Serialisable Semantics for Abstract Argumentation

Lars BENGEL and Matthias THIMM
Artificial Intelligence Group, University of Hagen, Germany

Abstract. We investigate the recently proposed notion of *serialisability* of semantics for abstract argumentation frameworks. This notion describes semantics where the construction of extensions can be serialised through iterative addition of minimal non-empty admissible sets. We investigate general relationships between serialisability and other principles from the literature. We also investigate the novel *unchallenged semantics* as a new instance of a serialisable semantics and, in particular, analyse it in terms of satisfied principles and computational complexity.

Keywords. abstract argumentation, serialisability, principles, computational complexity

1. Introduction

Abstract argumentation frameworks [1] are a simple, yet powerful formalism for representing argumentative scenarios and investigating matters regarding the acceptability of arguments. They consist simply of a set of arguments and an attack relation between arguments and can thus be represented as a directed graph. Abstract argumentation semantics [2] are used to interpret abstract argumentation frameworks by appropriately constraining the space of possible outcomes of the underlying argumentation. In particular, extension-based semantics define when a set of arguments (called extension) represents a plausible constellation of arguments that makes "sense" given the attacks in a framework. While we also consider extension-based semantics in this paper, it is noteworthy to mention that there are also other approaches for semantics such as the labelling-based approach [3], ranking and gradual semantics [4], and probabilistic approaches [5].

In [6] it has been shown that many of the mainstream extension-based semantics can be *serialised*, meaning that there is non-deterministic construction principle that allows to iteratively construct extensions by selecting minimal non-empty admissible sets—called *initial sets* [7]—and moving to the reduct [8]. Individual semantics can be distinguished by the way they select initial sets (via a so-called *selection function*) and how they terminate the construction (via a so-called *termination function*). For example, preferred semantics can be serialised by selecting initial sets arbitrarily until no further initial sets can be found [6]. Satisfaction of this principle of serialisability by a semantics allows a deeper inspection of the reasons why certain arguments are contained in an extension and therefore facilitates the explanatory power of an argumentation semantics [9,10]. In this paper, we aim at a better understanding of the principle of serialisability, in particular with respect to its connections to other principles for argumentation semantics

[11,12]. As it turns out, serialisability is an independent principle that is neither implied by nor implies other similar properties such as directionality.

As it has already been mentioned above, a serialisable semantics is characterised by a selection and termination function. This parametrisation of a semantics allows the easy development of further semantics simply by defining those two components. In [6], a specific candidate for such a semantics has already been suggested, which we coin here the *unchallenged semantics*. This semantics is defined by exhaustively adding *unattacked* and *unchallenged* initial sets (the formal definitions of these terms will be introduced in Section 3) and it has some interesting connections to preferred and ideal semantics. We investigate unchallenged semantics in more depth, in particular wrt. its compliance to principles from [11,12,13] and in terms of computational complexity. As with regard to the latter, unchallenged semantics turns out to be highly intractable, with credulous and skeptical reasoning shown to be Σ_2^P- and Π_2^P-complete, respectively.

To summarise, the contributions of this paper are as follows.

1. We recall the principle of serialisability and analyse its relationship with other principles (Section 3).
2. We investigate unchallenged semantics as a new instance of a serialisable semantics wrt. to its compliance to principles (Section 4).
3. We analyse unchallenged semantics wrt. computational complexity (Section 5).

Section 2 presents the necessary background on abstract argumentation and Section 6 concludes. Proofs of technical results are omitted due to space restrictions but can be found in an online appendix.[1]

2. Preliminaries

Let $\mathfrak{A}$ denote a universal set of arguments. An *abstract argumentation framework* AF is a tuple $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ where $\mathsf{A} \subseteq \mathfrak{A}$ is a finite set of arguments and R is a relation $\mathsf{R} \subseteq \mathsf{A} \times \mathsf{A}$ [1]. Let $\mathfrak{AF}$ denote the set of all abstract argumentation frameworks. For two arguments $a, b \in \mathsf{A}$, the relation $a\mathsf{R}b$ means that argument a attacks argument b. For $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ and $\mathsf{AF}' = (\mathsf{A}', \mathsf{R}')$ we write $\mathsf{AF}' \sqsubseteq \mathsf{AF}$ iff $\mathsf{A}' \subseteq \mathsf{A}$ and $\mathsf{R}' = \mathsf{R} \cap (\mathsf{A}' \times \mathsf{A}')$. For a set $X \subseteq \mathsf{A}$, we denote by $\mathsf{AF}|_X = (X, \mathsf{R} \cap (X \times X))$ the projection of AF on X. For a set $S \subseteq \mathsf{A}$ we define

$$S^+_{\mathsf{AF}} = \{a \in \mathsf{A} \mid \exists b \in S : b\mathsf{R}a\} \qquad\qquad S^-_{\mathsf{AF}} = \{a \in \mathsf{A} \mid \exists b \in S : a\mathsf{R}b\}$$

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

Different semantics can be phrased by imposing constraints on admissible sets [2]. In particular, an admissible set E

- is a *complete* (co) extension iff for all $a \in \mathsf{A}$, if E defends a then $a \in E$,
- is a *grounded* (gr) extension iff E is complete and minimally so,

[1] `http://mthimm.de/misc/lbmt_uncsem_proofs.pdf`

- is a *stable* (st) extension iff $E \cup E^+_{\mathsf{AF}} = \mathsf{A}$,
- is a *preferred* (pr) extension iff E is maximal.
- is a *semi-stable* (sst) extension iff $E \cup E^+_{\mathsf{AF}}$ 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 $\sigma \in \{\mathsf{co}, \mathsf{gr}, \mathsf{st}, \mathsf{pr}, \mathsf{sst}, \mathsf{id}, \mathsf{sa}\}$ let $\sigma(\mathsf{AF})$ denote the set of σ-extensions of AF.

3. Initial Sets and Serialisability

Non-empty minimal admissible sets have been coined *initial sets* by Xu and Cayrol [7].

Definition 1. For $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$, a set $S \subseteq \mathsf{A}$ 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 $\mathsf{IS}(\mathsf{AF})$ denote the set of initial sets of AF.

Initial sets are not supposed to be used to solve the whole argumentation represented in an argumentation framework, but rather a single atomic conflict within the framework. We can also differentiate between three types of initial sets [6].

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

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

In the following, we will denote with $\mathsf{IS}^{\not\leftarrow}(\mathsf{AF})$, $\mathsf{IS}^{\not\leftrightarrow}(\mathsf{AF})$, and $\mathsf{IS}^{\leftrightarrow}(\mathsf{AF})$ the set of unattacked, unchallenged, and challenged initial sets, respectively.

In [6] the notion of *serialisability* has been introduced as a new approach for constructing admissible sets (and extensions of a variety of semantics) iteratively using initial sets. This approach relies also on the notion of the reduct [8].

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

The idea behind the approach of [6] to construct admissible sets is quite simple: We solve an atomic conflict in AF by selecting an initial set S. Afterwards, we move to the reduct AF^S which may reveal further conflicts and therefore new initial sets. This process is continued until some termination criterion is satisfied. In order to formalise this idea, we need a way to select initial sets in each step and also a criterion for determining if the construction of an admissible set is finished. The following concepts have been defined for this purpose.

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 $\alpha : 2^{2^{\mathfrak{A}}} \times 2^{2^{\mathfrak{A}}} \times 2^{2^{\mathfrak{A}}} \to 2^{2^{\mathfrak{A}}}$ with $\alpha(X, Y, Z) \subseteq X \cup Y \cup Z$ for all $X, Y, Z \subseteq 2^{\mathfrak{A}}$.

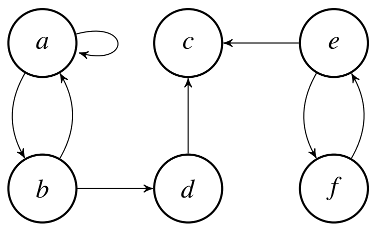

Figure 1. The argumentation framework AF_1 from Example 1.

We will apply a selection function α in the form $\alpha(\mathsf{IS}^{\not\leftarrow}(\mathsf{AF}), \mathsf{IS}^{\not\leftrightarrow}(\mathsf{AF}), \mathsf{IS}^{\leftrightarrow}(\mathsf{AF}))$ (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 $\beta(\mathsf{AF}, S) = 1$).

For some selection function α, consider the following transition rule:

$$(\mathsf{AF}, S) \xrightarrow{S' \in \alpha(\mathsf{IS}^{\not\leftarrow}(\mathsf{AF}), \mathsf{IS}^{\not\leftrightarrow}(\mathsf{AF}), \mathsf{IS}^{\leftrightarrow}(\mathsf{AF}))} (\mathsf{AF}^{S'}, S \cup S')$$

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 $\mathscr{E}^{\alpha,\beta}(\mathsf{AF})$ 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* if there exists a selection function α and a termination function β with $\sigma(\mathsf{AF}) = \mathscr{E}^{\alpha,\beta}(\mathsf{AF})$ for all AF. Then σ is also called the α, β-semantics.

In [6] it has already been shown that all of the standard admissible-based semantics adm, co, gr, pr and st as well as sa are serialisable. On the other hand, the semi-stable, ideal and eager semantics are not serialisable.

Example 1. As shown in [6], the preferred semantics can be serialised by the selection function $\alpha_{ad}(X, Y, Z) = X \cup Y \cup Z$ and the termination function

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

Consider the argumentation framework AF_1 in Figure 1. The initial sets of AF_1 are $\{b\}$, $\{e\}$ and $\{f\}$. In order to obtain the preferred extensions we start with the state $(\mathsf{AF}_1, \emptyset)$. According to the α_{ad} all three initial sets can be selected. Assume we select $\{b\}$ first, then we apply the transition rule as

$$(\mathsf{AF}_1, \emptyset) \xrightarrow{\{b\}} (\mathsf{AF}_1^{\{b\}}, \{b\}).$$

In this reduct $\mathsf{AF}_1^{\{b\}}$, we have two initial sets, namely $\{e\}$ and $\{f\}$. If we select $\{f\}$, the next transition would be

$$(\mathsf{AF}_1^{\{b\}}, \{b\}) \xrightarrow{\{f\}} (\mathsf{AF}_1^{\{b,f\}}, \{b,f\}).$$

This leaves us with one more possible transition via

$$(\mathsf{AF}_1^{\{b,f\}}, \{b,f\}) \xrightarrow{\{c\}} ((\emptyset,\emptyset), \{b,c,f\}).$$

Now, trivially, the termination function is true, since there is no initial set for the empty framework and $\{b,c,f\}$ is a preferred extension of AF_1. Similarly, we could have selected $\{e\}$ in the state $(\mathsf{AF}_1^{\{b\}}, \{b\})$. In that case, we also obtain an empty argumentation framework and the set $\{b,e\}$, which is the only other preferred extension of AF_1.

The principle of serialisability allows to define a semantics simply by specifying a selection function for initial sets and a termination function. In Section 4 we will define a completely new semantics using this approach and investigate its properties. However, in the remainder of this section we will analyse the principle of serialisability a bit deeper.

3.1. Relationship to Other Principles

In the following, we will look further at the serialisability principle and investigate its relationship with other principles from the literature [11]. First, we recall some basic definitions. Let $\mathsf{AF} = (\mathsf{A},\mathsf{R})$ be an argumentation framework. A set of arguments $U \subseteq \mathsf{A}$ is called *unattacked* if and only if $\nexists a \in (\mathsf{A} \setminus U) : a\mathsf{R}U$. The set of unattacked sets of AF is denoted as $\mathfrak{US}(\mathsf{AF})$. Furthermore, a set $S \subseteq \mathsf{A}$ is a *strongly connected component* of AF, if there is a directed path between any pair $a,b \in S$ in AF and there is no $S' \supset S$ with that property. Let $\mathsf{SCCs}_{\mathsf{AF}}$ be the set of strongly connected components of AF. For a set $S \subseteq \mathsf{A}$, we define $\mathrm{op}_{\mathsf{AF}}(S) = \{a \in \mathsf{A} \mid a \notin S \wedge a\mathsf{R}S\}$. In order to define the principle of *SCC-Recursiveness* [14], we need some additional concepts.

Definition 8. Given an argumentation framework $\mathsf{AF} = (\mathsf{A},\mathsf{R})$, a set $E \subseteq \mathsf{A}$ and a strongly connected component $S \in \mathsf{SCCs}_{\mathsf{AF}}$, we define:

- $D_{\mathsf{AF}}(S,E) = \{a \in S \mid (\mathrm{op}_{\mathsf{AF}}(S))\mathsf{R}a\}$,
- $P_{\mathsf{AF}}(S,E) = \{a \in S \mid (E \cap \mathrm{op}_{\mathsf{AF}}(S)) \not\mathsf{R} a \wedge \exists b \in (\mathrm{op}_{\mathsf{AF}}(S) \cap a^-_{\mathsf{AF}}) : E \not\mathsf{R} b\}$,
- $U_{\mathsf{AF}}(S,E) = S \setminus (D_{\mathsf{AF}}(S,E) \cup P_{\mathsf{AF}}(S,E))$.

Definition 9. Let $\mathsf{AF} = (\mathsf{A},\mathsf{R})$ be an argumentation framework and $C \subseteq \mathsf{A}$ is a set of arguments.

1. A function $\mathscr{BF}(\mathsf{AF},C)$ is called *base function*, if, given an argumentation framework $\mathsf{AF} = (\mathsf{A},\mathsf{R})$ such that $|\mathsf{SCCs}(\mathsf{AF})| = 1$ and a set $C \subseteq \mathsf{A}$, $\mathscr{BF}(\mathsf{AF},C) \subseteq 2^{\mathsf{A}}$.
2. Given a base function $\mathscr{BF}(\mathsf{AF},C)$ we define the function $\mathscr{GF}_{\mathscr{BF}}(\mathsf{AF},C) \subseteq 2^{\mathsf{A}}$ as follows: for any $E \subseteq \mathsf{A}, E \in \mathscr{GF}(\mathsf{AF},\mathscr{C})$ if and only if
    - in case $|\mathsf{SCCs}_{\mathsf{AF}}| = 1$, $E \in \mathscr{BF}(\mathsf{AF},C)$,
    - otherwise, $\forall S \in \mathsf{SCCs}_{\mathsf{AF}} : E \cap S \in \mathscr{GF}_{\mathscr{BF}}(\mathsf{AF}|_{S \setminus D_{\mathsf{AF}}(S,E)}, U_{\mathsf{AF}}(S,E) \cap C)$.

Definition 10. Let $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ be an argumentation framework and $E \subseteq \mathsf{A}$ is a set of arguments. We say that an argument $a \in \mathsf{A}$ is *strongly defended* by E (denoted as $sd(a,E)$) iff $\forall b \in \mathsf{A} : b\mathsf{R}a \implies \exists c \in E \setminus \{a\} : c\mathsf{R}b$ and $sd(c, E \setminus \{a\})$.

Finally, we recall the definitions of the different principles from the literature, that we considered in our analysis.

Definition 11. A semantics σ satisfies the principle of:

- *conflict-freeness [11]*, iff for every AF AF, every $E \in \sigma(\mathsf{AF})$ is conflict-free with respect to the attack relation.
- *admissibility [11]*, iff for every AF AF, every $E \in \sigma(\mathsf{AF})$ is conflict-free and defends itself in AF.
- *strong admissibility [11]*, iff for every AF AF, for every $E \in \sigma(\mathsf{AF})$ it holds that $a \in E$ implies that E strongly defends a.
- *reinstatement [11]*, iff for every AF $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ and $E \in \sigma(\mathsf{AF})$ we have: if E defends some $a \in \mathsf{A}$ then $a \in E$.
- *naivety [11]*, iff for every AF $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ and $E \in \sigma(\mathsf{AF})$ we have: E conflict-free and maximal among $cf(\mathsf{AF})$.
- *allowing abstention [15]*, iff for every AF AF and for every $a \in \mathsf{A}$, if there exist two extensions $E_1, E_2 \in \sigma(\mathsf{AF})$ such that $a \in E_1$ and $a \in E_2^+$ then there exists an extension $E_3 \in \sigma(\mathsf{AF})$ such that $a \notin (E_3 \cup E_3^+)$.
- *I-maximality [11]*, iff for every AF AF and $E_1, E_2 \in \sigma(\mathsf{AF})$, $E_1 \subseteq E_2 \implies E_1 = E_2$.
- *SCC-recursiveness [14]*, iff there is a base function $\mathscr{BF}_\sigma$ such that for every AF $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ we have that $\sigma(\mathsf{AF}) = \mathscr{GF}_{\mathscr{BF}_\sigma}(\mathsf{AF}, \mathsf{A})$.
- *directionality [11]*, iff for every AF $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ and $\forall U \in \mathfrak{US}(\mathsf{AF})$ we have that $\sigma(\mathsf{AF}, U) = \sigma(\mathsf{AF}|_U)$ with $\sigma(\mathsf{AF}, U) = \{E \cap U \mid E \in \sigma(\mathsf{AF})\}$.
- *modularization [16]*, iff for every AF AF we have: $E_1 \in \sigma(\mathsf{AF})$ and $E_2 \in \sigma(\mathsf{AF}^{E_1})$ implies $E_1 \cup E_2 \in \sigma(\mathsf{AF})$.
- *reduct-admissibility [13]*, iff for every AF AF and $E \in \sigma(\mathsf{AF})$, we have that $\forall a \in E$: if b attacks a then $b \notin \bigcup \sigma(\mathsf{AF}^E)$.
- *semi-qualified-admissibility [13]*, iff for every AF AF and $E \in \sigma(\mathsf{AF})$, we have that $\forall a \in E$, if b attacks a and $b \in \bigcup \sigma(\mathsf{AF})$ then $\exists c \in E$ s.t. c attacks b.

The principle of serialisability is intrinsically linked with admissibility since the building blocks of constructed extensions are the initial sets of an argumentation framework. By design, every extension constructed by the transition system for some α and β satisfies admissibility and thus also conflict-freeness. In other words, admissibility and conflict-freeness are necessary criteria for serialisability. Interestingly, the recently introduced principle of modularization [16] is also implied by serialisability.

Two of the more prominent principles from the literature are directionality and SCC-recursiveness. Like serialisability, the SCC-recursiveness principle can also be used to characterise existing semantics or define new semantics [14]. That raises the question if there exists a connection between these principles.

Interestingly, the principles of directionality and serialisability are independent of each other. The same holds true for SCC-recursiveness. While the above mentioned serialisable semantics are all SCC-recursive, the unchallenged semantics, which is investigated further in the following section, is not SCC-recursive. The relevant results are summarised in the following theorem.

Theorem 1. *Let σ be any semantics.*

- *If σ satisfies serialisability then it satisfies conflict-freeness.*
- *If σ satisfies serialisability then it satisfies admissibility.*
- *If σ satisfies serialisability then it satisfies modularization.*
- *Directionality does not imply serialisability and vice versa.*
- *SCC-recursiveness does not imply serialisability and vice versa.*

For all other mentioned principles, we could not find any relationships to serialisability. We will now take a closer look on the principle of directionality.

3.2. A Closer Look on Directionality

We now specify some additional property called $\alpha\beta$*-closure* that allows us to relate serialisability and directionality. This property captures whether or not every path of the transition system for α_σ and β_σ of the semantics σ eventually terminates for all argumentation frameworks $\mathsf{AF} \in \mathfrak{AF}$, i. e., every path leads to some σ-extension of AF.

Definition 12. Let σ be serialisable with α_σ and β_σ. We say that σ is $\alpha\beta$*-closed* for all argumentation frameworks $\mathsf{AF} \in \mathfrak{AF}$ if and only if, for every state (AF', S') with $(\mathsf{AF}, \emptyset) \rightsquigarrow^{\alpha_\sigma} (\mathsf{AF}', S')$ we have that, there exists some $\mathsf{AF}'' \in \mathfrak{AF}$ and some $S'' \subseteq \mathfrak{A}$ such that $(\mathsf{AF}', S') \rightsquigarrow^{\alpha_\sigma, \beta_\sigma} (\mathsf{AF}'', S'')$.

The property of $\alpha\beta$-closure is satisfied by most of the existing serialisable semantics. Only the transition system for the stable semantics does not terminate for all paths. Due to space limitations we do not recall the corresponding selection and termination functions but we refer to [6].

Theorem 2. *The* adm, co, gr, pr *and* sa *semantics are* $\alpha\beta$*-closed, while the* st *semantics is not, wrt. the selection and termination functions defined in [6].*

The fact that stable semantics is not closed wrt. its transition system is no coincidence since it is also the only semantics of the above that is not directional. In fact, if a semantics σ is serialisable and also $\alpha\beta$-closed, then it follows that σ must also be directional.

Theorem 3. *If a semantics σ is serialisable via α_σ and β_σ and is $\alpha_\sigma\beta_\sigma$-closed, then σ satisfies directionality.*

4. Unchallenged Semantics

The notion of serialisability allows to define completely new semantics by defining only a selection and a termination function. One aspect behind the initial sets is that they represent sets of arguments that solve a local conflict. We also have the differentiation between unattacked, unchallenged, and challenged initial sets, essentially distinguishing how convincing these sets solve their local conflict. In general, the grounded semantics can be considered to represent a minimal consensus, i. e., a set of arguments that everyone can agree on. The serialised characterisation of the grounded semantics shows us that this

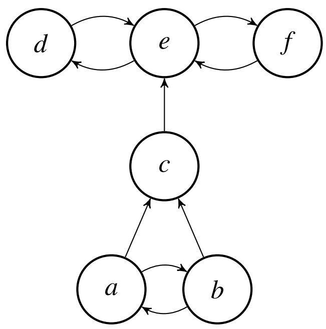

Figure 2. The argumentation framework AF_2 from Example 2.

is achieved by only considering unattacked initial sets in the selection function σ_{gr}. This is formalised by the selection function $\alpha_{gr}(X,Y,Z) = X$ and the termination function

$$\beta_{gr}(\mathsf{AF},S) = \begin{cases} 1 \text{ if } \mathsf{IS}^{\not\leftarrow}(\mathsf{AF}) = \emptyset \\ 0 \text{ otherwise} \end{cases}$$

However, from the perspective of local conflicts, an unchallenged initial set S also resolves its conflict while being uncontested by any acceptable argument. Therefore, it is reasonable to accept the arguments in an unchallenged initial set S as part of a consensus, since there exists no competing acceptable solution to the conflict S is concerned with. A natural approach to address this concern would now be to consider a semantics which allows for unattacked as well as unchallenged initial sets to be selected until no further unattacked or unchallenged initial sets exist. This means, we do not allow challenged initial sets to be included since there is at least one other set of arguments that solves the same local conflict, i. e., there is no consensual solution to this conflict. This approach has already been suggested in [6] but we will now investigate it in-depth. The approach can be implemented by the selection function α_{uc} defined via

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

and the termination function β_{uc} defined via

$$\beta_{uc}(\mathsf{AF},S) = \begin{cases} 1 \text{ if } \mathsf{IS}^{\not\leftarrow}(\mathsf{AF}) \cup \mathsf{IS}^{\not\leftrightarrow}(\mathsf{AF}) = \emptyset \\ 0 \text{ otherwise} \end{cases}$$

Essentially, this approach amounts to exhaustively adding unattacked and unchallenged initial sets. In light of this aspect, we also call the α_{uc},β_{uc}-semantics the *unchallenged semantics* (uc) where $\mathsf{uc}(\mathsf{AF}) = \{E \mid (\mathsf{AF},\emptyset) \rightsquigarrow^{\alpha_{uc},\beta_{uc}} (\mathsf{AF}',E)\}$ denotes the set of *unchallenged extensions*.

Example 2. Consider AF_2 depicted in Figure 2. There are four preferred extensions E_1, E_2, E_3, and E_4 in AF_2 defined via

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

while the grounded and ideal extensions are empty. However, there is one unchallenged extension $E_5 = \{d,f\}$. The reason for that is that both $\{d\}$ and $\{f\}$ are unchallenged initial sets in AF_2 (and once one is selected the other becomes an unattacked initial set of the respective reduct and can be selected as well).

The unchallenged semantics is more skeptical than the preferred semantics but less skeptical than the ideal semantics as has already been observed in [6].

Theorem 4. *For every $E \in uc(\mathsf{AF})$:*

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

Also clear is the following observation:

Proposition 1. *For every AF, $uc(\mathsf{AF}) \neq \emptyset$.*

In Definition 12 we introduced the property of $\alpha\beta$-closure for serialisable semantics. This property is also satisfied by the unchallenged semantics.

Theorem 5. *Unchallenged semantics is $\alpha_{uc}\beta_{uc}$-closed.*

In light of Theorem 3 this directly implies that the unchallenged semantics is directional. In addition to the above characterisation via the selection and termination functions, the unchallenged semantics can also be characterised in a different manner. The following theorem gives a recursive definition of the unchallenged semantics based on the notion of the reduct, but without use of the transition rule.

Theorem 6. *Let $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ be an abstract argumentation framework and $E \subseteq \mathsf{A}$. E is an unchallenged extension if and only if either*

- *$E = \emptyset$ and $IS^{\not\leftarrow} \cup IS^{\not\leftrightarrow}(\mathsf{AF}) = \emptyset$ or*
- *$E = E_1 \cup E_2$, $E_1 \in IS^{\not\leftarrow} \cup IS^{\not\leftrightarrow}(\mathsf{AF})$ and E_2 is an unchallenged extension in AF^{E_1}.*

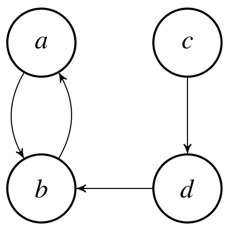

Figure 3. The argumentation framework AF_3 from Example 3.

Example 3. Consider the argumentation framework AF_3 in Figure 3. The initial sets of AF_3 are $\{c\}$ and $\{a\}$. Here, $\{c\}$ is unattacked and $\{a\}$ is considered unchallenged. Therefore, both sets are valid in terms of the selection function α_{uc}. Assume we select the set $\{c\}$, we transition to the framework $\mathsf{AF}_3^{\{c\}} = (\{a,b\},\{(a,b),(b,a)\})$. In this argumentation framework we have two initial sets $\{a\}$ and $\{b\}$. Both of which are challenged by each other. This means, the termination function β_{uc} is satisfied and therefore $\{c\}$ is an unchallenged extension of AF_3.

On the other hand, if we select $\{a\}$ as the first transition, we arrive at the argumentation framework $\mathsf{AF}_3^{\{a\}} = (\{c,d\},\{(c,d)\})$. Here, $\{c\}$ is the only initial set and it is unattacked, just like it was in AF_3 itself. After the transition step, we obtain $\mathsf{AF}_3^{\{a,c\}} = (\emptyset, \emptyset)$, which means we are in a terminal state since we have that $\beta_{uc}(\mathsf{AF}_3^{\{a,c\}}, \{a,c\}) = 1$.

All in all, $\{c\}$ and $\{a,c\}$ are the unchallenged extensions of AF_3.

In the following, we further investigate the compliance of the unchallenged semantics with principles from the literature. The unchallenged semantics satisfies conflict-freeness and admissibility by design. It also satisfies the more recently introduced principle of modularization as well as the reinstatement principle. Furthermore, the unchallenged semantics also satisfies the more complex principle of directionality.

Theorem 7. *The unchallenged semantics satisfies the following principles: Conflict-Freeness, Admissbility, Reduct Admissibility, Semi-Qualified Admissibility, Reinstatement, Directionality, Modularization and Serialisability.*

On the other hand, the unchallenged semantics does not satisfy strong admissibility. Like most admissible-based semantics, it does not satisfy the "allowing abstention" principle. As we have seen in Example 3 the unchallenged semantics does not satisfy I-maximality.

Interestingly, the SCC-recursiveness property is also not satisfied by this semantics. The reason for that stems from the inclusion of unchallenged initial sets. This allows for situations like in Example 3 where an unchallenged initial set can become challenged in some reduct of AF, but it can still be part of the extension if selected in an earlier transition step. Therefore, the unchallenged semantics serves as an example to show that not all serialisable semantics must necessarily be SCC-recursive.

Theorem 8. *The unchallenged semantics does not satisfy the following principles: Strong Admissibility, Naivety, Allowing Abstention, I-Maximality and SCC-Recursiveness.*

5. Computational Complexity

We assume familiarity with basic concepts of computational complexity and basic complexity classes such as P, NP, coNP, see [17] for an introduction. We also require knowledge of the classes Σ_2^P, Π_2^P, and $\mathsf{P}_{\|}^{\mathsf{NP}}$. The class $\Sigma_2^P = \mathsf{NP}^{\mathsf{NP}}$ is the class of decision problems that can be solved in polynomial time by a non-deterministic algorithm that has access to an NP-oracle, i. e., in every step of the algorithm it can immediately obtain the answer to an NP-complete problem. The class $\Pi_2^P = \mathsf{co}\Sigma_2^P = \mathsf{noNP}^{\mathsf{NP}}$ is the complement of Σ_2^P. The class $\mathsf{P}_{\|}^{\mathsf{NP}}$ [18] is the class of decision problems that can be solved by a deterministic polynomial-time algorithm that can make polynomially many *non-adaptive* (or *parallel*) queries to an NP-oracle. Note that $\mathsf{P}_{\|}^{\mathsf{NP}}$ is sometimes denoted by Θ_2^P and is equal to $\mathsf{P}^{\mathsf{NP}[log]}$, i. e., the class of decision problems solvable by a deterministic polynomial-time algorithm that can make logarithmically many *adaptive* NP-oracle calls [17].

We consider the following computational tasks, cf. [19]:

Ver_{uc}	Given $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ and $E \subseteq \mathsf{A}$, decide whether $E \in \mathsf{uc}(\mathsf{AF})$.
$Exists_{\mathsf{uc}}^{\neg\emptyset}$	Given $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$,

	decide whether there is an $E \in \mathsf{uc}(\mathsf{AF})$ with $E \neq \emptyset$.
Skept$_{\mathsf{uc}}$	Given $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ and $a \in \mathsf{A}$, decide whether for all $E \in \mathsf{uc}(\mathsf{AF})$, $a \in E$.
Cred$_{\mathsf{uc}}$	Given $\mathsf{AF} = (\mathsf{A}, \mathsf{R})$ and $a \in \mathsf{A}$, decide whether there is $E \in \mathsf{uc}(\mathsf{AF})$ with $a \in E$.

Note that we do not consider the problem *Exists*$_{\mathsf{uc}}$, which asks whether some unchallenged extension exists, since this problem is trivial due to Proposition 1.

The results of our analysis are as follows.

Theorem 9.

1. *Ver$_{uc}$ is in Σ_2^P and $\mathsf{P}_{\|}^{NP}$-hard.*
2. *Exists$_{uc}^{\neg\emptyset}$ is $\mathsf{P}_{\|}^{NP}$-complete.*
3. *Skept$_{uc}$ is Π_2^P-complete.*
4. *Cred$_{uc}$ is Σ_2^P-complete.*

As can be seen, the exact computational complexity of the verification task is still an open problem (which is a bit surprising since we have exact characterisations for the more "complex" problems). However, all results are in line with our previous observation that unchallenged semantics is somehow "in-between" ideal and preferred semantics, cf. Theorem 4. While most tasks related to ideal semantics are $\mathsf{P}_{\|}^{\mathsf{NP}}$-complete [20], skeptical reasoning with preferred semantics is Π_2^P-complete [21]. But in difference to preferred semantics, both skeptical and credulous reasoning is on the second level of the polynomial hierarchy for unchallenged semantics. As before, the proof of Theorem 9 can be found in the online appendix.[2] While the proofs of items 1 and 2 from Theorem 9 follow quite easily from existing results, in particular from [6], the hardness proofs of items 3 and 4 require quite a different reduction technique as, e. g., the Π_2^P-hardness proof for skeptical reasoning with preferred semantics [21].

6. Summary and Conclusion

We investigated the principle of *serialisability* in-depth, in particular wrt. its relationships to other principles from the literature [11,12,13]. While serialisability implies conflict-freeness, admissibility, and modularization, it is independent of similar principles like directionality and SCC-recursiveness. However, if a serialisable semantics is $\alpha\beta$-closed, it is also directional. We also analysed unchallenged semantics, a specific instance of a serialisable semantics, in terms of satisfied principles and computational complexity. This semantics is $\alpha_{uc}\beta_{uc}$-closed and thus directional. It also satisfies reinstatement, but interestingly it is not SCC-recursive, in contrast to all other serialisable semantics. We have also implemented a general serialisable reasoner as well as reasoners for all existing serialisable semantics[3].

In future work, we intent to further investigate serialisability. That includes defining and analysing completely new semantics with more sophisticated selection and termina-

[2] `http://mthimm.de/misc/lbmt_uncsem_proofs.pdf`

[3] Link to implementation: `https://tinyurl.com/serialisableReasoner`

tion functions. We will also consider applying the concept of serialisability to other types of semantics such as naive- or weak-admissible-based semantics. Regarding the unchallenged semantics, the question of whether there exists a non-recursive characterisation is also subject to future work.

Acknowledgements The research reported here was partially supported by the Deutsche Forschungsgemeinschaft (grant 375588274).

References

[1] Dung PM. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. Artificial Intelligence. 1995;77(2):321-58.

[2] Baroni P, Caminada M, Giacomin M. Abstract Argumentation Frameworks and Their Semantics. In: Baroni P, Gabbay D, Giacomin M, van der Torre L, editors. Handbook of Formal Argumentation. College Publications; 2018. p. 159-236.

[3] Caminada MWA, Gabbay DM. A Logical Account of Formal Argumentation. Studia Logica. 2009;93(2–3):109-45.

[4] Bonzon E, Delobelle J, Konieczny S, Maudet N. A Comparative Study of Ranking-based Semantics for Abstract Argumentation. In: Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI'16); 2016. p. 914-20.

[5] Hunter A, Thimm M. Probabilistic Reasoning with Abstract Argumentation Frameworks. Journal of Artificial Intelligence Research. 2017 August;59:565-611.

[6] Thimm M. Revisiting initial sets in abstract argumentation. Argument & Computation. 2022 July.

[7] Xu Y, Cayrol C. Initial Sets in Abstract Argumentation Frameworks. In: Proceedings of the 1st Chinese Conference on Logic and Argumentation (CLAR'16). vol. 1811; 2016. p. 72-85.

[8] Baumann R, Brewka G, Ulbricht M. Revisiting the foundations of abstract argumentation–semantics based on weak admissibility and weak defense. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34; 2020. p. 2742-9.

[9] Amgoud L, Prade H. Using arguments for making and explaining decisions. Artificial Intelligence. 2009 March;173(3-4):413-36.

[10] Rago A, Cocarascu O, Bechlivanidis C, Lagnado DA, Toni F. Argumentative explanations for interactive recommendations. Artificial Intelligence. 2021;296:103506.

[11] Baroni P, Giacomin M. On principle-based evaluation of extension-based argumentation semantics. In: Artificial Intelligence. vol. 171. Elsevier; 2007. p. 675-700.

[12] van der Torre L, Vesic S. The Principle-Based Approach to Abstract Argumentation Semantics. In: Handbook of formal argumentation, Vol. 1. College Publications; 2018. p. 2735-78.

[13] Dauphin J, Rienstra T, van der Torre L. A Principle-Based Analysis of Weakly Admissible Semantics. In: Proceedings of COMMA 2020; 2020. p. 167-78.

[14] Baroni P, Giacomin M, Guida G. SCC-recursiveness: a general schema for argumentation semantics. Artificial Intelligence. 2005;168(1–2):162-210.

[15] Baroni P, Caminada M, Giacomin M. An introduction to argumentation semantics. The knowledge engineering review. 2011;26(4):365-410.

[16] Baumann R, Brewka G, Ulbricht M. Comparing Weak Admissibility Semantics to their Dung-style Counterparts–Reduct, Modularization, and Strong Equivalence in Abstract Argumentation. In: Proceedings of KR 2020; 2020. p. 79-88.

[17] Papadimitriou C. Computational Complexity. Addison-Wesley; 1994.

[18] Eiter T, Gottlob G. The Complexity Class Theta2p: Recent Results and Applications in AI and Modal Logic. In: Proceedings of FCT'97. Berlin, Heidelberg; 1997. p. 1-18.

[19] Dvorák W, Dunne PE. Computational problems in formal argumentation and their complexity. In: Handbook of formal argumentation; 2018. p. 631-87.

[20] Dunne PE. The computational complexity of ideal semantics. Artificial Intelligence. 2009 December;173(18):1559-91.

[21] Dunne PE, Bench-Capon TJM. Coherence in finite argument systems. Artificial Intelligence. 2002;141(1/2):187-203. Available from: `https://doi.org/10.1016/S0004-3702(02)00261-8`.