$, S_1 is an initial set of AF and every S_i, $i = 2, \ldots, n$ is an initial set of $\mathsf{AF}^{S_1 \cup \cdots \cup S_{i-1}}$. Furthermore, only non-empty admissible sets S can be written in such a fashion.*

Example 6. Consider again the argumentation framework AF_0 from Fig. 1 and recall that

$$S_1 = \{b, e, f, h, i\}$$

is an admissible set of AF_0 (and actually a preferred extension). S_1 can be written as

$$S_1 = \{h\} \cup \{f\} \cup \{e, i\} \cup \{b\}$$

where

- $\{h\}$ is an initial set of AF_0 (Fig. 1),
- $\{f\}$ is an initial set of $\mathsf{AF}_0^{\{h\}}$ (Fig. 3(a)),
- $\{e, i\}$ is an initial set of $\mathsf{AF}_0^{\{h, f\}}$ (Fig. 3(b)), and
- $\{b\}$ is an initial set of $\mathsf{AF}_0^{\{e, i, h, f\}}$ (Fig. 3(c)).

Note that a decomposition $S = S_1 \cup \cdots \cup S_n$ of an admissible set S from Corollary 1 is not necessarily uniquely determined. In the previous example, S_1 could also have been constructed by selecting, e. g., $\{f\}$ first.

Let us now discuss the wider significance of Theorem 1 and Corollary 1. For that, recall the standard approach to compute and justify the (uniquely determined) grounded extension of an argumentation framework $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ [19], cf. also the discussion in [38]. Basically, the grounded extension E of AF can be computed by selecting any non-attacked argument $a \in \mathsf{A}$, add it to E, remove a and all arguments attacked by a from AF (so move from AF to $\mathsf{AF}^{\{a\}}$), and continue the process until no further

---

[1]Note that this direction can also be shown by using the fact that admissible "semantics" is fully decomposable [4], but we explicitly prove it for matters of simplicity.
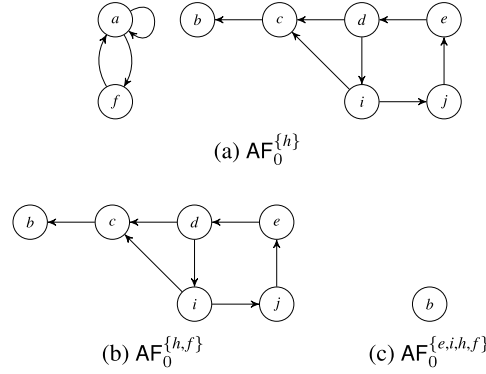
(a) $\mathsf{AF}_0^{\{h\}}$

(b) $\mathsf{AF}_0^{\{h,f\}}$ (c) $\mathsf{AF}_0^{\{e,i,h,f\}}$

Fig. 3. Reducts obtained from AF_0 for the construction of $S_1 = \{b, e, f, h, i\}$.

unattacked argument can be found. Observe the similarity of this procedure to the procedure indicated by Theorem 1: in order to construct any admissible set S of AF, first select any initial set S' of AF, add it to S (which is initially empty), remove S' and all arguments attacked by S' from AF, and continue the process. Therefore, initial sets allow us to *serialise* the construction of any admissible set into smaller steps, each of these steps solving a single conflict in the framework under consideration. Depending on how initial sets are selected at each step and how we end the process, we can also recover different semantics. Let us now formalise these ideas. For that, we first need two functions that define the selection mechanism of initial sets and the termination criterion.

Definition 4. A *state* T is a tuple $T = (\mathsf{AF}, S)$ with $\mathsf{AF} \in \mathfrak{AF}$ and $S \subseteq \mathfrak{A}$.

Definition 5. A *selection function* α is any function $\alpha : 2^{\mathfrak{A}} \times 2^{\mathfrak{A}} \times 2^{\mathfrak{A}} \to 2^{\mathfrak{A}}$ with $\alpha(X, Y, Z) \subseteq X \cup Y \cup Z$ for all $X, Y, Z \subseteq \mathfrak{A}$.

We will apply a selection function α in the form $\alpha(\mathsf{IS}^{\leftarrow}(\mathsf{AF}), \mathsf{IS}^{\leftrightarrow}(\mathsf{AF}), \mathsf{IS}^{\leftrightarrow}(\mathsf{AF}))$ (for some AF), so α selects a subset of the initial sets as eligible to be selected in the construction process. We explicitly differentiate the different types of initial sets as parameters here as a technical convenience.

Definition 6. A *termination function* β is any function $\beta : \mathfrak{AF} \times 2^{\mathfrak{A}} \to \{0, 1\}$.

A termination function β is used to indicate when a construction of an admissible set is finished (this will be the case if $\beta(\mathsf{AF}, S) = 1$).

We will now define a transition system on states that makes use of a selection and a termination function to constrain the construction of admissible sets. For some selection function α, consider the following transition rule:

$$(\mathsf{AF}, S) \xrightarrow{S' \in \alpha\left(\mathsf{IS}^{\leftarrow}(\mathsf{AF}), \mathsf{IS}^{\leftrightarrow}(\mathsf{AF}), \mathsf{IS}^{\leftrightarrow}(\mathsf{AF})\right)} \left(\mathsf{AF}^{S'}, S \cup S'\right) \tag{1}$$

If (AF', S') can be reached from (AF, S) via a finite number of steps (this includes no steps at all) with the above rule we write $(\mathsf{AF}, S) \rightsquigarrow^{\alpha} (\mathsf{AF}', S')$. If, in addition, the state (AF', S') also satisfies the termination criterion of β, i. e., $\beta(\mathsf{AF}', S) = 1$, then we write $(\mathsf{AF}, S) \rightsquigarrow^{\alpha, \beta} (\mathsf{AF}', S')$.

Given concrete instances of α and β, let $\mathcal{E}^{\alpha, \beta}(\mathsf{AF})$ be the set of all S with $(\mathsf{AF}, \emptyset) \rightsquigarrow^{\alpha, \beta} (\mathsf{AF}', S)$ (for some AF').

Definition 7. A semantics σ is *serialisable* with a selection function α and a termination function β if $\sigma(\mathsf{AF}) = \mathcal{E}^{\alpha,\beta}(\mathsf{AF})$ for all AF.

A direct consequence of Corollary 1 is the following.

Theorem 2. *Admissible semantics[2] is serialisable with*

$$\alpha_{adm}(X, Y, Z) = X \cup Y \cup Z \qquad \beta_{adm}(\mathsf{AF}, S) = 1$$

In other words, any admissible set can be constructed by not constraining the selection of initial sets at all (α_{adm}) and accepting every reachable state (β_{adm}). We can also characterise most of the admissibility-based semantics from Section 2 through serialisation using specific selection and termination functions, as the following results show.

The following observation has been hinted at in [40], but not formally proven. Using the notions of selection and termination functions, we can make the construction principle of complete extensions explicit.

Theorem 3. *Complete semantics is serialisable with* α_{adm} *and*

$$\beta_{co}(\mathsf{AF}, S) = \begin{cases} 1 & \text{if } IS^{\leftharpoondown}(\mathsf{AF}) = \emptyset \\ 0 & \text{otherwise} \end{cases}$$

Note that β_{co} only accepts those states, which cannot be extended by already defended arguments (which are those appearing in $\mathsf{IS}^{\leftharpoondown}(\mathsf{AF})$). The above theorem also allows us to characterise complete extensions in a similar manner as admissible sets in Theorem 1.

Corollary 2. *Let* $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ *be an abstract argumentation framework and* $S \subseteq \mathsf{A}$. *S is complete if and only if either*

- $S = \emptyset$ *and* $IS^{\leftharpoondown}(\mathsf{AF}) = \emptyset$ *or*
- $S = S_1 \cup S_2$, $S_1 \in IS(\mathsf{AF})$ *and* S_2 *is complete in* AF^{S_1}.

Grounded semantics has the same termination criterion as complete semantics, but constrains the selection of initial sets to those in $\mathsf{IS}^{\leftharpoondown}(\mathsf{AF})$.

Theorem 4. *Grounded semantics is serialisable with*

$$\alpha_{gr}(X, Y, Z) = X$$

and β_{co}.

Note that α_{gr} and β_{co} formalise the algorithm to compute the grounded extension sketched before. Therefore, the non-deterministic algorithm realised by the transition rule (1) is a generalisation of this algorithm. Similarly as Corollary 2 we obtain the following characterisation of grounded semantics in terms of the reduct.

---

[2]Although admissible sets are usually not regarded as a semantics, we can treat the function $\mathsf{adm}(\cdot)$ as such.

Corollary 3. *Let* $AF = (A, R)$ *be an abstract argumentation framework and* $S \subseteq A$. *S is grounded if and only if either*

- $S = \emptyset$ *and* $IS^{\leftharpoonup}(AF) = \emptyset$ *or*
- $S = S_1 \cup S_2$, $S_1 \in IS^{\leftharpoonup}(AF)$ *and* S_2 *is grounded in* AF^{S_1}.

For stable semantics, we do not need to constrain the selection of initial sets but only ensure that all arguments are either included in or attacked by the constructed extension.

Theorem 5. *Stable semantics is serialisable with* α_{adm} *and*

$$\beta_{st}(AF, S) = \begin{cases} 1 & \text{if } AF = (\emptyset, \emptyset) \\ 0 & \text{otherwise} \end{cases}$$

As before, we obtain the following characterisation of stable semantics in terms of the reduct.

Corollary 4. *Let* $AF = (A, R)$ *be an abstract argumentation framework and* $S \subseteq A$. *S is stable if and only if either*

- $S = \emptyset$ *and* $A = \emptyset$ *or*
- $S = S_1 \cup S_2$, $S_1 \in IS(AF)$ *and* S_2 *is stable in* AF^{S_1}.

For preferred semantics, we simply have to apply transitions exhaustively. Note that this result strengthens Proposition 3 from [40].

Theorem 6. *Preferred semantics is serialisable with* α_{adm} *and*

$$\beta_{pr}(AF, S) = \begin{cases} 1 & \text{if } IS(AF) = \emptyset \\ 0 & \text{otherwise} \end{cases}$$

Corollary 5. *Let* $AF = (A, R)$ *be an abstract argumentation framework and* $S \subseteq A$. *S is preferred if and only if either*

- $S = \emptyset$ *and* $IS(AF) = \emptyset$ *or*
- $S = S_1 \cup S_2$, $S_1 \in IS(AF)$ *and* S_2 *is preferred in* AF^{S_1}.

Our final positive result is about strong admissibility, which follows quite easily from the construction of the grounded extension.

Theorem 7. *Strong admissibility semantics is serialisable with* α_{gr} *and* β_{adm}.

Corollary 6. *Let* $AF = (A, R)$ *be an abstract argumentation framework and* $S \subseteq A$. *S is strongly admissible if and only if either*

- $S = \emptyset$ *or*
- $S = S_1 \cup S_2$, $S_1 \in IS^{\leftharpoonup}(AF)$ *and* S_2 *is strongly admissible in* AF^{S_1}.

A related result to the above observation is given by Baumann et. al in [9] (Definition 7 and Proposition 2). There, a strongly admissible extension E is characterised through the existence of pairwise disjoint sets $A_1, \ldots, A_n$ such that $E = A_1 \cup \ldots A_n$, A_1 is a set of unattacked arguments in AF and $A_1 \cup \cdots \cup A_j$ defends A_{j+1} for all $1 \leqslant j \leqslant n-1$. Our construction above also implies the existence of these sets $A_1, \ldots, A_n$ with the same properties, but with the additional feature that all $A_i, i = 1, \ldots, n$ are singleton sets (since unattacked initial sets are always singleton sets).

Missing from our results so far are the ideal and semi-stable semantics. Both of them cannot be serialised.

Theorem 8. *Ideal semantics is not serialisable.*

The proof of the above theorem is given by the following counterexample.

Example 7. Consider the two argumentation frameworks AF_2 and AF_3 in Fig. 4. We have

$$\mathsf{id}(\mathsf{AF}_2) = \{b\} \qquad \mathsf{id}(\mathsf{AF}_3) = \{b, e\}$$

and

$$\mathsf{IS}^{\leftarrow}(\mathsf{AF}_2) = \mathsf{IS}^{\leftarrow}(\mathsf{AF}_3) = \emptyset$$
$$\mathsf{IS}^{\leftrightarrow}(\mathsf{AF}_2) = \mathsf{IS}^{\leftrightarrow}(\mathsf{AF}_3) = \big\{\{b\}, \{e\}\big\}$$
$$\mathsf{IS}^{\leftrightarrow}(\mathsf{AF}_2) = \mathsf{IS}^{\leftrightarrow}(\mathsf{AF}_3) = \emptyset$$

So no selection function α can distinguish these frameworks and should return either

(1) $\alpha(\emptyset, \{\{b\}, \{e\}\}, \emptyset) = \emptyset$ or
(2) $\alpha(\emptyset, \{\{b\}, \{e\}\}, \emptyset) = \{\{b\}, \{e\}\}$.

Note also that α cannot return just one of the two initial sets as they cannot be distinguished. In case 1, the ideal extensions of both AF_2 and AF_3 cannot be constructed. So assume case 2 and select $\{e\}$ in
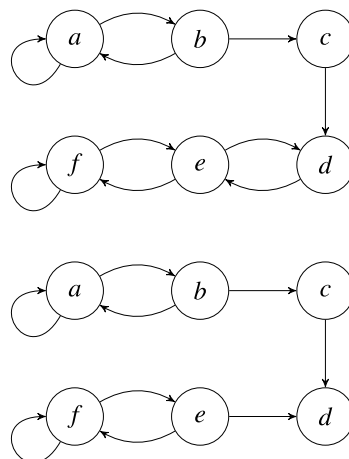

Fig. 4. The argumentation frameworks AF_2 (top) and AF_3 (bottom) from Example 7.

the first transition. Observe that $AF_2^{\{e\}} = AF_3^{\{e\}}$, so further constructions will be identical. Since the ideal extensions of AF_2 and AF_3 differ and no β can distinguish these cases, ideal semantics is not serialisable.

The above behaviour of ideal semantics is insofar surprising since the concept of unchallenged initial sets is closely related to the general idea of the ideal extension. Recall that the initial extension is the maximal admissible set contained in each preferred extension. This basically means that the arguments in the ideal extension are compatible with each admissible set and no admissible set attacks any argument in the ideal extension. On the other hand, an unchallenged initial set is likewise an undisputed admissible set, as there is no other initial set that attacks it. However, as the above example shows, unchallenged initial sets are not sufficient to characterise the ideal extension. We will discuss the relationship between unchallenged initial sets and the ideal extension a bit more below and in Section 5.

Likewise negative, but for somewhat different reasons, is the following result about semi-stable semantics.

Theorem 9. *Semi-stable semantics is not serialisable.*

The reason that semi-stable semantics is not serialisable is that it needs some form of "global" view on candidate extensions to judge whether a set of arguments is indeed a semi-stable extension. We illustrate this in the next example (which also serves as the proof of Theorem 9).

Example 8. Consider the three argumentation frameworks AF_4, AF_5, AF_6 in Fig. 5. Observe that

$$\mathsf{sst}(AF_4) = \big\{\{a, c\}, \{b\}\big\}$$
$$\mathsf{sst}(AF_5) = \big\{\{b\}\big\}$$
$$\mathsf{sst}(AF_6) = \big\{\{a\}, \{b\}\big\}$$

However,

$$\mathsf{IS}^{\leftarrow}(AF_4) = \mathsf{IS}^{\leftarrow}(AF_5) = \mathsf{IS}^{\leftarrow}(AF_6) = \emptyset$$
$$\mathsf{IS}^{\leftrightarrow\!\rightarrow}(AF_4) = \mathsf{IS}^{\leftrightarrow\!\rightarrow}(AF_5) = \mathsf{IS}^{\leftrightarrow\!\rightarrow}(AF_6) = \emptyset$$
$$\mathsf{IS}^{\leftrightarrow}(AF_4) = \mathsf{IS}^{\leftrightarrow}(AF_5) = \mathsf{IS}^{\leftrightarrow}(AF_6) = \big\{\{a\}, \{b\}\big\}$$

First, no selection function α can distinguish these frameworks (as it only operates on the sets $\mathsf{IS}^{\leftarrow}(\cdot)$, $\mathsf{IS}^{\leftrightarrow\!\rightarrow}(\cdot)$, and $\mathsf{IS}^{\leftrightarrow}(\cdot)$) and would either return
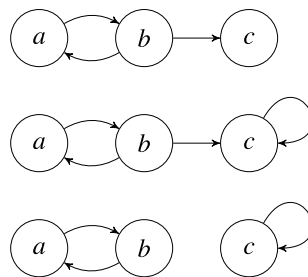

Fig. 5. The argumentation frameworks AF_4 (top), AF_5 (middle), and AF_6 (bottom) from Example 8.

(1) $\alpha(\emptyset, \emptyset, \{\{a\}, \{b\}\}) = \emptyset$,
(2) $\alpha(\emptyset, \emptyset, \{\{a\}, \{b\}\}) = \{\{a\}\}$,
(3) $\alpha(\emptyset, \emptyset, \{\{a\}, \{b\}\}) = \{\{b\}\}$, or
(4) $\alpha(\emptyset, \emptyset, \{\{a\}, \{b\}\}) = \{\{a\}, \{b\}\}$.

In cases 1–3, not all semi-stable extensions can be constructed for, e. g., AF_4. For case 4 and AF_5, a valid transition would then produce the state $T_1 = ((\{c\}, \{(c, c)\}), \{a\})$ from which no further transition is possible. A β function for semi-stable semantics should now determine that T_1 is not a terminating state (since $\{a\}$ is not a semi-stable extension of AF_5). However, the same state T_1 can also be reached for AF_6, but here $\{a\}$ is a semi-stable extension. Since β cannot distinguish these two scenarios at T_1, there is no such β.

The same argument from above can also be used to show that *eager semantics* is not serialisable. The *eager extension* is the maximal (wrt. set inclusion) admissible set contained in every semi-stable extension, see e. g. [5]. In Example 8, the eager extension of AF_5 is $\{b\}$ while it is $\emptyset$ for AF_4 and AF_6. No pair (α, β) can be defined to distinguish these cases as well.

The results of this section show that many admissibility-based semantics can be characterised through the notion of initial sets and a simple non-deterministic algorithm based on selecting initial sets at each step. This brings a new perspective on the rationality of admissibility-based semantics, as their basic construction principles are made explicit via an operational mechanism. This is similar in spirit to the purpose of *discussion games* [12] which model acceptability problems of individual arguments as an operational mechanism as well (here a dialogue between a proponent and an opponent). However, here we focused on the construction of whole extensions and not on acceptability problems of individual arguments.

The notion of serialisability also allows to define completely new semantics by defining only a selection and a termination function. For example, a straightforward idea for that would be the selection function α_0 defined via

$$\alpha_0(X, Y, Z) = X \cup Y$$

and the termination function β_0 defined via

$$\beta_0(\mathsf{AF}, S) = \begin{cases} 1 & \text{if } \mathsf{IS}^{\leftarrow}(\mathsf{AF}) \cup \mathsf{IS}^{\leftrightarrow}(\mathsf{AF}) = \emptyset \\ 0 & \text{otherwise} \end{cases}$$

which amounts to exhaustively adding unattacked and unchallenged initial sets. This yields a semantics that is more skeptical than the preferred semantics but less skeptical than the ideal semantics as the following result shows.

Theorem 10. *For every E with* $(\mathsf{AF}, \emptyset) \rightsquigarrow^{\alpha_0, \beta_0} (\mathsf{AF}', E)$

(1) $E \subseteq E'$ *for some preferred extension E' and*
(2) $E_{id} \subseteq E$ *for the ideal extension E_{id}.*

Consider the following examples showing the difference between the above semantics and ideal semantics.
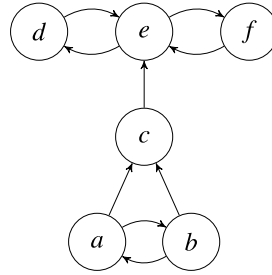
Fig. 6. The argumentation framework AF_7 from Example 10.

Example 9. Consider again AF_2 from Example 7 and depicted in Fig. 4. There are two extensions E_1 and E_2 wrt. to the serialisable semantics defined by α_0 and β_0:

$$E_1 = \{b\} \qquad E_2 = \{b, e\}$$

where the extension E_2 can be constructed by first selecting the initial set $\{e\}$ (which is unchallenged in AF_2) and then $\{b\}$.

Example 10. Consider AF_7 depicted in Fig. 6. There are four preferred extensions E_1, E_2, E_3, and E_4 in AF_7 defined via

$$E_1 = \{a, e\} \qquad E_2 = \{a, d, f\} \qquad E_3 = \{b, e\} \qquad E_4 = \{b, d, f\}$$

and the ideal extension E_{id} is empty:

$$E_{id} = \emptyset$$

However, there is one extension E_5 wrt. to the serialisable semantics defined by α_0 and β_0:

$$E_5 = \{d, f\}$$

The reason for that is that both $\{d\}$ and $\{f\}$ are unchallenged initial sets in AF_7 (and once one is selected the other becomes an unattacked initial set and can be selected as well).

For future work, we intend to investigate the above and further possibilities for serialisable semantics in more detail.

5. Computational complexity

In order to round up our investigation of initial sets, we now analyse them wrt. computational complexity. We assume familiarity with basic concepts of computational complexity and basic complexity classes such as P, NP, coNP, see [30] for an introduction. We also require knowledge of the "non-standard" classes DP (and its complement coDP), P^{NP}, and $P_{\parallel}^{NP}$. The class DP is the class of decision problems that are a conjunction of a problem in NP and a problem in coNP, i. e., in language notation DP $= \{L_1 \cap L_2 \mid L_1 \in NP, L_2 \in coNP\}$. The class P^{NP} is the class of decision problems that

can be solved by a deterministic polynomial-time algorithm that can make polynomially many *adaptive* queries[3] to an NP-oracle. The class $P_\parallel^{NP}$ [24] is the class of decision problems that can be solved by a deterministic polynomial-time algorithm that can make polynomially many *non-adaptive* queries to an NP-oracle. Note that $P_\parallel^{NP}$ is sometimes denoted by Θ_2^P and is equal to $P^{NP[\log]}$, i. e., the class of decision problems solvable by deterministic polynomial-time algorithm that can make logarithmically many *adaptive* NP-oracle calls [30]. Observe also that $DP \subseteq P_\parallel^{NP}$ and $coDP \subseteq P_\parallel^{NP}$ as well as $P_\parallel^{NP} \subseteq P^{NP}$.

We consider the following computational tasks for $\sigma \in \{IS, IS^{\leftarrow}, IS^{\leftrightarrow}, IS^{\leftrightarrow}\}$, cf. [22]:

Ver_σ Given $AF = (A, R)$ and $S \subseteq A$,
     decide whether $S \in \sigma(AF)$.
$Exists_\sigma$ Given $AF = (A, R)$,
     decide whether $\sigma(AF) \neq \emptyset$.
$Unique_\sigma$ Given $AF = (A, R)$,
     decide whether $|\sigma(AF)| = 1$.
$Cred_\sigma$ Given $AF = (A, R)$ and $a \in A$,
     decide whether there is $S \in \sigma(AF)$ with $a \in S$.
$Skept_\sigma$ Given $AF = (A, R)$ and $a \in A$,
     decide whether for all $S \in \sigma(AF)$, $a \in S$.

Note that we do not consider the problems $Exists_\sigma^{\neg\emptyset}$ [22] (which asks whether there is a non-empty (initial) set) as these are equivalent to $Exists_\sigma$ due to the non-emptiness of all types of initial sets.

Table 1 summarises our results on the complexity of the above tasks, all proofs can be found in the appendix.

The results for IS mirror the results on admissible sets [22], with a few exceptions. While the problem of deciding whether a set S is admissible can be solved in logarithmic space with polynomial time [22], we only showed that the problem of deciding whether a set S is an initial set can be solved in polynomial time. It is unlikely (though a formal proof is missing at the moment) that we can strengthen these results in the same way as for admissible sets, since the minimality condition of an initial set suggest that subsets (of linear size) have to be constructed in an algorithm. Moreover, while the problems *Exists* and *Skept* are trivial for admissible sets [22] (since the empty set is always admissible), they are intractable for initial sets. The problem $Unique_{IS}$ is DP-complete (it is coNP-complete for admissible sets) since existence of initial sets is not guaranteed.

We get some different complexity characterisations for the different types of initial sets. All tasks for unattacked initial sets are tractable, as only unattacked arguments have to be considered. All tasks

Table 1

Complexity of computational tasks related to initial sets. An attached "-h" refers to hardness for the class and an attached "-c" refers to completeness for the class. All hardness results are wrt. polynomial reductions

σ	IS	$IS^{\leftarrow}$	$IS^{\leftrightarrow}$	$IS^{\leftrightarrow}$
Ver_σ	in P	in P	coNP-c	NP-c
$Exists_\sigma$	NP-c	in P	$P_\parallel^{NP}$-c	NP-c
$Unique_\sigma$	DP-c	in P	in $P_\parallel^{NP}$, DP-h	trivial
$Cred_\sigma$	NP-c	in P	$P_\parallel^{NP}$-c	NP-c
$Skept_\sigma$	coNP-c	in P	$P_\parallel^{NP}$-c	coNP-c

---

[3]A query is adaptive if it may depend on a previous query; queries are non-adaptive if they can be posed in any order or in parallel.

become harder (under standard complexity-theoretic assumptions) when only unchallenged initial sets are considered. In particular, Ver_σ is coNP-complete as it has to be verified that there is no other initial set attacking the set that is to be verified. Moreover, Theorem 10 already showed that there is some conceptual relation between unchallenged initial sets and the ideal extension. This is strengthened by our observation of the computational complexity of the other tasks pertaining to unchallenged initial sets, as all non-trivial tasks on ideal semantics are $P_\parallel^{NP}$-complete (under *randomised reductions*) [21]. We showed $P_\parallel^{NP}$-completeness (under *polynomial reductions*) for most of those tasks as well, with the exception being the problem $Unique_{IS \nleftrightarrow}$, where we only showed DP-hardness and a $P_\parallel^{NP}$-hardness proof remains an open problem.

The complexity of the tasks for challenged initial sets are similar as in the general case, with two exceptions. Verification of challenged initial sets is intractable as another initial set has to found that attacks the set under consideration. Moreover, $Unique_{IS \leftrightarrow}$ has the trivial answer NO for all instances as the existence of one challenged initial set S_1 implies the existence of another challenged initial set S_2 that attacks (and is attacked by) S_1.

6. Discussion

In this paper, we revisited the notion of initial sets, i. e., non-empty minimal admissible sets. We investigated their general properties and used them as basic building blocks to construct any admissible set. We have characterised many admissibility-based semantics via this approach and concluded our analysis with some notes on computational complexity.

Initial sets allow us to concisely explain why a certain argument can be accepted (e. g., whether it is contained in a preferred extension). We can deconstruct an extension via the characterisation of Corollary 1, justify the inclusion of initial sets along this characterisation – e. g., by pointing to the conflicts that had to be solved –, and arrive step by step at the argument in question.

A recent paper that addresses a similar topic as we do is [10]. There, Baumann and Ulbricht introduce *explanation schemes* as a way to explain the construction of extensions wrt. complete, admissible, and strongly admissible semantics. Let us consider the case of complete semantics. Given $AF = (A, R)$, the construction follows basically three steps:[4] (1) Determine the grounded extension E_0 of AF, (2) select a conflict-free set E_1 from the set of arguments appearing in some even-length cycle in the reduct AF^{E_0}, and (3) determine the grounded extension E_2 of the reduct $AF^{E_0 \cup E_1}$. If $E = E_0 \cup E_1 \cup E_2$ defends E_1 then E is a complete extension (and every complete extension can be decomposed in such a way). The overall construction is similar to our approach, in particular, it consists of a series of steps where we select a set of arguments and move to the reduct of the framework wrt. the arguments accumulated so far. However, our approach provides a more fine-grained way to construct extensions. In each step, we solve a single issue by selecting a single initial set. Step 2 of Baumann and Ulbricht's approach possibly solves a series of different conflicts all at once. This may actually diffuse the goal to provide an explanation why a certain extension is constructed as it is, as a selection of a conflict-free set of arguments from all even cycles may not clearly show, which conflicts are actually resolved. Furthermore, we do not need to explicitly use the notion of grounded semantics (and the general possibility to include defended arguments) in our construction, as it arises naturally through selecting (unattacked) initial sets and moving to the reduct immediately.

---

[4]Although an abbreviated two-step procedure is also discussed.

As a by-product of our work, we introduced a new *principle* [37] for abstract argumentation semantics: *serialisability*. For future work, we aim at investigated relationships of this new principle to other existing principles listed in [37]. Another avenue for future work to investigate whether our characterisations of admissibility-based semantics can be exploited for algorithmic purposes [14]. It is clear, that there is no obvious advantage in terms of computational complexity by computing (for example) a preferred extension via our transition system as it involves the computation of initial sets at each step (which is an intractable problem as *Exists*$_{IS}$ is already NP-complete). However, our characterisation of initial sets through strongly connected components (see Section 3) could be exploited to devise a parallel algorithm – see also [15] –, as initial sets can be calculated independently of each other in each strongly connected component.

Acknowledgements

Appendix. Proofs of technical results

Proposition 1. *If S is an initial set of* AF *then there is* $AF' = (A', R') \in SCC(AF)$ *s.t.* $S \subseteq A'$.

Proof. If S is a singleton the claim is trivially true. So let S be an initial set of AF with $|S| > 1$. Assume there are two different $AF' = (A', R') \in SCC(AF)$ and $AF'' = (A'', R'') \in SCC(AF)$ with $S \cap A' \neq \emptyset$ and $S \cap A'' \neq \emptyset$ and also $S = (S \cap A') \cup (S \cap A'')$ (the following proof generalises easily if S is spanned across more than two SCCs). If $(S \cap A')^- \subseteq (S \cap A')^+$ then $S \cap A'$ is admissible, contradicting the fact that S is an initial set. So there is at least one $a \in (S \cap A')^-$ with $a \in (S \cap A'')^+$ since the complete set S is admissible. With the same reasoning there must be a $b \in (S \cap A'')^-$ with $b \in (S \cap A')^+$. Then there is a closed circuit: from a there is an edge to an element c in $(S \cap A')$. Since $(S \cap A')$ is part of an SCC, there is a path from c to any other argument in $(S \cap A')$, in particular, also to an attacker d of b. From d we can go to b and then to an element in $(S \cap A'')$. Again through the SCC of $(S \cap A'')$ we can reach a. This contradicts the assumption, so S is contained in a single SCC. □

Proposition 2. *S is an initial set of* AF *if and only if S is an initial set of* $SCC(S) = (A', R')$ *and $S^- \subseteq A'$.*

Proof. Let S be an initial set of AF. As S is admissible, all arguments $a \in E^-$ are attacked by some $b \in E$. Proposition 1 already established that $S \subseteq A'$. Since each $a \in S^-$ attacks and is attacked by S, $S^- \subseteq A'$ as well. It follows that S is an initial set of SCC(S) as well as $S \cup S^- \subseteq A'$ and verifying whether a set is initial only needs to consider the relationships of those arguments. This also proves the other direction. □

Proposition 3. *Let $S \in IS(AF)$ and $SCC(S) = (A', R')$.*

  (1) *If S is unattacked then $|A'| = 1$.*
  (2) *If S is challenged or unchallenged then $|A'| > 1$.*
  (3) *If S is challenged and $S' \in conflicts(S, AF)$ then $SCC(S) = SCC(S')$.*

Proof. Let $S \in \mathsf{IS}(\mathsf{AF})$ and $\mathsf{SCC}(S) = (\mathsf{A}', \mathsf{R}')$.

(1) Since $S = \{a\}$ is not attacked there is no other argument b with a path to a. It follows directly $|\mathsf{A}'| = 1$.
(2) As S is admissible and there is at least one b that attacks S, one argument in S must attack b. It follows $S \cup \{b\} \in \mathsf{SCC}(S)$ and therefore $|\mathsf{A}'| > 1$.
(3) Let $S' \in \mathsf{conflicts}(S, \mathsf{AF})$. As $S\mathsf{R}S'$ and $S'\mathsf{R}S$ as well as both S and S' are completely in one SCC, it follows $\mathsf{SCC}(S) = \mathsf{SCC}(S')$. $\quad\square$

Proposition 4. *Let* $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ *be an abstract argumentation framework and* $S, S' \in \mathsf{IS}(\mathsf{AF})$ *with* $S \neq S'$.

(1) *If* $S' \in \mathsf{IS}^{\leftarrow}(\mathsf{AF})$ *then* $S' \in \mathsf{IS}^{\leftarrow}(\mathsf{AF}^S)$
(2) *If* $S' \in \mathsf{conflicts}(S, \mathsf{AF})$ *then* $S' \notin \mathsf{IS}(\mathsf{AF}^S)$
(3) *If* $S' \notin \mathsf{conflicts}(S, \mathsf{AF})$ *then* $S' \cap \bigcup \mathsf{IS}(\mathsf{AF}^S) \neq \emptyset$

Proof. Let $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ be an abstract argumentation framework and $S, S' \in \mathsf{IS}(\mathsf{AF})$ with $S \neq S'$

(1) Let $S' \in \mathsf{IS}^{\leftarrow}(\mathsf{AF})$, so $S = \{a\}$. As a is not attacked in AF, it is also not attacked by S and it follows $a \in \mathsf{A}'$ for $\mathsf{AF}^S = (\mathsf{A}', \mathsf{R}')$. It also follows that a is not attacked in AF^S and therefore $\{a\} \in \mathsf{IS}^{\leftarrow}(\mathsf{AF}^S)$.
(2) If $S' \in \mathsf{conflicts}(S, \mathsf{AF})$ then there is $a \in S$ and $b \in S'$ with $a\mathsf{R}b$. It follows $b \notin \mathsf{A}'$ for $\mathsf{AF}^S = (\mathsf{A}', \mathsf{R}')$. So $S' \notin \mathsf{IS}(\mathsf{AF}^S)$.
(3) As $S' \notin \mathsf{conflicts}(S, \mathsf{AF})$ it follows that $S \cup S'$ is conflict-free and therefore $S' \subseteq \mathsf{A}'$ for $\mathsf{AF}^S = (\mathsf{A}', \mathsf{R}')$. Furthermore, since S' is admissible in AF, $S' \cap \mathsf{A}'$ remains admissible in AF^S. By definition, it follows that there is an initial set S'' of AF^S with $S'' \subseteq S'$, proving the claim. $\quad\square$

Corollary 1. *Every non-empty admissible set S can be written as* $S = S_1 \cup \cdots \cup S_n$ *with pairwise disjoint* S_i, $i = 1, \ldots, n$, S_1 *is an initial set of* AF *and every* S_i, $i = 2, \ldots, n$ *is an initial set of* $\mathsf{AF}^{S_1 \cup \cdots \cup S_{i-1}}$. *Furthermore, only non-empty admissible sets S can be written in such a fashion.*

Proof. This follows by iterative application of Theorem 1. $\quad\square$

Theorem 2. *Admissible semantics is serialisable with*

$$\alpha_{adm}(X, Y, Z) = X \cup Y \cup Z \qquad \beta_{adm}(\mathsf{AF}, S) = 1$$

Proof. Follows from Corollary 1. $\quad\square$

Theorem 3. *Complete semantics is serialisable with* α_{adm} *and*

$$\beta_{co}(\mathsf{AF}, S) = \begin{cases} 1 & \text{if } \mathsf{IS}^{\leftarrow}(\mathsf{AF}) = \emptyset \\ 0 & \text{otherwise} \end{cases}$$

Proof. We have to show that E is a complete extension if and only if $(\mathsf{AF}, \emptyset) \rightsquigarrow^{\alpha_{adm}, \beta_{co}} (\mathsf{AF}', E)$ for some AF'. Let $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$.

- "⇒":
  Let E be a complete extension. By Corollary 1 and the fact that α_{adm} does not constrain the selection of initial sets, it is clear that there is AF' with $(\mathsf{AF}, \emptyset) \leadsto^\alpha (\mathsf{AF}', E)$. It remains to show that $\beta_{\text{co}}(\mathsf{AF}', E) = 1$. As E is complete, there is no argument $a \in \mathsf{A}$ s.t. all attackers of a (in AF) are contained in E. This is equivalent to stating that there is no unattacked argument a in AF^E and therefore $\mathsf{IS}^{\leftarrow}(\mathsf{AF}^E) = \emptyset$. As $\mathsf{AF}^E = \mathsf{AF}'$ the claim follows.
- "⇐":
  Let $(\mathsf{AF}, \emptyset) \leadsto^{\alpha_{\text{adm}}, \beta_{\text{co}}} (\mathsf{AF}', E)$ for some $\mathsf{AF}' = (\mathsf{A}', \mathsf{R}')$. By Corollary 1, E is admissible. Since $\mathsf{IS}^{\leftarrow}(\mathsf{AF}') = \emptyset$ there is no argument $a \in \mathsf{A}'$ that is defended by E. Therefore, E is complete. □

Corollary 2. *Let $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ be an abstract argumentation framework and $S \subseteq \mathsf{A}$. S is complete if and only if either*

- *$S = \emptyset$ and $\mathsf{IS}^{\leftarrow}(\mathsf{AF}) = \emptyset$ or*
- *$S = S_1 \cup S_2$, $S_1 \in \mathsf{IS}(\mathsf{AF})$ and S_2 is complete in AF^{S_1}.*

Proof. If $\mathsf{IS}^{\leftarrow}(\mathsf{AF}) = \emptyset$ then $S = \emptyset$ is obviously a complete extension. Let S be any non-empty complete extension. Due to Theorem 1 there are sets S_1, S_2 with $S = S_1 \cup S_2$, $S_1 \in \mathsf{IS}(\mathsf{AF})$ and S_2 is admissible in AF^{S_1}. Assume S_2 is not complete in AF^{S_1}. Then there is an argument a in AF^{S_1} that is defended by S_2 but $a \notin S_2$. But then $a \notin S_1$ and $a \notin S_1^+$ (otherwise it would not be in AF^{S_1}) and S defends a as well in AF. This contradicts the assumption that S is complete. So S_2 is complete in AF^{S_1}.

For the other direction, let S be any set with $S = S_1 \cup S_2$, $S_1 \in \mathsf{IS}(\mathsf{AF})$ and S_2 is complete in AF^{S_1}. Assume that S is not complete in AF. Then there is an argument a that is defended by S but $a \notin S$. Since a is not contained in S_1 nor attacked by it, it follows that a is contained in AF^{S_1}. Let T be the attackers of a in AF. Since S attacks all arguments in T let $T_1 \subseteq T$ be those arguments in T attacked by some $b \in S_1$ and $T_2 = T \setminus T_1$. It follows that no argument of T_1 is contained in AF^{S_1} but all arguments in T_2 are contained in AF^{S_1} and are necessarily attacked by S_2. It follows that a is defended by S_2 in AF^{S_1}, contradicting the fact that S_2 is complete in AF^{S_1}. It follows that S is indeed complete. □

Theorem 4. *Grounded semantics is serialisable with*

$$\alpha_{gr}(X, Y, Z) = X$$

and β_{co}.

Proof. Let E_{gr} be the grounded extension of $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ and let E_k be such that

$$(\mathsf{AF}, \emptyset) \to^{\alpha_{gr}} (\mathsf{AF}_1, E_1) \to^{\alpha_{gr}} \cdots \to^{\alpha_{gr}} (\mathsf{AF}_k, E_k)$$

so that $\beta_{\text{co}}(\mathsf{AF}_k, E_k) = 1$. We first show $E_i \subseteq E_{gr}$, for $i = 1, \ldots, k$, by induction on i.

- $i = 1$: By definition of α_{gr} we have $E_1 = \{a\}$ for some unattacked argument $a \in \mathsf{A}$. Since E_{gr} is complete, it contains all unattacked arguments of A, therefore $E_1 \subseteq E_{gr}$.
- $i > 1$: By definition of α_{gr} we have $E_i = \{a\} \cup E_{i-1}$ for some unattacked argument $a \in \mathsf{A}$ in AF_i. Since a is unattacked in AF_i, all attackers of a in AF must be attacked by E_{i-1}. By assumption, $E_{i-1} \subseteq E_{gr}$ and a is defended by E_{gr} as well. Since E_{gr} is complete it follows $E_i \subseteq E_{gr}$.

Since Theorem 3 already established that E_k is complete and E_{gr} is the smallest complete extension, from $E_k \subseteq E_{gr}$ is follows $E_k = E_{gr}$ and therefore the claim. □

Corollary 3. *Let* $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ *be an abstract argumentation framework and* $S \subseteq \mathsf{A}$. *S is grounded if and only if either*

- $S = \emptyset$ *and* $\mathsf{IS}^{\leftarrow}(\mathsf{AF}) = \emptyset$ *or*
- $S = S_1 \cup S_2$, $S_1 \in \mathsf{IS}^{\leftarrow}(\mathsf{AF})$ *and* S_2 *is grounded in* AF^{S_1}.

Proof. If $\mathsf{IS}^{\leftarrow}(\mathsf{AF}) = \emptyset$ then $S = \emptyset$ is obviously the grounded extension. Assume S is the non-empty grounded extension of AF. Then there must be an argument $a \in S$ that is not attacked in AF. It follows $\{a\} \in \mathsf{IS}^{\leftarrow}(\mathsf{AF})$. Let S' be the grounded extension of $\mathsf{AF}^{\{a\}}$. Then $S' \cup \{a\}$ is complete in AF:

(1) $S' \cup \{a\}$ is admissible due to Theorem 1.
(2) a is defended by $S' \cup \{a\}$ as it is not attacked.
(3) every $b \in S'$ is defended by $S' \cup \{a\}$ as all attackers of b are either attacked by a (and therefore not in the reduct $\mathsf{AF}^{\{a\}}$) or some argument in S' (since S' is grounded in $\mathsf{AF}^{\{a\}}$).

Assume there is a proper subset $S'' \subset S' \cup \{a\}$ that is complete. Since a is not attacked, $a \in S''$. It can easily be seen that $S' \cap S''$ would be complete in $\mathsf{AF}^{\{a\}}$ and $S' \cap S'' \subset S'$, contradicting the fact that S' is the grounded extension of $\mathsf{AF}^{\{a\}}$. It follows that $S = S' \cup \{a\}$.
The other direction is analogous. □

Theorem 5. *Stable semantics is serialisable with* α_{adm} *and*

$$\beta_{st}(\mathsf{AF}, S) = \begin{cases} 1 & \text{if } \mathsf{AF} = (\emptyset, \emptyset) \\ 0 & \text{otherwise} \end{cases}$$

Proof. Let S be a stable extension. By Theorem 2 it is clear (since any stable extension is admissible) that there is AF' with $(\mathsf{AF}, \emptyset) \leadsto^{\alpha_{adm}} (\mathsf{AF}', S)$. As $S \cup S^+ = \mathsf{A}$ it follows $\mathsf{AF}' = (\emptyset, \emptyset)$ and therefore $\beta_{\mathsf{AF}',S} = 1$. Furthermore, for any S with $(\mathsf{AF}, \emptyset) \leadsto^{\alpha_{adm},\beta_{st}} ((\emptyset, \emptyset), S)$ it follows that S is admissible and there is no $a \in \mathsf{A}$ with $a \notin S^+$ or $a \in S$. This is equivalent to stating that S is stable. □

Corollary 4. *Let* $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ *be an abstract argumentation framework and* $S \subseteq \mathsf{A}$. *S is stable if and only if either*

- $S = \emptyset$ *and* $\mathsf{A} = \emptyset$ *or*
- $S = S_1 \cup S_2$, $S_1 \in \mathsf{IS}(\mathsf{AF})$ *and* S_2 *is stable in* AF^{S_1}.

Proof. If $\mathsf{A} = \emptyset$ then $S = \emptyset$ is obviously the (only) stable extension. Assume S is a non-empty stable extension of AF. Due to Theorem 1 there are sets S_1, S_2 with $S = S_1 \cup S_2$, $S_1 \in \mathsf{IS}(\mathsf{AF})$ and S_2 is admissible in AF^{S_1}. Assume S_2 is not stable in $\mathsf{AF}^{S_1} = (\mathsf{A}', \mathsf{R}')$. Then there is $a \in \mathsf{A}'$ that is not attacked by S_2. Since $a \in \mathsf{A}'$, a is also not attacked by S_1. It follows that S is not a stable extension, in contradiction to the assumption. It follows that S_2 is stable in AF^{S_1}.

For the other direction, let $S = S_1 \cup S_2$, $S_1 \in \mathsf{IS}(\mathsf{AF})$ and S_2 is stable in $\mathsf{AF}^{S_1} = (\mathsf{A}', \mathsf{R}')$. By Theorem 1, S is admissible. Assume S is not stable, then there is $a \in \mathsf{A}$ that is not attacked by S. It follows that $a \in \mathsf{A}'$, so a is also not attacked by S_2 in AF^{S_1}, in contradiction to the assumption that S_2 is stable in AF^{S_1}. It follows that S is stable in AF. □

Theorem 6. *Preferred semantics is serialisable with α_{adm} and*

$$\beta_{pr}(\mathsf{AF}, S) = \begin{cases} 1 & \text{if } \mathsf{IS}(\mathsf{AF}) = \emptyset \\ 0 & \text{otherwise} \end{cases}$$

Proof. Let S be preferred. Due to Theorem 2 it follows that there is AF' with $(\mathsf{AF}, \emptyset) \leadsto^{\alpha_{adm}} (\mathsf{AF}', S)$. If $\mathsf{IS}(\mathsf{AF}') \neq \emptyset$ then there is another admissible set S' with $(\mathsf{AF}', S) \leadsto^{\alpha_{adm}} (\mathsf{AF}'', S')$ and $S \subsetneq S'$, in contradiction to the assumption that S is preferred. It follows $\mathsf{IS}(\mathsf{AF}') = \emptyset$ and therefore $(\mathsf{AF}, \emptyset) \leadsto^{\alpha_{adm}, \beta_{pr}} (\mathsf{AF}', S)$.

For the other direction, let $(\mathsf{AF}, \emptyset) \leadsto^{\alpha_{adm}, \beta_{pr}} (\mathsf{AF}', S)$. By Theorem 2 it is clear that S is admissible. Assume S is not preferred, so there is admissible S' with $S \subsetneq S'$. Define $S'' = S' \setminus S$. We show now that S'' is an admissible set of $\mathsf{AF}' = \mathsf{AF}^S$. First, since S' is conflict-free so is S''. Let now $a \in S''$. As S' is admissible we have $a^- \subseteq (S')^+$. Let $a^- = X_1 \cup X_2$ with disjoint sets X_1 and X_2 such that $X_1 = a^- \cap S^+$ and $X_2 = a^- \setminus X_1$. As $\mathsf{AF}' = \mathsf{AF}^S$, those attackers of a in X_1 are not present in AF' anymore, so there is no need to defend against them. However, since S' is admissible, $X_2 \subseteq (S')^+$ and it follows $X_2 \subseteq (S'')^+$ (as arguments in X_2 are not attacked by arguments in S). So S'' defends a and S'' is therefore admissible. As S'' is a (non-empty) admissible set of AF', it follows that $\mathsf{IS}(\mathsf{AF}') \neq \emptyset$. This contradicts the assumption that $(\mathsf{AF}, \emptyset) \leadsto^{\alpha_{adm}, \beta_{pr}} (\mathsf{AF}', S)$ and it follows that S is indeed a preferred extension. $\square$

Corollary 5. *Let $\mathsf{AF} = (A, R)$ be an abstract argumentation framework and $S \subseteq A$. S is preferred if and only if either*

- *$S = \emptyset$ and $\mathsf{IS}(\mathsf{AF}) = \emptyset$ or*
- *$S = S_1 \cup S_2$, $S_1 \in \mathsf{IS}(\mathsf{AF})$ and S_2 is preferred in AF^{S_1}.*

Proof. If $\mathsf{IS}(\mathsf{AF}) = \emptyset$ then $S = \emptyset$ is obviously the only preferred extension, since it is the only admissible set. Assume S is a non-empty preferred extension of AF. Due to Theorem 1 there are sets S_1, S_2 with $S = S_1 \cup S_2$, $S_1 \in \mathsf{IS}(\mathsf{AF})$ and S_2 is admissible in AF^{S_1}. Assume S_2 is not preferred in AF^{S_1}. Then there is admissible S_2' with $S_2 \subsetneq S_2'$. By Theorem 1, $S_1 \cup S_2'$ is admissible in AF and $S = S_1 \cup S_2 \subsetneq S_1 \cup S_2'$, contradicting the fact that S is preferred. It follows that S_2 is preferred in AF^{S_1}.

For the other direction, let $S = S_1 \cup S_2$, $S_1 \in \mathsf{IS}(\mathsf{AF})$ and S_2 is preferred in AF^{S_1}. By Theorem 1, S is admissible. Assume S is not preferred, then there is admissible S' with $S \subsetneq S'$. Since $S_1 \subseteq S'$, $S' \setminus S_1$ must be completely contained in AF^{S_1} (otherwise S' would not be conflict free) and $S_2 \subsetneq S' \setminus S_1$. $S' \setminus S_1$ is also necessarily admissible in AF^{S_1}, contradicting the fact that S_2 is preferred in AF^{S_1}. It follows that S is preferred. $\square$

Theorem 7. *Strong admissibility semantics is serialisable with α_{gr} and β_{adm}.*

Proof. Let S be a strongly admissible set. We show $(\mathsf{AF}, \emptyset) \leadsto^{\alpha_{gr}} (\mathsf{AF}', S)$ (note that we do not have to consider β_{adm} as this function always returns 1) by induction on the size of S.

(1) $|S| = 0$: trivial as $(\mathsf{AF}, \emptyset) \leadsto^{\alpha} (\mathsf{AF}, \emptyset)$ for any α via zero steps.
(2) $|S| = n$: Let $a \in S$ such that $S' = S \setminus \{a\}$ is strongly admissible (the existence of such a is guaranteed as a direct corollary of Theorem 5 in [13] and the definition of strong admissibility). By induction hypothesis $(\mathsf{AF}, \emptyset) \leadsto^{\alpha_{gr}} (\mathsf{AF}', S')$. As S is strongly admissible, it follows that a cannot

be attacked in AF' (otherwise a is not defended by (a subset of) S'). So we have $\{a\} \in \mathsf{IS}^{\leftarrow}(\mathsf{AF}')$ and $(\mathsf{AF}, \emptyset) \rightsquigarrow^{\alpha_{\mathsf{gr}}} (\mathsf{AF}', S') \rightsquigarrow^{\alpha_{\mathsf{gr}}} (\mathsf{AF}'', S)$.

For the other direction, let $(\mathsf{AF}, \emptyset) \rightsquigarrow^{\alpha_{\mathsf{gr}}} (\mathsf{AF}', S)$. We show that S is strongly admissible by induction on the size of S.

(1) $|S| = 0$: the empty set is by definition strongly admissible.
(2) $|S| = n$: Consider the final step in the construction of S, i. e., $(\mathsf{AF}, \emptyset) \rightsquigarrow^{\alpha_{\mathsf{gr}}} (\mathsf{AF}', S') \rightsquigarrow^{\alpha_{\mathsf{gr}}} (\mathsf{AF}'', S)$. As α_{gr} returns only singleton sets, we have $S = S' \cup \{a\}$ for some $a \in \mathsf{A}$. As a is unattacked in AF'', S' defends a in AF. By induction hypothesis, S' is strongly admissible, showing that S is strongly admissible. $\square$

Corollary 6. *Let* $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ *be an abstract argumentation framework and* $S \subseteq \mathsf{A}$. *S is strongly admissible if and only if either*

- $S = \emptyset$ *or*
- $S = S_1 \cup S_2$, $S_1 \in \mathsf{IS}^{\leftarrow}(\mathsf{AF})$ *and* S_2 *is strongly admissible in* AF^{S_1}.

Proof. Since $\emptyset$ is always strongly admissible, we only consider the second case. Assume S is a non-empty strongly admissible set. Since $S \subseteq E_{gr}$, where E_{gr} is the grounded extension of AF, there is an unattacked $a \in S$ and $\{a\} \in \mathsf{IS}^{\leftarrow}(\mathsf{AF})$. Let $S_2 = S \setminus \{a\}$ and $b \in S_2$. Since S is strongly admissible in AF, there is $S'' \subseteq S \setminus \{b\}$ that is strongly admissible and defends b. If $a \in S''$ then $S'' \setminus \{a\}$ is strongly admissible in $\mathsf{AF}^{\{a\}}$ and defends b. If $a \notin S''$ then S'' remains strongly admissible in $\mathsf{AF}^{\{a\}}$ and defends b. In any case, it follows that S_2 is strongly admissible in $\mathsf{AF}^{\{a\}}$.

The other direction is analogous. $\square$

Theorem 10. *For every E with* $(\mathsf{AF}, \emptyset) \rightsquigarrow^{\alpha_0, \beta_0} (\mathsf{AF}', E)$

(1) $E \subseteq E'$ *for some preferred extension E' and*
(2) $E_{id} \subseteq E$ *for the ideal extension E_{id}.*

Proof. (1) Note that $\alpha_0(X, Y, Z) \subseteq \alpha_{\mathsf{adm}}(X, Y, Z)$ for all X, Y, Z. So if $(\mathsf{AF}, \emptyset) \rightsquigarrow^{\alpha_0} (\mathsf{AF}', E)$ then $(\mathsf{AF}, \emptyset) \rightsquigarrow^{\alpha_{\mathsf{adm}}} (\mathsf{AF}', E)$. Furthermore, $(\mathsf{AF}', E) \rightsquigarrow^{\alpha_{\mathsf{adm}}, \beta_{\mathsf{pr}}} (\mathsf{AF}'', E')$ eventually with a preferred extension E' due to Theorem 6. This shows $E \subseteq E'$.

(2) Let E_{id} be the ideal extension of AF and let $E_{\mathsf{id}} = S_1 \cup \cdots \cup S_n$ with S_i being an initial set of $\mathsf{AF}^{S_1 \cup \cdots \cup S_{i-1}}$ for all $i = 1, \ldots, n$ (this representation exists due to Corollary 1 and the fact that E_{id} is admissible).

Let E with $(\mathsf{AF}, \emptyset) \rightsquigarrow^{\alpha_0, \beta_0} (\mathsf{AF}^E, E)$. Assume that $E_{\mathsf{id}} \nsubseteq E$ and let $k \in \{1, \ldots, n\}$ be the smallest integer such that $S_k \nsubseteq E$. Let $\hat{S}_k = E \setminus S_k$. We show now that $\hat{S}_k$ is admissible in $\mathsf{AF}^E = (\mathsf{A}', \mathsf{R}')$:

(a) $\hat{S}_k \subseteq \mathsf{A}'$: For the sake of contradiction, assume there is $a \in \hat{S}_k$ with $a \notin \mathsf{A}'$. Due to $\hat{S}_k = E \setminus S_k$ it follows that a is attacked by E. Then the admissible set E attacks E_{id}, contradicting the fact that E_{id} is the ideal extension.
(b) $\hat{S}_k$ is conflict-free: clear since $\hat{S}_k \subseteq E_{\mathsf{id}}$.
(c) $\hat{S}_k$ defends all its elements (in AF^E): recall that $S_1 \cup \cdots \cup S_k$ is an admissible set in AF (due to Corollary 1) and that $S_1 \cup \cdots \cup S_{k-1} \cup (S_k \setminus \hat{S}_k) \subseteq E$. Let a be an attacker of $\hat{S}_k$ in AF and $b \in S_1 \cup \cdots \cup S_k$ that attacks a. Then either $b \in E$ (meaning that $a \notin \mathsf{A}'$ and $\hat{S}_k$ does not need to defend against a in AF^E) or $b \in \hat{S}_k$ (meaning that $\hat{S}_k$ defends against a in AF^E).

It follows that $\hat{S}_k$ is admissible in AF^E. Then there must be an initial set $\hat{S}'_k \subseteq \hat{S}_k$. Assume $\hat{S}'_k$ is challenged by another initial set T. Then $T \cup E$ would be an admissible set that attacks E_{id}. It follows $\hat{S}'_k$ is unattacked or unchallenged in AF^E. This contradicts the fact that $(\mathsf{AF}, \emptyset) \leadsto^{\alpha_0, \beta_0} (\mathsf{AF}^E, E)$. Therefore we have $E_{\mathsf{id}} \subseteq E$. $\quad\square$

Lemma 1. *Let S be conflict-free and $a \in S$. Deciding whether there is an admissible set $S' \subseteq S$ with $a \in S'$ can be decided in polynomial time.*

Proof. In polynomial time we can check first whether S is already admissible. If not, define S_1 via

$$S_1 = F_{\mathsf{AF}}(S) \cap S$$

Note that $S_1 \subsetneq S$ (if $S_1 = S$ then S would already have been admissible). Furthermore, all $a \in S \setminus S_1$ are not defended by arguments in S and can therefore not be a member of any admissible set $S' \subseteq S$. It follows that, if there is an admissible set $S' \subseteq S$ with $a \in S'$ then $S' \subseteq S_1$. So if $a \notin S_1$ or ($a \in S_1$ and S_1 is admissible), we are finished. Otherwise define S_2 via

$$S_2 = F_{\mathsf{AF}}(S_1) \cap S_1$$

and continue as before. Note that moving from S_i to S_{i+1} at least one argument is discarded (otherwise we have found an admissible set). So we have to compute at maximum $S_1, \ldots, S_{|S|}$ and all computations are polynomial. $\quad\square$

Proposition 5. *Ver_{IS} is in P.*

Proof. In polynomial time we can check first whether the input S is admissible. Then, for each $a, b \in S$ with $a \neq b$ we can test whether $S \setminus \{b\}$ contains an admissible set including a (see Lemma 1). If this is the case for one pair a, b, S cannot be an initial set. If this is not the case for any a, b then S is an initial set. All checks are in polynomial time. $\quad\square$

Proposition 6. *$Exists_{\mathsf{IS}}$ is NP-complete.*

Proof. Equivalence of $Exists_{\mathsf{IS}}$ and $Exists^{\neg\emptyset}_{\mathsf{adm}}$ follows from the fact that every non-empty admissible set contains an initial set [38]. $Exists^{\neg\emptyset}_{\mathsf{adm}}$ is NP-complete [22]. $\quad\square$

Proposition 7. *$Unique_{\mathsf{IS}}$ is DP-complete.*

Proof. Let AF be the input argumentation framework. Note that $Unique_{\mathsf{IS}}$ can be solved by solving the two problems:

(1) decide whether AF has at least one initial set and
(2) decide whether AF has at most one initial set.

Problem 1 is $Exists_{\mathsf{IS}}$ and therefore NP-complete. The complement of problem 2 can be solved by non-deterministically guessing two different sets S_1 and S_2 and verifying that both are initial sets. Problem 2 is therefore in coNP and this shows DP-membership of $Unique_{\mathsf{IS}}$.

For hardness, we provide a reduction from the problem $Unique_{st}$, i. e., the problem of deciding whether an argumentation framework has a unique stable extension, cf. [18,22]. For that, we directly use construction Tr_4 from [23], which translates (with polynomial overhead) an argumentation framework AF into an argumentation framework $Tr_4(AF)$ such that $st(AF) = adm(Tr_4(AF)) \setminus \{\emptyset\}$, cf. Lemma 5 and Theorem 4 in [23]. Since for every pair of stable extensions E_1, E_2 it holds $E_1 \nsubseteq E_2$ and $E_2 \nsubseteq E_1$, it follows that $E_1, E_2 \in adm(Tr_4(AF)) \setminus \{\emptyset\}$, $E_1 \nsubseteq E_2$ and $E_2 \nsubseteq E_1$, and therefore $adm(Tr_4(AF)) \setminus \{\emptyset\} = IS(Tr_4(AF))$. It follows $st(AF) = IS(Tr_4(AF))$ and $|st(AF)| = 1$ if and only if $|IS(Tr_4(AF))| = 1$ and therefore the claim. $\square$

Proposition 8. *$Cred_{IS}$ is NP-complete.*

Proof. For NP-membership consider the following algorithm. Upon input $a \in A$ we guess a set $S \subseteq A$ with $a \in S$ and verify in polynomial time that S is an initial set (see Proposition 5). It follows that a is credulously accepted wrt. initial sets. This shows NP-membership.

For NP-hardness we do a reduction from $Cred_{st}$, i. e., the problem of deciding whether an argument is credulously accepted wrt. stable semantics. We use the same reduction as in the proof of Proposition 7, namely the construction Tr_4 from [23]. We already established in the proof of Proposition 7 that $st(AF) = IS(Tr_4(AF))$. It follows that an argument a is credulously accepted wrt. stable semantics in AF if and only if it is credulously accepted wrt. initial sets in $Tr_4(AF)$. $\square$

Proposition 9. *$Skept_{IS}$ is coNP-complete.*

Proof. For coNP-membership consider the following algorithm, which solves the complement problem in NP. Upon input $a \in A$ we guess a set $S \subseteq A$ with $a \notin S$ and verify in polynomial time that S is an initial set (see Proposition 5). It follows that a is not skeptically accepted wrt. initial sets. This shows coNP-membership for $Skept_{IS}$.

For coNP-hardness we do a reduction from $Skept_{st}$, i. e., the problem of deciding whether an argument is skeptically accepted wrt. stable semantics. We use the same reduction as in the proof of Proposition 7, namely the construction Tr_4 from [23]. We already established in the proof of Proposition 7 that $st(AF) = IS(Tr_4(AF))$. It follows that an argument a is skeptically accepted wrt. stable semantics in AF if and only if it is skeptically accepted wrt. initial sets in $Tr_4(AF)$. $\square$

Proposition 10. *$Ver_{IS^{\leftarrow}}$, $Exists_{IS^{\leftarrow}}$, $Unique_{IS^{\leftarrow}}$, $Cred_{IS^{\leftarrow}}$, and $Skept_{IS^{\leftarrow}}$ are in P.*

Proof. Note that all these problems only have to make a simple check on the input:

- $Ver_{IS^{\leftarrow}}$: Verifying whether a single argument is unattacked is in P.
- $Exists_{IS^{\leftarrow}}$: Checking whether there is an unattacked argument in a given input framework AF is in P.
- $Unique_{IS^{\leftarrow}}$: Checking whether there is a single unattacked argument in a given input framework AF is in P.
- $Cred_{IS^{\leftarrow}}$: this is equivalent to $Ver_{IS^{\leftarrow}}$ with input AF and $\{a\}$.
- $Skept_{IS^{\leftarrow}}$: checking whether a is the only unattacked argument in AF is in P. $\square$

Proposition 11. *$Ver_{IS^{\leftrightarrow}}$ is coNP-complete.*

Proof. For coNP-membership, we consider the complement problem of verifying that input S is *not* an unchallenged initial set in NP. We first check in polynomial time whether the input S is an initial set at all, see Proposition 5. Then we check whether S is an unattacked initial set, see Proposition 10. If it is not, we guess another set S' with $S'RS$ and verify in polynomial time that S' is an initial set. This shows that S is not an unchallenged initial set in NP.

For coNP-hardness, we do a reduction from 3UNSAT, i. e., the problem of deciding whether a propositional formula in conjunctive normal form with exactly three literals per clause is unsatisfiable. For that, we extend Reduction 3.6 from [22]. Let ϕ be an instance of 3UNSAT in set notation, i. e., $\phi = \{C_1, \ldots, C_n\}$ and $C_i = \{l_{1,i}, l_{2,i}, l_{3,i}\}$ with literals $l_{1,i}, l_{2,i}, l_{3,i}$ from a set of atoms At, for $i = 1, \ldots, n$. Define an abstract argumentation framework $\mathsf{AF}'_\phi = (\mathsf{A}'_\phi, \mathsf{R}'_\phi)$ via

$$\mathsf{A}'_\phi = \{\phi, \tilde{\phi}, \psi\} \cup \{C_1, \ldots, C_n\} \cup \{a, \neg a \mid a \in At\}$$

$$\begin{aligned}
\mathsf{R}'_\phi = {} & \big\{(C_1, \phi), \ldots, (C_n, \phi)\big\} \\
& \cup \big\{(l, C_i) \mid l \in C, i \in \{1, \ldots, n\}\big\} \\
& \cup \big\{(a, \neg a), (\neg a, a) \mid a \in At\big\} \\
& \cup \big\{(\tilde{\phi}, a), (\tilde{\phi}, \neg a) \mid a \in At\big\} \\
& \cup \big\{(\phi, \tilde{\phi}), (\phi, \psi), (\psi, \phi)\big\}
\end{aligned}$$

Figure 7 shows an example of the reduction. We first show that there is an initial set containing ϕ if and only if ϕ is satisfiable. For that, assume first that ϕ is satisfiable and let $I : At \to \{\mathsf{true}, \mathsf{false}\}$ be a model of ϕ. Consider the set

$$S_I = \big\{a \mid I(a) = \mathsf{true}\big\} \cup \big\{\neg a \mid I(a) = \mathsf{false}\big\} \cup \{\phi\}$$

First observe that S_I is conflict-free: as I is an interpretation it is not the case that $a, \neg a \in S$ for some $a \in At$. Furthermore, there are no attacks between any argument from $\{a, \neg a \mid a \in At\}$ to ϕ and vice versa. Now observe that S_I is admissible (in fact S_I is stable):

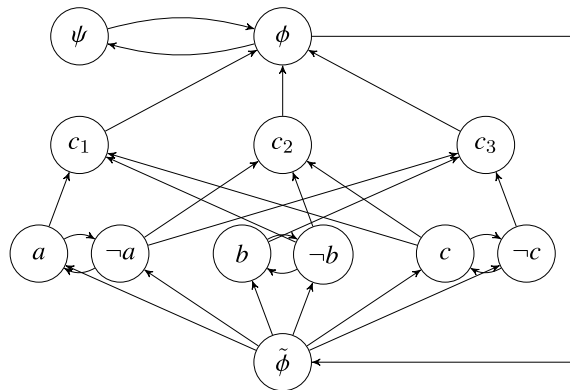(1) each $a \in S$ defends itself against $\neg a$ (for $a \in At$),

Fig. 7. The argumentation framework AF'_ϕ for $\phi = \{\{a, \neg b, c\}, \{\neg a, \neg b, c\}, \{\neg a, b, \neg c\}\}$.

(2) each $\neg a \in S$ defends itself against a (for $a \in At$),
(3) each C_i is attacked by some $a \in S$ or $\neg a \in S$ (since I is a model, every clause is satisfied), and
(4) $\tilde{\phi}$ is attacked by ϕ.

However, S_I is not necessarily an initial set. Consider the example in Fig. 7 again: here, I with $I(a) = I(b) = I(c) =$ false is a model of ϕ and we have

$$S_I = \{\neg a, \neg b, \neg c, \phi\}$$

Furthermore, $S'_I = \{\neg a, \neg b, \phi\} \subseteq S_I$ is also admissible (this happens when the truth value of one or more atoms does not matter for satisfiability). However, due to the fact that every non-empty admissible set contains an initial set, there is always an initial set $S'_I \subseteq S_I$ and S'_I must always contain ϕ as this is the only argument defending all arguments in $\{a, \neg a \mid a \in At\}$. It follows that if ϕ is satisfiable then there is an initial set containing ϕ. The other direction is analogous.

We now claim that ϕ is unsatisfiable if and only if $\{\psi\}$ is an unchallenged initial set. First, it is clear that $\{\psi\}$ is an initial set since ψ counterattacks the only attack. Moreover, we established above that ϕ is satisfiable if and only if there is an initial set S containing ϕ. So if ϕ is satisfiable $\{\psi\}$ is challenged by S. If ϕ is unsatisfiable then $\{\psi\}$ is clearly unchallenged. $\square$

Proposition 12. *$Exists_{IS^{\leftrightarrow}}$ is in $\mathsf{P}_{\parallel}^{\mathsf{NP}}$-complete.*

Proof. In order to show $\mathsf{P}_{\parallel}^{\mathsf{NP}}$-completeness, we use a characterisation of $\mathsf{P}_{\parallel}^{\mathsf{NP}}$ from [17].[5] More precisely, Theorem 9 of [17] establishes that a problem X is $\mathsf{P}_{\parallel}^{\mathsf{NP}}$-complete if and only if

(1) $X \in \mathsf{P}_{\parallel}^{\mathsf{NP}}$,
(2) X is NP-hard,
(3) X is coNP-hard,
(4) Two problem instances i_1 and i_2 of X can be polynomially reduced to a problem instance i_3 of X such that i_3 is a positive instance if and only if both i_1 and i_2 are positive instances.
(5) A set of problem instances $I = \{i_1, \ldots, i_k\}$ of X can be polynomially reduced to a problem instance i of X such that i is a positive instance if and only if there is at least one positive instance in I.

We now show that properties 1–5 above hold for the problem $Exists_{IS^{\leftrightarrow}}$.

(1) For $\mathsf{P}_{\parallel}^{\mathsf{NP}}$-membership, consider the following algorithm.[6] Let $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ be the input argumentation framework.

  1. For each argument $a \in \mathsf{A}$, check whether a is not attacked by an initial set
  2. For each argument $a \in \mathsf{A}$, check whether a is contained in an initial set
  3. Let M be the set of arguments for which both checks 1 and 2 were positive
  4. Remove all unattacked arguments from M, yielding a new set M'
  5. Compute the maximal admissible set M'' in M' (which is uniquely determined)
  6. If $M'' \neq \emptyset$ return YES, otherwise return NO

  First observe that the above algorithm runs in $\mathsf{P}_{\parallel}^{\mathsf{NP}}$. All checks in steps 1 and 2 can be solved by an NP-oracle:

---

[5]I am very grateful to an anonymous reviewer for pointing out that characterisation.
[6]Note that this algorithm is inspired by an algorithm for determining the ideal extension, cf. [20,21].

(a) The check in step 1 is a decision problem in coNP as the complement problem (i. e., checking whether a is attacked by an initial set) can be solved by guessing a set S that attacks a and verifying in polynomial time (see Proposition 5) that it is an initial set.

(b) The check in step 2 is a decision problem in NP: it can be solved by guessing a set S that contains a and verifying in polynomial time (see Proposition 5) that it is an initial set.

All checks in step 1 and 2 are non-adaptive, so they can be done in parallel (and there are linearly many of them). For step 4, observe that identifying unattacked arguments can be done in polynomial time. For step 5, note that M' is conflict-free (if there would be two arguments $a, b \in M'$ with $a\mathsf{R}b$, it means that both a and b are in some (possibly different) initial sets and that b is in some initial set that is attacked by some other initial set, which cannot be due to step 1). Determining the maximal admissible set in M' can be done similarly as in the proof of Lemma 1 in polynomial time.

Now we claim that the above algorithm returns YES if and only if the answer to $Exists_{\mathsf{IS}^{\leftrightarrow}}$ upon input AF is YES.

- "$\Rightarrow$": Let $M'' \neq \emptyset$ be the admissible set computed in step 5 and let $S \subseteq M''$ be an initial set contained in M'' (which necessarily exists since M'' is admissible and non-empty). Observe that S is not an unattacked initial set as we removed all unattacked arguments in step 4. Assume S is a challenged initial set. Then there exists another initial set S' which attacks some argument $b \in S$. This is in contradiction to the fact that $S \subseteq M' \subseteq M$ and M contains only arguments that are not attacked by an initial set (see step 1). So S is an unchallenged initial set and the answer to $Exists_{\mathsf{IS}^{\leftrightarrow}}$ upon input AF is YES.

- "$\Leftarrow$": Assume the answer to $Exists_{\mathsf{IS}^{\leftrightarrow}}$ upon input AF is YES. Then there exists an unchallenged initial set S. By definition, every argument $a \in S$ is contained in an initial set and not attacked by an initial set, so $S \subseteq M$ in step 3 of the above algorithm. Since S is unchallenged (but not unattacked), every argument in S is attacked and we have $S \subseteq M'$ in step 4. As S is admissible (and in no conflict with any other argument in M') we also have $S \subseteq M''$ in step 5 of the above algorithm. As $S \neq \emptyset$ it follows $M'' \neq \emptyset$ as well and the algorithm return YES.

In conclusion, $Exists_{\mathsf{IS}^{\leftrightarrow}}$ is in $\mathsf{P}_{\|}^{\mathsf{NP}}$.

(2) We show DP-hardness (which entails NP-hardness) instead. For that we use the same reduction from $Unique_{\mathsf{st}}$ as in the proof of Proposition 7, i. e., the construction Tr_4 from [23]. In particular, observe that if $\mathrm{Tr}_4(\mathsf{AF})$ has exactly one initial set S then S is unchallenged. Furthermore, if $\mathrm{Tr}_4(\mathsf{AF})$ has no initial sets then it obviously also has no unchallenged initial sets. If $\mathrm{Tr}_4(\mathsf{AF})$ has at least two initial sets, then all these initial sets are challenged (as they are all stable extensions of the original framework and two different stable extensions necessarily attack each other). This shows DP-hardness of $Exists_{\mathsf{IS}^{\leftrightarrow}}$.

(3) Since $Exists_{\mathsf{IS}^{\leftrightarrow}}$ is DP-hard (see above) it is also coNP-hard.

(4) Let $\mathsf{AF}_1 = (\mathsf{A}_1, \mathsf{R}_1)$ and $AF_2 = (\mathsf{A}_2, \mathsf{R}_2)$ be two argumentation frameworks and assume $\mathsf{A}_1 \cap \mathsf{A}_2 = \emptyset$ (otherwise rename arguments accordingly). Construct $\mathsf{AF}_3 = (\mathsf{A}_3, \mathsf{R}_3)$ as follows (let a_1, a_2 be fresh arguments):

$$\mathsf{A}_3 = \mathsf{A}_1 \cup \mathsf{A}_2 \cup \{a_1, a_2\}$$
$$\mathsf{R}_3 = \mathsf{R}_1 \cup \mathsf{R}_2 \cup \big\{(a_1, a_1), (a_2, a_2)\big\}$$

$$\cup \big\{(a, a_1) \mid a \in \mathsf{A}_1, a^- \neq \emptyset\big\} \cup \big\{(a_1, b) \mid b \in \mathsf{A}_2, b^- \neq \emptyset\big\}$$
$$\cup \big\{(b, a_2) \mid b \in \mathsf{A}_2, b^- \neq \emptyset\big\} \cup \big\{(a_2, a) \mid a \in \mathsf{A}_1, b^- \neq \emptyset\big\}$$

The intuition behind the above construction is that the two frameworks AF_1 and AF_2 are arranged in a circle where every (already attacked) argument of AF_1 attacks a_1, a_1 attacks every (already attacked) argument of AF_2, which in turn all attack a_2, which in turn attacks all (already attacked) arguments of AF_1. Obviously, the construction of AF_3 is polynomial in the size of AF_1 and AF_2.

We now show that AF_3 has an unchallenged initial set if and only if both AF_1 and AF_2 have unchallenged initial sets. Let M be an unchallenged initial set of AF_3. Since a_1 and a_2 attack themselves, $a_1, a_2 \notin M$. Assume $M \subseteq \mathsf{A}_1$. Since M is unchallenged and not unattacked, all $a \in M$ are attacked. It follows that a_2 attacks M. Since only arguments in AF_2 attack a_2, M cannot defend itself from a_2. It follows that $\mathsf{A}_2 \cap M \neq \emptyset$. For the same reason and using a_1 instead of a_2, it follows $\mathsf{A}_1 \cap M \neq \emptyset$. Let $M_1 = M \cap \mathsf{A}_1$ and $M_2 = M \cap \mathsf{A}_2$. Assume M_1 is challenged in AF_1 and let M_1' be an initial set in conflict with M_1 in AF_1. Then $M_1' \cup M_2$ is also an initial set of AF_3 and in conflict with M, contradicting the assumption that M is unchallenged. It follows that both M_1 and (with the same argument) M_2 are unchallenged in AF_1 and AF_2, respectively. For the other direction, given that M_1 and M_2 are unchallenged initial sets of AF_1 and AF_2, respectively, it is also clear that $M_1 \cup M_2$ is an unchallenged initial set of AF_3.

(5) Let $\mathsf{AF}_1 = (\mathsf{A}_1, \mathsf{R}_1), \ldots, \mathsf{AF}_n = (\mathsf{A}_n, \mathsf{R}_n)$ be argumentation frameworks and assume $\mathsf{A}_i \cap \mathsf{A}_j = \emptyset$ for all $i, j = 1, \ldots, n$ and $i \neq j$ (otherwise rename arguments accordingly). Observe that $AF = (\mathsf{A}_1 \cup \cdots \cup \mathsf{A}_n, \mathsf{R}_1 \cup \cdots \cup \mathsf{R}_n)$ has an unchallenged initial set M if and only if at least one of $\mathsf{AF}_1, \ldots, \mathsf{AF}_n$ has an unchallenged initial set (since necessarily $M \subseteq \mathsf{A}_i$ for some i due to the disconnectedness of AF). $\square$

Proposition 13. *$Unique_{IS^{\leftrightarrow}}$ is in $\mathsf{P}_{\parallel}^{\mathsf{NP}}$ and DP-hard.*

Proof. For $\mathsf{P}_{\parallel}^{\mathsf{NP}}$-membership, we use a similar algorithm as in the proof of Proposition 12. Let $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ be the input argumentation framework.

1. For each argument $a \in \mathsf{A}$, check whether a is not attacked by an initial set
2. For each argument $a \in \mathsf{A}$, check whether a is contained in an initial set
3. Let M be the set of arguments for which both checks 1 and 2 were positive
4. Remove all unattacked arguments from M, yielding a new set M'
5. Compute the maximal admissible set M'' in M' (which is uniquely determined)
6. If $M'' = \emptyset$ return NO
7. For each argument $a \in M''$, let $M_a = M'' \setminus \{a\}$
8. For each argument $a \in M''$, let M_a' be the maximal admissible set in M_a (which is uniquely determined)
9. If for all $a \in M''$, $M_a' \neq \emptyset$, return NO
10. Let $S = \{a \mid M_a' = \emptyset\}$
11. If S is an initial set return YES, otherwise return NO

First observe that the above algorithm runs in $\mathsf{P}_{\parallel}^{\mathsf{NP}}$. Steps 1–6 run in $\mathsf{P}_{\parallel}^{\mathsf{NP}}$ as already shown in the proof of Proposition 12. Furthermore, steps 7–11 run in (deterministic) polynomial time (in particular, step 8 runs in polynomial time by leveraging a similar algorithm as in the proof of Lemma 1 and step 11 because of Proposition 5).

We now claim that the above algorithm returns YES if and only if AF has a unique unchallenged initial set.

- "⇒": If the algorithm returns YES, we have obviously found an initial set S in step 11. As $S \subseteq M''$ we also have that S is an unchallenged initial set (see the proof of Proposition 12). Assume there exists an unchallenged initial set S' with $S' \neq S$. Note that both $S \nsubseteq S'$ and $S' \nsubseteq S$ since both are initial sets and $S' \subseteq M''$ (again, see the proof of Proposition 12). Let $x \in S \setminus S'$. Then $M'_x \neq \emptyset$ as M_x completely contains S'. This contradicts the fact that $x \in S$ due to step 10. It follows that AF has the unique unchallenged initial set S.
- "⇐": Assume AF has the unique unchallenged initial set M. By the argumentation in the proof of Proposition 12 we have that $M \subseteq M''$ in step 6. Then for all $a \in M''$ we have that

  * $M'_a = \emptyset$ if and only if $a \in M$ since $M_a \neq \emptyset$ would imply that there is another unchallenged initial set contained in M_a.
  * $M'_a \neq \emptyset$ if and only if $a \notin M$ as M is contained in M_a.

  By the definition of S in step 10 it follows $S = M$. As M is an initial set, the algorithm returns YES in step 11.

For showing DP-hardness we use the same reduction from $Unique_{st}$ as in the proof of Proposition 7, i. e., the construction Tr_4 from [23]. In particular, note that AF has exactly one stable extension if and only if $Tr_4(AF)$ has exactly one (unchallenged) initial set. □

Proposition 14. *$Cred_{IS^{\leftrightarrow}}$ is $P_\parallel^{NP}$-complete.*

Proof. In order to show $P_\parallel^{NP}$-completeness, we again use the characterisation of $P_\parallel^{NP}$ from [17] (see also Proposition 12). So we show that is $Cred_{IS^{\leftrightarrow}}$ is $P_\parallel^{NP}$-complete by showing that

(1) $Cred_{IS^{\leftrightarrow}} \in P_\parallel^{NP}$,
(2) $Cred_{IS^{\leftrightarrow}}$ is NP-hard,
(3) $Cred_{IS^{\leftrightarrow}}$ is coNP-hard,
(4) Two problem instances (AF_1, a_1) and (AF_2, a_2) of $Cred_{IS^{\leftrightarrow}}$ can be polynomially reduced to a problem instance (AF_3, a_3) of $Cred_{IS^{\leftrightarrow}}$ such that (AF_3, a_3) is a positive instance if and only if both (AF_1, a_1) and (AF_2, a_2) are positive instances.
(5) A set of problem instances $I = \{(AF_1, a_1), \ldots, (AF_k, a_k)\}$ of $Cred_{IS^{\leftrightarrow}}$ can be polynomially reduced to a problem instance (AF, a) of $Cred_{IS^{\leftrightarrow}}$ such that (AF, a) is a positive instance if and only if there is at least one positive instance in I.

We now show that properties 1–5 above hold:

(1) For showing $P_\parallel^{NP}$-membership, we use a similar algorithm as in the proof of Proposition 12. Let $AF = (A, R)$ be the input argumentation framework and x the input argument..

  1. For each argument $a \in A$, check whether a is not attacked by an initial set
  2. For each argument $a \in A$, check whether a is contained in an initial set
  3. Let M be the set of arguments for which both checks 1 and 2 were positive
  4. Remove all unattacked arguments from M, yielding a new set M'
  5. Compute the maximal admissible set M'' in M' (which is uniquely determined)
  6. If there is an initial set $S \subseteq M''$ with $x \in S$, return YES, otherwise return NO

First observe that the above algorithm runs in $\mathsf{P}^{\mathsf{NP}}_{\parallel}$ using two consecutive rounds of parallel calls (which is still in $\mathsf{P}^{\mathsf{NP}}_{\parallel}$ due to Proposition 2.1 of [24]): steps 1–5 run in $\mathsf{P}^{\mathsf{NP}}_{\parallel} \subseteq \mathsf{P}^{\mathsf{NP}}$ as already shown in the proof of Proposition 12, and step 6 can be solved by one further NP-oracle call (non-deterministically guess $S \subseteq M''$ and verify that it is an initial set with $x \in S$). The above algorithm also solves $Cred_{\mathsf{IS}^{\leftrightarrow}}$ as M'' contains all (and only) unchallenged initial sets.

(2) We show DP-hardness (which entails NP-hardness) instead. For that, we use the same reduction from $Unique_{\mathsf{st}}$ as in the proof of Proposition 7, i. e., the construction Tr_4 from [23], but first augment the input argumentation framework $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ with a fresh argument a (without any additional attacks), yielding an argumentation framework $\hat{\mathsf{AF}} = (\mathsf{A} \cup \{a\}, \mathsf{R})$. Note that E is a stable extension of AF if and only if $E \cup \{a\}$ is a stable extension of $\hat{\mathsf{AF}}$. By the same argumentation as in the proof of Proposition 7, AF has a unique stable extension E if and only if $\mathrm{Tr}_4(\hat{\mathsf{AF}})$ has the unique (and unchallenged) initial set $E \cup \{a\}$. So AF has a unique stable extension if and only if a is credulously accepted wrt. unchallenged initial sets in $\mathrm{Tr}_4(\hat{\mathsf{AF}})$.

(3) Since $Cred_{\mathsf{IS}^{\leftrightarrow}}$ is DP-hard (see above) it is also coNP-hard.

(4) Let $\mathsf{AF}_1 = (\mathsf{A}_1, \mathsf{R}_1)$ and $AF_2 = (\mathsf{A}_2, \mathsf{R}_2)$ be two argumentation frameworks and assume $\mathsf{A}_1 \cap \mathsf{A}_2 = \emptyset$ (otherwise rename arguments accordingly) and let (AF_1, a_1) and (AF_2, a_2) be two instances for $Cred_{\mathsf{IS}^{\leftrightarrow}}$. Construct an instance (AF_3, a_3) with $\mathsf{AF}_3 = (\mathsf{A}_3, \mathsf{R}_3)$ for $Cred_{\mathsf{IS}^{\leftrightarrow}}$ as follows (let a_3 be a fresh argument):

$$\mathsf{A}_3 = \mathsf{A}_1 \setminus \{a_1\} \cup \mathsf{A}_2 \setminus \{a_2\} \cup \{a_3\}$$

$$\mathsf{R}_3 = \big\{(b, c) \in \mathsf{R}_1 \mid b \neq a_1, c \neq a_1\big\}$$
$$\big\{(b, c) \in \mathsf{R}_2 \mid b \neq a_2, c \neq a_2\big\}$$
$$\big\{(a_3, c) \mid (a_1, c) \in \mathsf{R}_1\big\} \cup \big\{(b, a_3) \mid (b, a_1) \in \mathsf{R}_1\big\}$$
$$\big\{(a_3, c) \mid (a_2, c) \in \mathsf{R}_2\big\} \cup \big\{(b, a_3) \mid (b, a_2) \in \mathsf{R}_2\big\}$$

Informally speaking, AF_3 is simply the union of AF_1 and AF_2 where the two arguments a_1 (from AF_1) and a_2 (from AF_2) are merged into a new argument a_3 that retains all previous attacks of a_1 and a_2.

We now show that a_3 is credulously accepted wrt. unchallenged initial sets in AF_3 if and only if both a_1 and a_2 are credulously accepted wrt. unchallenged initial sets in AF_1 and AF_2, respectively. Without loss of generality, we assume that both a_1 and a_2 are attacked in AF_1 and AF_2, respectively (otherwise the problem trivialises since $\{a_1\}$ and/or $\{a_2\}$ are then unattacked initial sets). Assume now that a_3 is credulously accepted wrt. unchallenged initial sets in AF_3 and let M be an unchallenged initial set with $a_3 \in M$. Consider $M_1 = M \cap \mathsf{A}_1 \cup \{a_1\}$ and observe

- M_1 is admissible in AF_1: let $c \in \mathsf{A}_1$ be attacking M_1. Then c also attacks M in AF_3 and M defends itself in AF_3 through some $d \in M$. Since there are no attacks between AF_1 and AF_2 either $d = a_3$ or $d \in \mathsf{A}_1$. In the first case, $a_1 \in M_1$ then attacks c. In the second case $d \in M_1$ attacks c. It follows that M_1 is admissible in AF_1.
- M_1 is an initial set in AF_1: Assume there exists non-empty initial $M_1' \subsetneq M_1$. If $a_1 \notin M_1'$ then $M_1' \subsetneq M$ as well, contradicting the fact that M is initial. If $a_1 \in M_1'$, it is also easy to see that $M_1' \setminus \{a_1\} \cup \{a_3\} \cup (M \cap \mathsf{A}_2) \subsetneq M$ must be an initial set of AF_3.
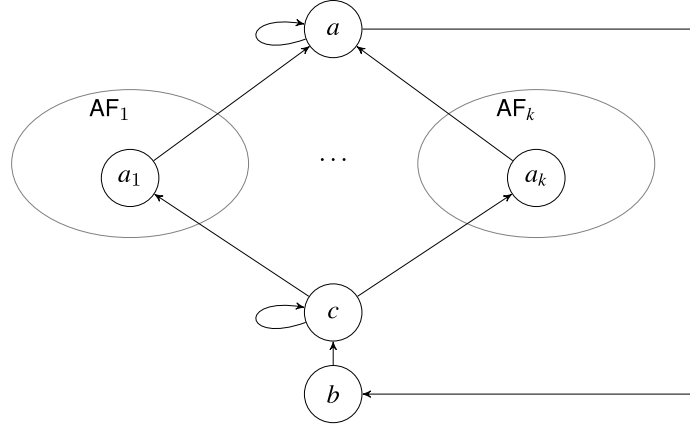
Fig. 8. A sketch of the argumentation framework AF from the proof of Proposition 14.

- M_1 is unchallenged: Suppose there is another initial set M_1' in conflict with M_1 in AF_1. If $a_1 \notin M_1'$ then M_1' also is in conflict with M in AF_3, contradicting the fact that M is unchallenged. The same follows for $a_1 \in M_1'$ with the same argumentation as above.

For the same reason it follows that $M_2 = M \cap \mathsf{A}_2 \cup \{a_2\}$ is an unchallenged initial set in AF_2 and, therefore, both a_1 and a_2 are credulously accepted wrt. unchallenged initial sets in AF_1 and AF_2, respectively. The other direction is analogous.

(5) Let $\mathsf{AF}_1 = (\mathsf{A}_1, \mathsf{R}_1), \ldots, \mathsf{AF}_n = (\mathsf{A}_n, \mathsf{R}_n)$ be argumentation frameworks and assume $\mathsf{A}_i \cap \mathsf{A}_j = \emptyset$ for all $i, j = 1, \ldots, n$ and $i \neq j$ (otherwise rename arguments accordingly). Let $I = \{(\mathsf{AF}_1, a_1), \ldots, (\mathsf{AF}_k, a_k)\}$ be a set of instances of $Cred_{\mathsf{IS}^{\leftrightarrow}}$. Construct $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ as follows (let a, b, c be fresh arguments):

$$\mathsf{A} = \mathsf{A}_1 \cup \cdots \cup \mathsf{A}_k \cup \{a, b, c\}$$

$$\mathsf{R} = \mathsf{R}_1 \cup \cdots \cup \mathsf{R}_k \cup \big\{(a, a), (a, b), (b, c), (c, c), (a_1, a), \ldots, (a_k, a), (c, a_1), \ldots, (c, a_k)\big\}$$

A sketch of the construction is shown in Fig. 8. We now show that b is credulously accepted wrt. unchallenged initial sets in AF if and only if there is a_i $(i = 1, \ldots, k)$ that is credulously accepted wrt. unchallenged initial sets in AF_i. Again, without loss of generality, we assume that all a_i $(i = 1, \ldots, k)$ are attacked in AF_i, respectively. Assume that b is credulously accepted wrt. unchallenged initial sets in AF and let M be an unchallenged initial set with $b \in M$. Since M is admissible and b is attacked by a, there must be $a_i \in M$ for some $i = 1, \ldots, k$. Let $M_i = M \setminus \{b\}$. Then M_i is necessarily an unchallenged initial set in AF_i (if, e. g., M_i is challenged in AF_i then M would also be challenged in AF). It follows that a_i is credulously accepted wrt. unchallenged initial sets in AF_i. The other direction is analogous. $\square$

Proposition 15. *$Skept_{\mathsf{IS}^{\leftrightarrow}}$ is $P_\parallel^{NP}$-complete.*

Proof. In order to show $P_\parallel^{NP}$-completeness, we use a characterisation of $P_\parallel^{NP}$ from [25], in particular Corollary 8. This allows us to show $P_\parallel^{NP}$-completeness of $Skept_{\mathsf{IS}^{\leftrightarrow}}$ by showing that

(1) $Skept_{\mathsf{IS}^{\leftrightarrow}} \in \mathsf{P}_{\|}^{\mathsf{NP}}$,
(2) $Skept_{\mathsf{IS}^{\leftrightarrow}}$ is coDP-hard,
(3) A set of problem instances $I = \{(\mathsf{AF}_1, a_1), \ldots, (\mathsf{AF}_k, a_k)\}$ of $Skept_{\mathsf{IS}^{\leftrightarrow}}$ can be polynomially reduced to a problem instance (AF, a) of $Skept_{\mathsf{IS}^{\leftrightarrow}}$ such that (AF, a) is a positive instance if and only if all instances in I are positive.

We now show that properties 1–3 above hold:

(1) For $\mathsf{P}_{\|}^{\mathsf{NP}}$-membership, we use a similar algorithm as in the proof of Proposition 12. Let $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ be the input argumentation framework and x the input argument.

   1. For each argument $a \in \mathsf{A}$, check whether a is not attacked by an initial set
   2. For each argument $a \in \mathsf{A}$, check whether a is contained in an initial set
   3. Let M be the set of arguments for which both checks 1 and 2 were positive
   4. Remove all unattacked arguments from M, yielding a new set M'
   5. Compute the maximal admissible set M'' in M' (which is uniquely determined)
   6. If $M'' = \emptyset$ return YES
   7. Let $M_x = M'' \setminus \{x\}$
   8. Let M'_x be the maximal admissible set in M_x (which is uniquely determined)
   9. If $M'_x = \emptyset$ return YES, otherwise return NO

   First observe that the above algorithm runs in $\mathsf{P}_{\|}^{\mathsf{NP}}$. Steps 1–6 run in $\mathsf{P}_{\|}^{\mathsf{NP}}$ as already shown in the proof of Proposition 12. Furthermore, steps 7–9 run in (deterministic) polynomial time (in particular, step 8 runs in polynomial time by leveraging a similar algorithm as in the proof of Lemma 1).
   We now claim that the above algorithm returns YES if and only if x is skeptically accepted wrt. unchallenged initial sets in AF.

   - "$\Rightarrow$": Assume the algorithm returns YES. First, assume the algorithm terminates in step 6. $M'' = \emptyset$ means that AF does not contain any unchallenged initial set (see the proof of Proposition 12). Then all arguments (including x) are trivially skeptically accepted. Now, assume the algorithm terminates in step 9. Since M'' is non-empty it contains unchallenged initial sets $S_1, \ldots, S_n$ (which are all unchallenged initial sets of AF, cf. the proof of Proposition 12). Since $M'_x = \emptyset$ it follows $S_i \nsubseteq M_x$ for all $i = 1, \ldots, n$. It follows $x \in S_i$ for all $i = 1, \ldots, n$ and x is skeptically accepted.
   - "$\Leftarrow$": Assume x is skeptically accepted wrt. unchallenged initial sets. Consider the following case differentiation:
     * There is no unchallenged initial set: in that case $M'' = \emptyset$ in step 6 (see again the proof of Proposition 12) and the algorithm returns YES in step 6.
     * There are unchallenged initial sets $S_1, \ldots, S_n$: since $x \in S_i$ for all $i = 1, \ldots, n$ we have that $M'_x = \emptyset$ (all unchallenged initial sets are "broken" by removing x). Then the algorithm returns YES in step 9.

(2) For showing coDP-hardness, we provide a reduction from the problem $\neg Unique_{\mathsf{st}}$, i. e., the problem of deciding whether an argumentation framework AF does *not* have a unique stable extension (which is naturally coDP-complete as its complement $Unique_{\mathsf{st}}$ is DP-complete). We use a similar approach as in the proof of Proposition 14, see also the proof of Proposition 7, again using the construction Tr_4 from [23]. For an input argumentation framework $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ and a fresh argument

a, we construct

$$\tilde{\mathsf{AF}} = \big(\mathsf{A} \cup \{a\}, \mathsf{R} \cup \big\{(b, a) \mid b \in \mathsf{A}\big\}\big)$$

In other words, we add an argument *a* and attacks from each original argument to *a*. Note that *E* is a stable extension of AF if and only if *E* is a stable extension of $\tilde{\mathsf{AF}}$. We now claim that an input argumentation framework AF does not possess a unique stable extension if and only if *a* is skeptically accepted wrt. unchallenged initial sets in $\tilde{\mathsf{AF}}'$.

- Assume AF has no stable extension. Then $\mathrm{Tr}_4(\tilde{\mathsf{AF}})$ has no initial set (see the proof of Proposition 7) and also no unchallenged initial set. Then every argument (including *a*) is trivially skeptically accepted wrt. unchallenged initial sets.
- Assume AF has a unique stable extension *E*. Then $\mathrm{Tr}_4(\tilde{\mathsf{AF}})$ has the unique (and unchallenged) initial set *E* that does not contain *a*. So *a* is not skeptically accepted, as desired.
- Assume AF has more than one stable extension. Then $\mathrm{Tr}_4(\tilde{\mathsf{AF}})$ has the same sets as (challenged) initial sets (as all these sets attack each other) and no unchallenged initial set. Then every argument (including *a*) is trivially skeptically accepted wrt. unchallenged initial sets.

For the other direction, assume *a* is skeptically accepted wrt. unchallenged initial sets. As observed above, this can only happen if there are no unchallenged sets and this can only happen if AF does not have a unique stable extension. This shows coDP-hardness of $Skept_{\mathsf{IS}^{\leftrightarrow}}$.

(3) Let $\mathsf{AF}_1 = (\mathsf{A}_1, \mathsf{R}_1), \ldots, \mathsf{AF}_n = (\mathsf{A}_n, \mathsf{R}_n)$ be argumentation frameworks and assume $\mathsf{A}_i \cap \mathsf{A}_j = \emptyset$ for all $i, j = 1, \ldots, n$ and $i \neq j$ (otherwise rename arguments accordingly). Let $I = \{(\mathsf{AF}_1, a_1), \ldots, (\mathsf{AF}_k, a_k)\}$ be a set of instances of $Skept_{\mathsf{IS}^{\leftrightarrow}}$. We again assume that each a_i is attacked in each AF_i. Consider the argumentation framework AF following the construction in the proof of Proposition 14, item (5), see also Fig. 8. We claim that *b* is skeptically accepted wrt. unchallenged initial sets in AF if and only if $a_1, \ldots, a_k$ are skeptically accepted wrt. unchallenged initial sets in $\mathsf{AF}_1, \ldots, \mathsf{AF}_k$, respectively. Without loss of generality, assume that a_1 is not skeptically accepted wrt. unchallenged initial sets in AF_1. Then there is an unchallenged initial set *M* in AF_1 with $a_1 \notin M$. Observe that *M* is also necessarily an unchallenged initial set in AF (since $a_1 \notin M$ there is no interaction with the rest of AF). Since $b \notin M$ it follows that *b* is also not skeptically accepted wrt. unchallenged initial sets in AF. The argument generalises naturally to all $i = 1, \ldots, k$. The other direction is analogous. $\square$

Proposition 16. *$Ver_{\mathsf{IS}^{\leftrightarrow}}$ is NP-complete.*

Proof. For NP-membership, on input AF and S_1 we guess a set S_2 that attacks S_1 and verify in polynomial time that S_2 is an initial set, cf. Proposition 5. This shows that S_1 is a challenged initial set.

To show NP-hardness, we use the same reduction as in the proof of Proposition 11 but using an instance of 3SAT. It is easy to see that an input instance ϕ is satisfiable if and only if $\{\psi\}$ is a challenged initial set in AF'_ϕ. $\square$

Proposition 17. *$Exists_{\mathsf{IS}^{\leftrightarrow}}$ is NP-complete.*

Proof. For NP-membership, we guess two sets S_1 and S_2 that attack each other and verify in polynomial time that both are initial sets, cf. Proposition 5. This shows actually that both S_1 and S_2 are challenged initial sets.

To show NP-hardness, we use the same reduction as in the proof of Proposition 11 but using an instance of 3SAT. It is easy to see that an input instance ϕ is satisfiable if and only if there is a challenged initial set (concretely, $\{\psi\}$) in AF'_ϕ. □

Proposition 18. *$Cred_{IS\leftrightarrow}$ is NP-complete.*

Proof. For NP-membership, on input AF and a we guess two sets S_1 and S_2 with $a \in S_1$, S_1 and S_2 attack each other and verify in polynomial time that both are initial sets, cf. Proposition 5. This shows actually that both S_1 and S_2 are challenged initial sets and a is credulously accepted wrt challanged initial sets (as $a \in S_1$).

To show NP-hardness, we use the same reduction as in the proof of Proposition 11 but using an instance of 3SAT. It is easy to see that an input instance ϕ is satisfiable if and only if ψ is credulously accepted wrt. challenged initial sets in AF'_ϕ. □

Proposition 19. *$Skept_{IS\leftrightarrow}$ is coNP-complete.*

Proof. For coNP-membership, we show that the complement problem, i. e., the problem of deciding whether an input argument a is *not* skeptically accepted wrt. initial sets, is in NP. For that we guess two sets S_1 and S_2 with $a \notin S_1$, S_1 and S_2 attack each other and verify in polynomial time that both are initial sets, cf. Proposition 5. This shows that a is not skeptically accepted and therefore $Skept_{IS\leftrightarrow}$ is in coNP.

To show coNP-hardness, we use the same reduction as in the proof of Proposition 11. It is easy to see that an input instance ϕ is unsatisfiable if and only if ψ is skeptically accepted wrt. challenged initial sets in AF'_ϕ. □

References

[1] L. Amgoud and J. Ben-Naim, Ranking-based semantics for argumentation frameworks, in: *Proceedings of the 7th International Conference on Scalable Uncertainty Management (SUM 2013)*, W. Liu, V.S. Subrahmanian and J. Wijsen, eds, Lecture Notes in Artificial Intelligence, Vol. 8078, Washington, DC, USA, 2013.

[2] L. Amgoud and H. Prade, Using arguments for making and explaining decisions, *Artificial Intelligence* **173**(3–4) (2009), 413–436. doi:10.1016/j.artint.2008.11.006.

[3] K. Atkinson, P. Baroni, M. Giacomin, A. Hunter, H. Prakken, C. Reed, G.R. Simari, M. Thimm and S. Villata, Toward artificial argumentation, *AI Magazine* **38**(3) (2017), 25–36. doi:10.1609/aimag.v38i3.2704.

[4] P. Baroni, G. Boella, F. Cerutti, M. Giacomin, L. van der Torre and S. Villata, On the input/output behavior of argumentation frameworks, *Artificial Intelligence* **217** (2014), 144–197. doi:10.1016/j.artint.2014.08.004.

[5] P. Baroni, M. Caminada and M. Giacomin, Abstract argumentation frameworks and their semantics, in: *Handbook of Formal Argumentation*, P. Baroni, D. Gabbay, M. Giacomin and L. van der Torre, eds, College Publications, 2018, pp. 159–236.

[6] P. Baroni, D. Gabbay, M. Giacomin and L. van der Torre (eds), *Handbook of Formal Argumentation*, College Publications, 2018.

[7] P. Baroni, M. Giacomin and G. Guida, SCC-recursiveness: A general schema for argumentation semantics, *Artificial Intelligence* **168**(1–2) (2005), 162–210. doi:10.1016/j.artint.2005.05.006.

[8] R. Baumann, G. Brewka and M. Ulbricht, Revisiting the foundations of abstract argumentation – semantics based on weak admissibility and weak defense, in: *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI'20)*, 2020.

[9] R. Baumann, T. Linsbichler and S. Woltran, Verifiability of argumentation semantics, in: *Proceedings of the 6th International Conference on Computational Models of Argument (COMMA'16)*, 2016, pp. 83–94.

[10] R. Baumann and M. Ulbricht, Choices and their consequences – explaining acceptable sets in abstract argumentation frameworks, in: *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning (KR'21)*, 2021.

[11] N. Boudjani, A. Gouaïch and S. Kaci, CLEAR: Argumentation frameworks for constructing and evaluating deductive mathematical proofs, in: *Proceedings of the 7th International Conference on Computational Models of Argument (COMMA'18)*, 2018, pp. 281–288.

[12] M. Caminada, Argumentation semantics as formal discussion, in: *Handbook of Formal Argumentation*, P. Baroni, D. Gabbay, M. Giacomin and L. van der Torre, eds, College Publications, 2018.

[13] M. Caminada and P.E. Dunne, Strong admissibility revisited: Theory and applications, *Argument & Computation* **10**(3) (2019), 277–300. doi:10.3233/AAC-190463.

[14] F. Cerutti, S.A. Gaggl, M. Thimm and J.P. Wallner, Foundations of implementations for formal argumentation, in: *Handbook of Formal Argumentation*, P. Baroni, D. Gabbay, M. Giacomin and L. van der Torre, eds, College Publications, 2018, Chapter 15. Also appears in IfCoLog Journal of Logics and their Applications 4(8):2623–2706, October 2017.

[15] F. Cerutti, I. Tachmazidis, M. Vallati, S. Batsakis, M. Giacomin and G. Antoniou, Exploiting parallelism for hard problems in abstract argumentation, in: *Proceedings of the 29th AAAI Conference (AAAI'15)*, 2015.

[16] L.A. Chalaguine and A. Hunter, A persuasive chatbot using a crowd-sourced argument graph and concerns, in: *Proceedings of the 8th International Conference on Computational Models of Argument (COMMA'20)*, 2020, pp. 9–20.

[17] R. Chang and J. Kadin, On computing Boolean connectives of characteristic functions, *Mathematical systems theory* **28**(3) (1995), 173–198. doi:10.1007/BF01303054.

[18] Y. Dimopoulos and A. Torres, Graph theoretical structures in logic programs and default theories, *Theoretical Computer Science* **170**(1–2) (1996), 209–244. doi:10.1016/S0304-3975(96)80707-9.

[19] P.M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games, *Artificial Intelligence* **77**(2) (1995), 321–358. doi:10.1016/0004-3702(94)00041-X.

[20] P.M. Dung, P. Mancarella and F. Toni, Computing ideal sceptical argumentation, *Artificial Intelligence* **171**(10) (2007), 642–674. doi:10.1016/j.artint.2007.05.003.

[21] P.E. Dunne, The computational complexity of ideal semantics, *Artificial Intelligence* **173**(18) (2009), 1559–1591. doi:10.1016/j.artint.2009.09.001.

[22] W. Dvořák and P.E. Dunne, Computational problems in formal argumentation and their complexity, in: *Handbook of Formal Argumentation*, P. Baroni, D. Gabbay, M. Giacomin and L. van der Torre, eds, College Publications, 2018, Chapter 14.

[23] W. Dvořák and S. Woltran, On the intertranslatability of argumentation semantics, *Journal of Artificial Intelligence Research* **41** (2011), 45–475.

[24] T. Eiter and G. Gottlob, The complexity class Theta2p: Recent results and applications in AI and modal logic, in: *Proceedings of the 11th International Symposium on Fundamentals of Computation Theory (FCT'97)*, Springer-Verlag, Berlin, Heidelberg, 1997, pp. 1–18.

[25] F. Frei, E. Hemaspaandra and J. Rothe, Complexity of stability, *Journal of Computer and System Sciences* **123** (2022), 103–121. doi:10.1016/j.jcss.2021.07.001.

[26] A. Hunter and M. Thimm, Probabilistic reasoning with abstract argumentation frameworks, *Journal of Artificial Intelligence Research* **59** (2017), 565–611. doi:10.1613/jair.5393.

[27] B. Liao and L. van der Torre, Explanation semantics for abstract argumentation, in: *Proceedings of the 8th International Conference on Computational Models of Argument (COMMA'20)*, 2020.

[28] S. Modgil and H. Prakken, The ASPIC+ framework for structured argumentation: A tutorial, *Argument & Computation* **5** (2014), 31–62. doi:10.1080/19462166.2013.869766.

[29] A. Niskanen and M. Järvisalo, Smallest explanations and diagnoses of rejection in abstract argumentation, in: *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning (KR'20)*, 2020, pp. 667–671.

[30] C. Papadimitriou, *Computational Complexity*, Addison-Wesley, 1994.

[31] A. Rago, O. Cocarascu, C. Bechlivanidis and F. Toni, Argumentation as a framework for interactive explanations for recommendations, in: *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning (KR'20)*, 2020, pp. 805–815.

[32] Z.G. Saribatur, J.P. Wallner and S. Woltran, Explaining non-acceptability in abstract argumentation, in: *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI'20)*, 2020.

[33] D. Seselja and C. Straßer, Abstract argumentation and explanation applied to scientific debates, *Synthese* **190**(12) (2013), 2195–2217. doi:10.1007/s11229-011-9964-y.

[34] M. Thimm, Revisiting minimal admissible sets in abstract argumentation, in: *Proceedings of the Second Workshop on Explainable Logic-Based Knowledge Representation (XLoKR'21)*, 2021.

[35] F. Toni, A tutorial on assumption-based argumentation, *Argument & Computation* **5**(1) (2014), 89–117. doi:10.1080/19462166.2013.869878.

[36] M. Ulbricht and J.P. Wallner, Strong explanations in abstract argumentation, in: *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI'21)*, 2021.

[37] L. van der Torre and S. Vesic, The principle-based approach to abstract argumentation semantics, in: *Handbook of Formal Argumentation*, Vol. 1, College Publications, 2018.

[38] Y. Xu and C. Cayrol, Initial sets in abstract argumentation frameworks, in: *Proceedings of the 1st Chinese Conference on Logic and Argumentation (CLAR'16)*, 2016.

[39] Y. Xu and C. Cayrol, Initial sets in abstract argumentation frameworks, *Journal of Applied Non-Classical Logics* **28**(2–3) (2018), 260–279. doi:10.1080/11663081.2018.1457252.

[40] Y. Xu, L. Xu and C. Cayrol, On structural analysis of extension-based argumentation semantics, in: *The Second Chinese Conference on Logic and Argumentation (CLAR'18)*, 2018.