

INSTITUT FÜR INFORMATIONSSYSTEME
ABTEILUNG DATENBANKEN UND ARTIFICIAL INTELLIGENCE

On Rejected Arguments and Implicit Conflicts: The Hidden Power of Argumentation Semantics

DBAI-TR-2016-102

**Ringo Baumann
Thomas Linsbichler
Hannes Strass**

**Wolfgang Dvořák
Christof Spanring
Stefan Woltran**

Institut für Informationssysteme
Abteilung Datenbanken und
Artificial Intelligence
Technische Universität Wien
Favoritenstr. 9
A-1040 Vienna, Austria
Tel: +43-1-58801-18403
Fax: +43-1-58801-18493
sekret@dbai.tuwien.ac.at
www.dbai.tuwien.ac.at

DBAI TECHNICAL REPORT
2016



TECHNISCHE
UNIVERSITÄT
WIEN
Vienna University of Technology

On Rejected Arguments and Implicit Conflicts: The Hidden Power of Argumentation Semantics

Ringo Baumann¹ Wolfgang Dvořák²
Thomas Linsbichler³ Christof Spanring⁴ Hannes Strass⁵
Stefan Woltran⁶

Abstract. Abstract argumentation frameworks (AFs) are one of the most studied formalisms in AI and are formally simple tools to model arguments and their conflicts. The evaluation of an AF yields extensions (with respect to a semantics) representing alternative acceptable sets of arguments. For many of the available semantics two effects can be observed: there exist arguments in the given AF that do not appear in any extension (rejected arguments); there exist pairs of arguments that do not occur jointly in any extension, albeit there is no explicit conflict between them in the given AF (implicit conflicts). In this paper, we investigate the question whether these situations are only a side-effect of particular AFs, or whether rejected arguments and implicit conflicts contribute to the expressiveness of the actual semantics. We do so by introducing two subclasses of AFs, namely *compact* and *analytic* frameworks. The former class contains AFs that do not contain rejected arguments with respect to a semantics at hand; AFs from the latter class are free of implicit conflicts for a given semantics. Frameworks that are contained in both classes would be natural candidates towards normal forms for AFs since they minimize the number of arguments on the one hand, and on the other hand maximize the information on conflicts, a fact that might help argumentation systems to evaluate AFs more efficiently. Our main results show that under stable, preferred, semi-stable, and stage semantics neither of the classes is able to capture the full expressive power of these semantics; we thus also refute a recent conjecture by Baumann *et al.* on implicit conflicts. Moreover, we give a detailed complexity analysis for the problem of deciding whether an AF is compact, resp. analytic. Finally, we also study the signature of these subclasses for the mentioned semantics and shed light on the question under which circumstances an arbitrary framework can be transformed into an equivalent compact, resp. analytic, AF.

¹Leipzig University, Germany. baumann@informatik.uni-leipzig.de

²University of Vienna, Austria. wolfgang.dvorak@univie.ac.at

³TU Wien, Austria. linsbich@dbai.tuwien.ac.at

⁴TU Wien, Austria and University of Liverpool, UK. spanring@dbai.tuwien.ac.at

⁵Leipzig University, Germany. strass@informatik.uni-leipzig.de

⁶TU Wien, Austria. woltran@dbai.tuwien.ac.at

Acknowledgements: The authors are grateful to the anonymous reviewers of the preceding papers and the current article for their detailed reviews and helpful comments that led to an improved presentation. Furthermore, the authors are grateful to Paul Dunne for helpful discussions. This research has been supported by the German Research Foundation (DFG) under project BR 1817/7-1 and the Austrian Science Fund (FWF) under projects I1102, I2854, and P25518.

Copyright © 2016 by the authors

1 Introduction

In recent years argumentation has emerged to become one of the major fields of research in Artificial Intelligence [34, 11]. In particular, Dung’s well-studied abstract argumentation frameworks (AFs) [18] are a simple, yet powerful formalism for modeling and deciding argumentation problems that are integral to many advanced argumentation systems, see e.g. [12]. The evaluation of AFs in terms of finding reasonable positions with respect to a given framework is defined via so-called argumentation semantics (cf. [1] for an overview). Given an AF F , an argumentation semantics σ returns acceptable sets of arguments $\sigma(F)$, the (σ -)extensions of F . Several semantics have been introduced over the years [18, 39, 13, 2] with motivations ranging from the desired treatment of specific examples to fulfilling certain abstract principles. One important line of research in abstract argumentation is thus the systematic comparison of the different semantics available. Hereby, the behavior of extensions with respect to certain properties [4] has been analyzed and the expressive power of semantics [23, 26, 28, 36] has been studied by identifying the set of extension-sets achievable under certain semantics. On the other hand, subclasses of AFs such as acyclic, symmetric, odd-cycle-free or bipartite AFs, have been considered, since for some of these classes different semantics collapse [14, 19]. Beside these subclasses based on the graph structure there are also classes defined via properties of extensions. The probably most prominent such subclass is the class of coherent AFs [21], i.e. AFs where the stable and preferred semantics coincide. Further examples for subclasses that are defined via extensions can be found in [5, 25].

In this work we contribute to both streams of research. We introduce two new classes, which to the best of our knowledge have not received attention in the literature. The actual definition of these two classes is motivated by typical phenomena that can be observed for several semantics. First, there exist arguments in a given AF that do not appear in any extension. Since these so-called *rejected* arguments do not appear in the result of extension-based semantics, it is a natural question whether such arguments can be “removed” from the AF at hand without changing its outcome (in a certain way, this question is similar to the problem of simplifying propositional formulas by removing “don’t care” atoms). In order to have a handle for analyzing the effect of rejected arguments, we introduce the class of *compact* AFs: an AF is *compact* (with respect to a semantics σ), if each of its arguments appears in at least one σ -extension. Second, we are interested in the concept of *implicit conflicts*. An attack between two arguments represents an explicit conflict. By the nature of most argumentation semantics, conflicts can however also be implicit in the sense that some arguments do not occur together in any extension, although there is no attack between them. In order to understand the expressive power of implicit conflicts we introduce the class of *analytic* frameworks. Given a semantics σ , if every conflict between two arguments a, b in an AF F is explicit (i.e., for all arguments a, b , if $\{a, b\} \cap E \neq \{a, b\}$ for all σ -extensions E , then a attacks b in F or b attacks a in F) then F is called analytic. Both compact and analytic AFs thus yield a “semantic” subclass since their definitions rely on the actual extensions obtained via the chosen semantics.

The role of rejected arguments Although rejected arguments are natural ingredients in typical argumentation scenarios, it is of interest to understand in which ways rejected arguments contribute




Figure 1: Rejected argument x cannot be removed without changing the stable extensions.

to the “strength” of a particular semantics. Let us first have a brief look on the naive semantics, which is defined as subset-maximal conflict-free sets: Here, it is rather easy to see that any AF can be transformed into an equivalent compact AF by just removing all self-attacking arguments. In other words, the same outcome (in terms of the naive extensions) can be achieved by a simplified AF without rejected arguments. On the one hand, this can be seen as a general weakness of naive semantics, since any possible outcome can be equivalently achieved in the absence of rejected arguments. On the other hand, this shows that towards evaluating an AF under naive semantics, the transformation into a compact AF can provide a beneficial pre-processing step for computing the extensions (which afterwards should however be interpreted in terms of the original AF).

How is the situation with semantics that are considered more mature? We borrow an example from (author?) [22]. Consider the AF F_1 in Figure 1, where nodes represent arguments and directed edges represent attacks.

The *stable* extensions (conflict-free sets attacking all other arguments) of F_1 are given by the set $\mathbb{S} = \{\{a, b, c\}, \{a, b', c'\}, \{a', b, c'\}, \{a', b', c\}, \{a, b, c'\}, \{a', b, c\}, \{a, b', c\}\}$. Observe that x is rejected, i.e. x does not appear in any stable extension of F_1 . Hence, this framework is not compact for the stable semantics. Moreover, it was shown in [22] that there is no compact AF (in this case an AF not using argument x) that yields the same stable extensions as F_1 . By the necessity of conflict-freeness any such compact AF would only allow conflicts between arguments a and a' , b and b' , and c and c' , respectively. Moreover, there would have to be attacks in both directions for each of these conflicts in order to ensure stability. Hence any compact AF having the same stable extensions as F_1 necessarily yields $\{a', b', c'\}$ in addition. In other words, under the stable semantics particular outcomes (in the example the set \mathbb{S} of extensions) can only be achieved via AFs containing at least one rejected argument. Thus, the stable semantics makes proper use of rejected arguments. As we will see, all semantics under consideration (except naive semantics) show a similar behaviour.

The role of implicit conflicts As introduced earlier, implicit conflicts arise when two arguments are never jointly accepted although they do not attack each other. The AF F_2 in Figure 2 provides a simple example for this effect.

It can be seen that stable semantics yields two extensions $\{a, d\}$ and $\{b, c\}$ for F_2 . Since c and d do not occur together in an extension there is an implicit conflict and thus F_2 is not analytic (for stable semantics). However, the naive extensions of F_2 are given by $\{a, d\}, \{b, c\}, \{c, d\}$. Thus c and d are not in an implicit conflict here, and the AF is easily seen to be analytic for naive semantics.




Figure 2: AF illustrating an implicit conflict between c and d for stable semantics.

Indeed, by definition of naive semantics, two arguments occur together in a naive extension if and only if there is no attack between them and they are not self-attacking. Thus not every AF is analytic for naive semantics, but it is quite easy to see that every AF can be turned into an equivalent analytic one over the same arguments, by just connecting the self-attacking arguments to any other argument. Coming back to our example and to stable semantics, the question remains whether F_2 can be turned into an equivalent analytic one? This is quite an easy exercise. Just add an attack from c to d , or likewise from d to c . In fact, this addition does not change the set of extensions. However, it has been left as open question in [7] (stated as “Explicit Conflict Conjecture”) whether such a manipulation of an AF is always possible. In this work, we shall negatively answer this question showing that (i) for preferred and semi-stable semantics, there exist AFs such that there is no equivalent analytic AF; and (ii) for stable and stage semantics, there exist AFs such that there is no equivalent analytic AF, unless we are allowed to add rejected arguments.

Expressiveness of compact and analytic argumentation frameworks Before giving an overview of the obtained results, let us further illustrate some issues that come along with the subclasses of compact and analytic argumentation frameworks. One natural question is whether any AF F can be transformed to an equivalent AF G , i.e. $\sigma(F) = \sigma(G)$ for a given semantics σ , that is compact or analytic. In case the answer is no, we can conclude that the full range of expressiveness of σ indeed relies on the concepts of rejected arguments and implicit conflicts. Knowing which sets of extensions a semantics is able to express is of central interest in approaches of extension-based revision of AFs [16]. As the result of the revision may also be subject to certain syntactic constraints (e.g. a fixed set of arguments [15]) it is important to know about the role of rejected arguments and implicit conflicts. For instance, a revised AF might be required (e.g. in order to fulfill revision postulates) to have exactly the extension-set \mathbb{S} from above under stable semantics while consisting solely of the arguments $\{a, b, c, a', b', c'\}$. As we have already observed, and we will show in a more comprehensive and general manner in the paper, such a revision is not possible since getting \mathbb{S} under stable semantics would require an additional, rejected argument.

Implications for argumentation systems An even more promising application of our results lies in the field of concrete software systems for computing semantics of abstract argumentation frameworks. A considerable number of such systems (“solvers”) exist, as has been witnessed by the First International Competition on Computational Models of Argument (ICCMA 2015) [38].¹ Using instances from that competition and additional instances created according to the same graph model as the competition instances, we also performed an experimental evaluation on the theoret-

¹A total number of 18 solvers participated, see <http://argumentationcompetition.org>.

ical phenomena we study in this paper. The results can be found in A, and demonstrate the clear computational benefit of knowing about implicit conflicts in an argumentation framework. More precisely, once all implicit conflicts of an AF are made explicit, then the competition winners are able to compute the AF’s extensions (for stable and preferred semantics) much faster than before (without implicit conflicts made explicit). Thus it is a naturally arising research question whether information about implicit conflicts can be obtained “cheaply” in terms of computational cost, a question that we will also address in the paper. For knowing about rejected arguments, the computational gain is immediately clear, since the lower the number of arguments, the smaller is the search space a solver has to go through in order to find all extensions. Thus, preprocessing steps that remove rejected arguments might also be beneficial to solving runtime. Moreover, if an AF has no rejected arguments then all of its arguments are contained in at least one extension, and so credulous as well as skeptical reasoning become easy tasks [7].

Overall, the research question we are interested in is: how computationally costly is it to determine whether an AF can be simplified along the dimensions *rejected arguments* and *implicit conflicts*? Answering this question would be crucial towards the development of clever methods for preprocessing AFs before solving. However, more fundamental questions need to be addressed first. On the one hand, we analyse how hard it is to decide whether an AF is compact (resp. analytic); on the other hand, we ask whether any AF can be transformed into an AF that is compact (resp. analytic) and equivalent under a particular semantics. Unfortunately, the answers to both of these questions is in a certain sense negative for all of the semantics we consider: intuitively speaking, our complexity results will show that deciding whether simplification is applicable (having certain reasoning tasks in mind) is as expensive as solving the reasoning tasks themselves. Furthermore, we can even show that there are AFs that cannot be exhaustively simplified. (More formally, there are AFs that have “pathological” implicit conflicts that cannot be made explicit even if we allow arbitrary semantics-preserving changes in other parts of the AF.) This does not make our results less applicable to implementation of reasoning systems, however. These negative results help the solver development community to delineate what can and cannot be done in improving solver performance by intelligent preprocessing. That is, by our results, we know that computing *all* rejected arguments and implicit conflicts are not viable candidates for simplifying given argumentation frameworks.

Main contributions & structure of the paper The main contributions of this article are organized as follows. Recall that the semantics we mainly investigate are stable, preferred, semi-stable, stage, and naive semantics.

- In Section 3 we formally introduce the subclasses of compact and analytic AFs with respect to the considered semantics and investigate their relationship. For both classes the picture is similar: for instance, if an AF is compact (resp. analytic) for stable it also is for semi-stable (preferred, stage, and naive); but the other direction does not hold in general.
- Section 4 answers the question how hard it is to decide whether an AF is compact (resp. analytic). As it turns out, the complexity of this problem for a given semantics σ is the

same as credulous acceptance under σ . Thus, we obtain tractability for naive semantics, NP-completeness for stable and preferred semantics, and Σ_2^P -completeness for semi-stable and stage semantics.

- In Section 5 we refute the Explicit Conflict Conjecture [7] for σ being among stable, preferred, semi-stable and stage semantics. In fact, we provide AFs such that there is no AF equivalent under σ that contains solely explicit conflicts. On the other hand, we identify sufficient conditions guaranteeing equivalence-preserving translations to analytic AFs.
- The final collection of results in Section 6 is concerned with signatures for compact and analytic frameworks. Signatures as introduced in [23] plainly collect all possible sets of extensions AFs can deliver under a given semantics. For instance, it is shown in [23] that preferred and semi-stable semantics yield an equal signature Σ , while the signature of stage semantics is a proper subset of Σ . Compared to [23], we do not give exact characterizations of signatures for compact (resp. analytic) frameworks, but obtain a full picture of their relationship with respect to the different semantics. For instance, we show that in terms of compact AFs, the signatures for semi-stable and preferred semantics become incomparable, while for analytic AFs, the signature for semi-stable semantics is a proper subset of the signature for preferred semantics. Finally, we generalize some recent results on maximal numbers of extensions [9] to give some impossibility results for *compact realizability*.

In this work we consider several rather complex examples of argumentation frameworks, whose evaluation is a non-trivial task. Thus, for the reader's convenience, we provide encodings in the *.apx format*, which can be used to evaluate the AFs with systems like ASPARTIX [29]². These encodings can be either downloaded from <http://www.dbai.tuwien.ac.at/proj/adf/HiddenPowerAFs.zip> or directly accessed by clicking at the corresponding figure. (Depending on the actual pdf viewer, a right or double-click should initiate saving.)

This article is based on [7] and [32], but also contains several new results.

2 Preliminaries

In what follows, we briefly recall the necessary background on abstract argumentation and computational complexity. For an excellent overview on abstract argumentation and in particular on argumentation semantics, we refer to [1].

Abstract Argumentation

Throughout the paper we assume a countably infinite domain \mathfrak{A} of arguments. An argumentation framework (AF) is a pair $F = (A, R)$ where $A \subseteq \mathfrak{A}$ is a finite set of arguments and $R \subseteq A \times A$ is the attack relation. The collection of all AFs is given as $AF_{\mathfrak{A}}$. For an AF $F = (B, S)$ we use A_F and R_F to refer to B and S , respectively. We write $a \succ_F b$ for $(a, b) \in R_F$ and $S \succ_F a$

²A web front-end is available at <http://rull.dbai.tuwien.ac.at:8080/ASPARTIX/>.

(resp. $a \succrightarrow_F S$) if there exists some $s \in S$ such that $s \succrightarrow_F a$ (resp. $a \succrightarrow_F s$). Symmetric attacks $\{(a, b), (b, a)\} \subseteq R_F$ are occasionally denoted by $\langle a, b \rangle \in R_F$. For $S \subseteq A$, the *range* of S (w.r.t. F), denoted S_F^+ , is the set $S \cup \{b \mid S \succrightarrow_F b\}$. Moreover, $F|_S$ denotes the AF $(A_F \cap S, R \cap (S \times S))$.

Given an AF F , an argument $a \in A_F$ is *defended* (in F) by $S \subseteq A_F$ if for each $b \in A_F$, such that $b \succrightarrow_F a$, also $S \succrightarrow_F b$. A set T of arguments is defended (in F) by S if each $a \in T$ is defended by S (in F). A set $S \subseteq A_F$ is *conflict-free* (in F), if there are no arguments $a, b \in S$, such that $\langle a, b \rangle \in R_F$. We denote the set of all conflict-free sets in F as $cf(F)$. $S \in cf(F)$ is called *admissible* (in F) if S defends itself. We denote the set of admissible sets in F as $adm(F)$.

The terms semantics and extension are often used almost synonymously. Formally a semantics is a mapping, while extensions are concrete elements of its image. The semantics we study in this work are those characterized by the naive, stable, preferred, stage, and semi-stable extensions. Given an AF F they are defined as subsets of $cf(F)$ as follows:

- $S \in nai(F)$, if $\nexists T \in cf(F)$ with $T \supset S$;
- $S \in stb(F)$, if $S_F^+ = A_F$;
- $S \in prf(F)$, if $S \in adm(F)$ and $\nexists T \in adm(F)$ with $T \supset S$;
- $S \in stg(F)$, if $\nexists T \in cf(F)$ with $T_F^+ \supset S_F^+$;
- $S \in sem(F)$, if $S \in adm(F)$ and $\nexists T \in adm(F)$ with $T_F^+ \supset S_F^+$.

The following relations between these semantics are well-known to hold for any AF F :

$$stb(F) \subseteq sem(F) \subseteq prf(F)$$

$$stb(F) \subseteq stg(F) \subseteq nai(F)$$

Furthermore, apart from stable semantics all considered semantics guarantee the existence of at least one (possibly empty) extension as long as finite AFs are considered (cf. [8] for a detailed overview including the infinite case).

We will also make frequent use of the following concepts.

Definition 1. Given $\mathbb{S} \subseteq 2^{\mathcal{A}}$, $Args_{\mathbb{S}}$ denotes $\bigcup_{S \in \mathbb{S}} S$ and $Pairs_{\mathbb{S}}$ denotes $\{(a, b) \mid \exists S \in \mathbb{S} : \{a, b\} \subseteq S\}$. \mathbb{S} is called an *extension-set* (over \mathcal{A}) if $Args_{\mathbb{S}}$ is finite.

In words, $Args_{\mathbb{S}}$ stands for all arguments occurring in some element of \mathbb{S} and $Pairs_{\mathbb{S}}$ for all pairs of arguments occurring together in some element of \mathbb{S} . As is easily observed, for all semantics σ , $\sigma(F)$ is an extension-set for any AF F .

Example 1. Consider the AF F depicted in Figure 3. We have that a' , b' and f' are self-attacking, since all semantics considered build upon conflict-freeness these three arguments can thus not be included in any extension. Similarly we may accept only one argument from a and c as these two are mutually attacking each other; the same holds for b and d and also for c and d . Naive semantics generates maximal conflict-free sets, we thus get $nai(F) = \{\{a, b, e\}, \{a, d, e\}, \{b, c, e\}\}$.




Figure 3: Argumentation framework F used in Example 1.

Argument e is contained in each naive extension as it does not share any attacks with not self-attacking arguments. However e is attacked by f and can not defend itself against this attack. Thus the set $\{a, b, e\}$ is not admissible. Preferred semantics, as maximal admissible sets, then computes to $prf(F) = \{\{a, b\}, \{a, d, e\}, \{b, c, e\}\}$.

Now for stable extensions we need conflict-freeness as well as a partition of all arguments into accepted or attacked. Naturally this means that only maximal conflict-free sets are candidates. However neither of the naive sets has full range: $\{a, b, e\}$ does not attack f , $\{a, d, e\}$ does not attack b' and $\{b, c, e\}$ does not attack a' . Thus there is no stable extension, i.e. $stb(F) = \emptyset$.

Stage semantics can be seen as a less restrictive form of stable semantics in that we do not need to cover all arguments in range but want extensions to be conflict-free and range-maximal. As emphasized above all naive extensions have incomparable range (missing f , b' , or resp. a') and thus $stg(F) = nai(F)$. Similarly semi-stable extensions are those admissible sets that are range-maximal. And in this case also the preferred extensions all have incomparable range (missing f and e , b' , or resp. a') and thus $sem(F) = prf(F)$.

Now as for the concepts introduced in Definition 1 we have $Args$ which are all the arguments occurring in any extension; in this case for all semantics $\sigma \in \{nai, stg, sem, prf\}$ we get $Args_{\sigma(F)} = \{a, b, c, d, e\}$. And we have $Pairs$, all pairs of arguments that occur together in any extension; in this case as can easily be checked again for all semantics $\sigma \in \{naive, stg, sem, prf\}$ we get $Pairs_{\sigma(F)} = \{(a, b), (b, a), (a, e), (e, a), (b, e), (e, b), (a, d), (d, a), (d, e), (e, d), (b, c), (c, b), (c, e), (e, c), (a, a), (b, b), (c, c), (d, d), (e, e)\}$. \diamond

Computational Complexity

We assume the reader is familiar with standard complexity concepts, such as P, NP and completeness. Nevertheless we briefly recapitulate the concept of NP-oracle machines and the related complexity class Σ_2^P . By an NP-oracle machine we mean a Turing machine that can access an oracle that decides a given sub-problem from NP (or coNP) within one step. The class Σ_2^P (sometimes also denoted by NP^{NP}), contains the problems that can be decided in polynomial time by a nondeterministic NP-oracle machine.

The known complexity results for the argumentation semantics under consideration are summarized in Table 1 [13, 17, 19, 21, 27]. Here, Ver_σ refers to the problem of verifying that a given set is an extension of a given arbitrary AF F w.r.t. the semantics σ ; $Cred_\sigma$ refers to the problem of verifying that a given argument x is credulously accepted w.r.t. σ in F (there is at least one σ -

Table 1: Complexity of decision problems (\mathcal{C} -c denotes completeness for class \mathcal{C}).

	Ver_σ	$Cred_\sigma$	$Skept_\sigma$
<i>nai</i>	in P	in P	in P
<i>stb</i>	in P	NP-c	coNP-c
<i>adm</i>	in P	NP-c	trivial
<i>prf</i>	coNP-c	NP-c	Π_2^P -c
<i>stg</i>	coNP-c	Σ_2^P -c	Π_2^P -c
<i>sem</i>	coNP-c	Σ_2^P -c	Π_2^P -c

extension of F containing x); and $Skept_\sigma$ refers to the problem of verifying that a given argument x is skeptically accepted w.r.t. σ in F (x is contained in each σ -extension of F). For a more detailed discussion of the complexity results the interested reader is referred to [20, 24]. We only mention that the hardness results still hold if restricted to frameworks without self-attacking arguments, which we will make use of later on.

Later, for semantics σ , we will also need upper bounds for the problem $Cred_\sigma^2$ defined as follows: given AF F and arguments a and b , does there exist an extension $S \in \sigma(F)$ such that $\{a, b\} \subseteq S$ (see e.g. [19]). For the semantics under consideration, it is rather straightforward to see that membership for $Cred_\sigma$ carries over to $Cred_\sigma^2$. For $\sigma \in \{prf, stb, sem, stg\}$ this is witnessed by the standard NP-algorithm of guessing a set S containing a and b and apply an oracle for verifying whether S is a σ -extension. The complexity of the verification problem then yields the desired upper bound. Membership in P for the naive semantics can be decided by just checking whether a, b are neither self-attacking nor attacking each other. Indeed, in this case $\{a, b\}$ is conflict-free in the given AF F , and thus there must exist a naive extension of F containing both a and b .

3 Subclasses of Argumentation Frameworks

In this section, we formally introduce the two central subclasses of argumentation frameworks of this paper, namely compact and analytic frameworks. We study basic properties and relations within the classes first. At the end of the section we will compare the two classes.

3.1 Compact Argumentation Frameworks

The main idea behind compact argumentation frameworks is the absence of rejected arguments (w.r.t. a given semantics).

Definition 2. Given a semantics σ , an AF F is called *compact* for σ (or σ -compact) if $Args_{\sigma(F)} = A_F$. The set of all *compact argumentation frameworks* for σ is denoted by CAF_σ .




Figure 4: AF discussed in Example 2, which is *prf*-compact but neither *sem*-compact nor *stg*-compact.

Example 2. Let us consider the AF F depicted in Figure 4.³ The preferred extensions of F are $\text{prf}(F) = \{\{z\}, \{x_1, a_1\}, \{x_2, a_2\}, \{x_3, a_3\}, \{y_1, b_1\}, \{y_2, b_2\}, \{y_3, b_3\}\}$, meaning that F is *prf*-compact ($F \in \text{CAF}_{\text{prf}}$) since each argument occurs in at least one preferred extension. On the other hand observe that $\text{sem}(F) = \text{prf}(F) \setminus \{\{z\}\}$ and $\text{stg}(F) = \{\{x_i, a_i, b_j\}, \{y_i, b_i, a_j\} \mid 1 \leq i, j \leq 3\}$, i.e. z is not contained in any semi-stable or stage extension. Therefore F is neither compact for semi-stable nor compact for stage semantics (i.e. $F \notin \text{CAF}_{\text{sem}}$ and $F \notin \text{CAF}_{\text{stg}}$). \diamond

As indicated by Example 2, the contents of CAF_σ differ with respect to the semantics σ . Concerning relations between the classes of compact AFs we start with an easy observation. In the following result, the only requirement on a semantics σ is that extensions are subsets of the arguments in the framework, i.e. $\text{Args}_{\sigma(F)} \subseteq A_F$ for any AF F .

Lemma 1. *For any two semantics σ and θ such that for each AF F , for every $S \in \sigma(F)$ there is some $S' \in \theta(F)$ with $S \subseteq S'$, we have $\text{CAF}_\sigma \subseteq \text{CAF}_\theta$.*

Proof. Suppose $F \in \text{CAF}_\sigma$. By definition, $\text{Args}_{\sigma(F)} = A_F$. Now if for each $S \in \sigma(F)$ there is some $S' \in \theta(F)$ with $S \subseteq S'$, we have $\text{Args}_{\sigma(F)} \subseteq \text{Args}_{\theta(F)}$. Since $\text{Args}_{\theta(F)} \subseteq A_F$ by definition, $\text{Args}_{\theta(F)} = A_F$ follows. Hence, $F \in \text{CAF}_\theta$. \square

Note that the case where $\sigma(F) \subseteq \theta(F)$ holds for each AF F is a special case of the premise of Lemma 1. The next result provides a full picture of the relations between classes of compact AFs for the semantics we consider (see also Figure 5).

Theorem 2. *The following relations hold:*

1. $\text{CAF}_{\text{stb}} \subset \text{CAF}_\sigma \subset \text{CAF}_{\text{nai}}$ for $\sigma \in \{\text{prf}, \text{sem}, \text{stg}\}$;
2. $\text{CAF}_{\text{sem}} \subset \text{CAF}_{\text{prf}}$;
3. $\text{CAF}_{\text{stg}} \not\subseteq \text{CAF}_\theta$ and $\text{CAF}_\theta \not\subseteq \text{CAF}_{\text{stg}}$ for $\theta \in \{\text{prf}, \text{sem}\}$.

Proof. (1) Let $\sigma \in \{\text{prf}, \text{sem}, \text{stg}\}$. The \subseteq -relations are due to Lemma 1 together with following facts: (a) in any AF F , $\text{stb}(F) \subseteq \sigma(F)$; (b) each σ -extension E of an AF F is conflict-free in F , thus there exists a naive extension E' of F with $E \subseteq E'$.

$CAF_\sigma \subset CAF_{nai}$: The AF $(\{a, b\}, \{(a, b)\})$ is compact for naive semantics but not for σ .

$CAF_{stb} \subset CAF_\sigma$: Consider AF F from Figure 6a. We have $prf(F) = sem(F) = \{\{x_1, a_1\}, \{x_2, a_2\}, \{x_3, a_3\}, \{y_1, b_1\}, \{y_2, b_2\}, \{y_3, b_3\}\}$, and each of these extensions can be extended to a stage extension (the former three by adding one of the arguments b_1, b_2, b_3 the latter three by adding one of the arguments a_1, a_2, a_3), but $stb(F) = \emptyset$. Thus $A_F = Args_\sigma(F) \neq Args_{stb}(F) = \emptyset$, meaning that $F \in CAF_\sigma$ but $F \notin CAF_{stb}$.

(2) $CAF_{sem} \subseteq CAF_{prf}$ is by the fact that, in any AF F , $sem(F) \subseteq prf(F)$ (cf. Lemma 1). Properness is by the AF in Figure 4, which is (as discussed in Example 2) prf -compact but not sem -compact.

(3) First we show $CAF_{stg} \not\subseteq CAF_\theta$ for $\theta \in \{prf, sem\}$. To this end, consider the simple AF $F' = (\{a, b, c\}, \{(a, b), (b, c), (c, a)\})$. We have $stg(F') = \{\{a\}, \{b\}, \{c\}\}$, thus $F' \in CAF_{stg}$. On the other hand, $sem(F') = prf(F') = \{\emptyset\}$, thus $F' \notin CAF_\theta$.

$CAF_{prf} \not\subseteq CAF_{stg}$ follows by the observations in Example 2.

$CAF_{sem} \not\subseteq CAF_{stg}$: Consider the AF F'' in Figure 6b. One can check that this AF is sem -compact, but not stg -compact. In fact, argument a does not occur in any stage extension. Although $\{a, u_1, x_5\}, \{a, u_2, x_6\}, \{a, u_3, x_7\} \in sem(F'')$, the range of any conflict-free set containing a is a proper subset of the range of every stage extension of F'' : $stg(F'') = \{\{c, u_i, x_4\} \mid i \in \{1, 2, 3\}\} \cup \{\{b, u_i, s_j, x_{i+4}\} \mid i, j \in \{1, 2, 3\}\} \cup \{\{t_i, u_j, s_i, x_i\} \mid i, j \in \{1, 2, 3\}\}$. Hence $CAF_{sem} \not\subseteq CAF_{stg}$. \square

Finally note that every symmetric and irreflexive (i.e. no self-attacking arguments) AF is contained in CAF_{stb} , as already observed in [14, Proposition 6], and therefore also in each CAF_σ for all semantics σ under consideration. But already CAF_{stb} contains strictly more AFs than the class of symmetric and irreflexive AFs, which is, for instance, indicated by the AF $(\{a, b, c, d\}, \{(a, b), (b, c), (c, d), (d, a)\})$, which is clearly not symmetric but compact for the stable semantics. On the other hand observe that $CAF_{nai} \subset AF_{\neq}$, as every AF having self-attacking arguments is not contained in CAF_{nai} .

3.2 Analytic Argumentation Frameworks

In this section we deal with AFS containing no implicit conflicts, which we will call analytic. We differentiate between the concept of an attack (as a syntactical element) and the concept of a conflict (with respect to the evaluation under a given semantics).

³The construct in the lower part of the figure represents symmetric attacks between each pair of distinct arguments. We will make use of this style in illustrations throughout the paper.




Figure 5: Relations between classes of compact AFS (cf. Theorem 2).





Figure 6: AF used in the proof of Theorem 2 to show the incomparability of certain classes of compact AFs.

Definition 3. Given some AF $F = (A, R)$, a semantics σ and arguments $a, b \in A$. If $(a, b) \notin Pairs_{\sigma(F)}$, we say that a and b are in *conflict* in F for σ . If $(a, b) \in R$ or $(b, a) \in R$ we say that the conflict between a and b is *explicit*, otherwise the conflict is called *implicit* (with respect to σ).

Notice that Definition 3 does not require a and b to be different arguments. In particular, an argument that is not contained in any σ -extension is in conflict with itself. This conflict is explicit if the argument is self-attacking and implicit otherwise.

Definition 4. Given a semantics σ , an AF F is called *analytic* for σ (or σ -analytic) if all conflicts of F for σ are explicit in F . The set of all *analytic argumentation frameworks* for σ is denoted by XAF_{σ} .

Example 3. As a simple example consider the AF F_2 from the introduction, depicted in Figure 2. For $\sigma \in \{stb, prf, sem, stg\}$ we have $\sigma(F_2) = \{\{a, d\}, \{b, c\}\}$. Observe that there is an implicit conflict between arguments c and d , denoted by a dashed line in Figure 2. Hence F_2 is not σ -analytic, i.e. $F_2 \notin XAF_{\sigma}$. Observe however that $nai(F_2) = \sigma(F_2) \cup \{\{c, d\}\}$, which means that F_2 is analytic for naive semantics. \diamond

As indicated in Example 3 the sets of analytic AFs can differ for different semantics. Again, well-known relations between the extensions of certain semantics allow us to obtain \subseteq -relations between classes of analytic AFs.

Lemma 3. For any two semantics σ and θ such that for each AF F , for every $S \in \sigma(F)$ there is some $S' \in \theta(F)$ with $S \subseteq S'$, we have $XAF_{\sigma} \subseteq XAF_{\theta}$.




Figure 7: Relations between classes of analytic AFs (cf. Theorem 4).

Proof. Let $F \in XAF_\sigma$ and let there be a conflict between arguments $a, b \in A_F$ for θ , i.e. $(a, b) \notin Pairs_{\theta(F)}$. Now since for every $S \in \sigma(F)$ there is some $S' \in \theta(F)$ with $S \subseteq S'$ it follows that $Pairs_{\sigma(F)} \subseteq Pairs_{\theta(F)}$. Hence also $(a, b) \notin Pairs_{\sigma(F)}$. By the assumption that $F \in XAF_\sigma$ we know that there is an attack $a \rightarrow_F b$ or $b \rightarrow_F a$, hence also $F \in XAF_\theta$. \square

Similarly as for compact AFs, observe that every symmetric and irreflexive (i.e. no self-attacking arguments) AF is contained in XAF_σ for all semantics under consideration.

The next result provides a full picture of the relations between classes of analytic AFs for the semantics we consider (see also Figure 7). We will frequently use Lemma 3, with either the exact condition or the special case $\sigma(F) \subseteq \theta(F)$.

Theorem 4. *The following relations hold:*

1. $XAF_{stb} \subset XAF_\sigma \subset XAF_{nai}$ for $\sigma \in \{prf, sem, stg\}$;
2. $XAF_{sem} \subset XAF_{prf}$;
3. $XAF_{stg} \not\subseteq XAF_\theta$ and $XAF_\theta \not\subseteq XAF_{stg}$ for $\theta \in \{prf, sem\}$.

Proof. (1) Let $\sigma \in \{prf, sem, stg\}$. The \subseteq -relations are due to Lemma 3 together with following facts: (a) in any AF F , $stb(F) \subseteq \sigma(F)$; (b) each σ -extension E of an AF F is conflict-free in F , thus there exists a naive extension E' of F with $E \subseteq E'$.

$XAF_\sigma \subset XAF_{nai}$: The AF in Figure 2 is, as discussed in Example 3, *nai*-analytic but not σ -analytic.

$XAF_{stb} \subset XAF_\sigma$: Consider the AF F from Figure 8. It contains several kinds of complete subframeworks, in the sense that each member of such a subframework attacks each other member. Two complete subframeworks of nine arguments ($\{r_i, u_i, x_i \mid 1 \leq i \leq 3\}$ and $\{s_i, v_i, y_i \mid 1 \leq i \leq 3\}$) and three complete subframeworks of six arguments ($\{r_i, s_i \mid 1 \leq i \leq 3\}$, $\{u_i, v_i \mid 1 \leq i \leq 3\}$ and $\{x_i, y_i \mid 1 \leq i \leq 3\}$). Further there are three directed three-cycles (among $\{a_i \mid 1 \leq i \leq 3\}$, $\{b_i \mid 1 \leq i \leq 3\}$ and $\{c_i \mid 1 \leq i \leq 3\}$), and each argument from the complete subframeworks attacks exactly two arguments from one three-cycle, effectively activating the third one. Now observe that we have $stb(F) = \emptyset$, as at least one argument of a_i, b_i, c_i remains out of range due to conflict-freeness, i.e. a conflict-free set in F can have only one argument from each complete nine-component and thus leaves at least one of the three-cycles unattacked. Therefore there is an implicit conflict for *stb* for every pair of non-attacking arguments, hence $F \notin XAF_{stb}$. On the other hand we have $prf(F) = sem(F) = \{\{r_i, v_j, a_i, b_j\}, \{s_i, u_j, a_i, b_j\}, \{r_i, y_j, a_i, c_j\}, \{s_i, x_j, a_i, c_j\}, \{u_i, y_j, b_i, c_j\}, \{v_i, x_j, b_i, c_j\} \mid 1 \leq i, j \leq 3\}$ and $stg(F) = \{\{r_i, v_j, a_i, b_j, c_k\}, \{s_i, u_j, a_i, b_j, c_k\},$




Figure 8: AF F with $F \in XAF_\sigma$ for $\sigma \in \{prf, sem, stg\}$ and $F \notin XAF_{stb}$.




Figure 9: AF F with $F \in XAF_{prf}$ and $F \notin XAF_\sigma$ for $\sigma \in \{stb, sem, stg\}$.

$\{r_i, y_j, a_i, c_j, b_k\}, \{s_i, x_j, a_i, c_j, b_k\}, \{u_i, y_j, b_i, c_j, a_k\}, \{v_i, x_j, b_i, c_j, a_k\} \mid 1 \leq i, j, k \leq 3\}$, which allows to verify that all conflicts for σ are explicit in F , hence $F \in XAF_\sigma$.

(2) By Lemma 3 we get $XAF_{sem} \subseteq XAF_{prf}$. In order to obtain properness of this relation consider the AF F from Figure 9 and define a cyclic successor function s as $s(1) = 2, s(2) = 3, s(3) = 1$, and $s(4) = 5, s(5) = 6, s(6) = 4$. We have $sem(F) = \{\{x_i, y_j, z_{s(i)}, z_{s(j)}\} \mid i \in \{1, 2, 3\}, j \in \{4, 5, 6\} \text{ or } i \in \{4, 5, 6\}, j \in \{1, 2, 3\}\}$, yielding plenty of implicit conflicts, e.g. between x_i and y_i . Hence F is not analytic for semi-stable semantics. We further define $s(\{i\}) = s(i)$ and for $s(i) = j$ also $s(\{i, j\}) = s(j)$. Now on the other hand we have $prf(F) = sem(F) \cup \{\{x_i, y_j, z_{s(\{i,j\})}\} \mid i, j \in \{1, 2, 3\} \text{ or } i, j \in \{4, 5, 6\}\}$, witnessing $F \in XAF_{prf}$.

(3) $XAF_{stg} \not\subseteq XAF_\sigma$: Consider a directed cycle of five arguments F , $A_F = \{x_1, x_2, x_3, x_4, x_5\}$ and $R_F = \{(x_1, x_2), (x_2, x_3), (x_3, x_4), (x_4, x_5), (x_5, x_1)\}$. Here we have $stg(F) = \{\{x_1, x_3\}, \{x_1, x_4\}, \{x_2, x_4\}, \{x_2, x_5\}, \{x_3, x_5\}\}$ and thus $F \in XAF_{stg}$. On the other hand $sem(F) = prf(F) = \{\emptyset\}$, marking all pairs of arguments as being in conflict and thus for instance the conflict between x_1 and x_3 is implicit for prf and sem (and also stb).

$XAF_{prf} \not\subseteq XAF_{stg}$: The AF F in Figure 9 is, as argued in (2), explicit for prf , but not for sem .




Figure 10: AF F with $F \in XAF_{sem}$ for $F \notin XAF_{stg}$. Here F_X refers to the AF from Figure 8 and \bar{x} is in a symmetric attack relationship with all arguments from F_X .

However, it holds that $stg(F) = sem(F)$, hence also $F \notin XAF_{stg}$.

$XAF_{sem} \not\subseteq XAF_{stg}$: As witness of $XAF_{sem} \not\subseteq XAF_{stg}$ consider the AF F from Figure 10. This AF is composed of two subframeworks, F_X from Figure 8 and F_C from Figure 6b, and a connecting interface consisting of argument \bar{x} and its counterpart set $Y = \{\bar{s}_i, \bar{t}_i, \bar{u}_i \mid i \in \{1, 2, 3\}\}$. There are symmetric attacks between the members $\bar{\alpha}$ of Y and their counterparts α from F_C , between \bar{x} and all members of Y , and between \bar{x} and all arguments from F_X .

A key ingredient to this construction is that both, F_C and F_X , on their own do not provide stable extensions and thus at least one argument remains out of range for any stage or semi-stable extension. In addition observe that F_X is compact for both semi-stable and stage, while F_C is compact only for semi-stable, where a is the argument that does not occur in any $S \in stg(F_C)$.

Considering range-maximal (conflict-free or admissible) sets for F we first distinguish between sets S in relation to the argument \bar{x} . In case $\bar{x} \in S$ we have that all arguments from F_X are in range, Y is attacked and thus F_C needs to be evaluated on its own. In case $\bar{x} \notin S$, wlog. assume $Y \subseteq S$ and $a, x_5 \in S$, we have that all of F_C and Y are in range, \bar{x} is attacked and F_X needs to be evaluated on its own. This means that either some argument from F_C or some argument from F_X remains out of range of any semi-stable or stage extension in F and thus $stb(F) = \emptyset$. On a sidenote observe that for very similar reasons F is compact for both, semi-stable and stage semantics.

Recall that F_C is compact for semi-stable, but not for stage (cf. Theorem 2). This immediately means that for stage semantics there is an implicit conflict between \bar{x} and F_C (argument a to be precise). This also means that for semi-stable semantics there are no implicit conflicts between \bar{x} and any argument from F_C .

It remains to show that F indeed is analytic for semi-stable semantics. To this end we still need to investigate possible implicit conflicts between F_X and Y , between F_C and Y , as well as between F_X and F_C , and among arguments from F_C , as well as among arguments from Y .

As mentioned before the range of any semi-stable extension will cover Y and \bar{x} and either all of F_C or all of F_X . We start with extensions S with $Y \subseteq S$ and thus $\bar{x} \notin S$ and, wlog.

fix the evaluation of F_X and consider some arbitrary $S_X \in \text{sem}(F_X)$. First observe that this immediately means that Y does not contain any conflicts and, due to F_X being compact, there are also no conflicts between Y and F_X . As $Y \cup S_X \cup \{c, x_i\} \in \text{sem}(F)$ for $i \in \{1, 2, 3, 4\}$, and for $i \in \{5, 6, 7\}$ also $Y \cup S_X \cup \{a, x_i\} \in \text{sem}(F)$ as well as $Y \cup S_X \cup \{b, c, x_i\} \in \text{sem}(F)$, there are no conflicts between Y and $a, b, c, x_1 \dots x_7$, between c and $b, x_1 \dots x_7$, or between a, b and x_5, x_6, x_7 .

We now investigate extensions $S \in \text{sem}(F)$ that contain gradually less arguments from Y . In the following we will omit certain x_i from extensions, due to in F_C explicit conflicts, for instance x_2 as well as x_4 attack s_1 and t_1 . For $(Y \setminus \{\bar{s}_1\} \cup \{s_1\}) \subseteq S$ we can have $x_i \in S$ for $i \in \{1, 3\}$, and for $i \in \{5, 6, 7\}$ on the other hand $x_i, a \in S$ or $x_i, b \in S$. For $(Y \setminus \{\bar{t}_1\} \cup \{t_1\}) \subseteq S$ we can have $x_i, c \in S$ for $i \in \{1, 3\}$, or for $i \in \{5, 6, 7\}$ on the other hand $x_i, a \in S$. For $(Y \setminus \{\bar{u}_1\} \cup \{u_1\}) \subseteq S$ we can have $x_i, a \in S$ or $x_i, b, c \in S$ for $i \in \{5, 7\}$, or for $i \in \{1, 2, 3, 4\}$ on the other hand $x_i, c \in S$. Hence for symmetry reasons for $i \in \{1, 2, 3\}$ there are no implicit conflicts between arguments s_i, t_i, u_i on the one side and on the other side Y and arguments a, b, c, x_j for $j \in \{1, 2 \dots 7\}$. Here we can already conclude that there are no implicit but only explicit conflicts between F_C and Y in F .

For $i, j, k \in \{1, 2, 3\}$ fixed and $S_Y = Y \setminus \{\bar{s}_i, \bar{t}_j, \bar{u}_k\}$ we have that $S_X \cup S_Y \cup \{s_i, t_j, u_k, x_i\} \in \text{sem}(F)$. This means that there are no conflicts between s_i, t_j and u_k , and subsequently that the subframework F_C does not have any implicit conflicts in F .

Now finally, as elaborated on, each argument from F_C can appear in semi-stable extensions S of F that do not contain \bar{x} and thus contain some arbitrary F_X -extension S_X . This means that there are no conflicts between F_C and F_X , which closes the gaps and shows that F indeed is analytic for semi-stable semantics. \square

3.3 Relations between Compact and Analytic Frameworks

In the previous two subsections we have separately investigated relations between semantics for compact and analytic AFs respectively. It looks like the relations (Theorems 2 and 4) are not only similar but indeed equal. The question why we looked at the different classes of AFs separately and whether the equal subset relations are based on stronger similarities must be answered two-fold.

On the one hand the examples used for the different proofs share exploitation of similar properties for each semantics considered, and for instance Figure 10 actually builds upon fine-tuned relations between the properties of being compact or analytic. On the other hand in fact not a single example could be used in the other subsection. The compact AFs are not analytic or the analytic AFs are not compact. In what follows we draw some relations between the two classes. We start with similarities as observed in self-loop free AFs.

Proposition 5. *For any $F \in XAF_\sigma$ that is self-loop free, $F \in CAF_\sigma$ ($\sigma \in \{nai, stb, prf, sem, stg\}$).*

Proof. Observe that in Definition 3 we allow arguments in conflict to be equal. Hence for any semantics rejected arguments are in conflict with themselves, and rejected arguments in analytic AFs need to be self-attacking. If there is no self-loop in some analytic AF then naturally there is no rejected argument. \square

For naive semantics we can provide even stronger observations.

Proposition 6. For any self-loop free AF F we have $F \in CAF_{nai}$ and $F \in XAF_{nai}$.

Proof. Two self-loop free arguments where none is attacking the other form a conflict-free set. Since we are dealing with finite sets only this immediately means that there is a naive extension containing both arguments. \square

Proposition 7. $CAF_{nai} \subset XAF_{nai}$.

Proof. For an AF $F \in CAF_{nai}$ it holds that F is self-loop free, hence $F \in XAF_{nai}$ by Proposition 6. Properness is by the AF $(\{a\}, \{(a, a)\})$, which is *nai*-analytic, but not *nai*-compact. \square

However observe that still not every AF is analytic for naive semantics. To see this consider the AF $(\{a, b\}, \{(a, a)\})$. Here $\{b\}$ is the only naive extension, which means that a and b share an implicit conflict.

Finally we conclude this subsection with an observation on the missing relations. That is, we provide reasons why except for naive semantics the properties of being compact or analytic are sufficiently distinct, despite their similarities.

Proposition 8. For $\sigma \in \{stb, sem, prf, stg\}$, we have $CAF_\sigma \not\subseteq XAF_\sigma$ and $XAF_\sigma \not\subseteq CAF_\sigma$.

Proof. Consider the AF from Figure 2. We have as σ -extensions $\{a, d\}$ and $\{b, c\}$. Hence the AF is compact, but not analytic as the conflict between c and d is implicit only, resulting in $CAF_\sigma \not\subseteq XAF_\sigma$.

For $XAF_\sigma \not\subseteq CAF_\sigma$ consider the AF $(\{a, b\}, \{(a, b), (b, b)\})$. This AF consists of one accepted and one rejected argument only. It is analytic but not compact. \square

4 Complexity

When aiming for the simplification of an AF along the dimensions of rejected arguments and implicit conflicts the very first questions one has to face is whether there are any rejected arguments or implicit conflicts, in other words whether the AF is already compact, analytic resp., or there is potential for simplifications. That is, in the following we focus on the computational complexity of the following problems for the semantics σ under consideration: (1) decide whether a given AF is σ -compact or not and (2) to decide whether a given AF is σ -analytic or not. Note that the first problem can also be stated as a decision problem for fairness: given an AF, does each argument appear in at least one σ -extension? Further complexity issues for these two classes are mentioned at the end of the section.

As being compact means that each argument must be credulously accepted, this question is closely related to credulous reasoning (the decision problem $Cred_\sigma$ is defined by the question whether, given an AF F and an argument a , a is contained in at least one σ -extension of F , i.e. whether $a \in Args_{\sigma(F)}$ holds). Exploiting this observation we first give a generic upper bound for the computational complexity.

Theorem 9. For any argumentation semantics σ , with $\text{Cred}_\sigma \in \mathcal{C}$ for a complexity class \mathcal{C} closed under conjunction⁴, we have that deciding whether an AF is compact for σ is in \mathcal{C} .

Proof. By definition an AF $F = (A, R)$ is σ -compact if each $a \in A$ is credulously accepted w.r.t. σ . Hence to check whether F is compact we simply evaluate $\bigwedge_{a \in A} \text{Cred}(F, a)$, which is only of linear size and by assumption can be evaluated in \mathcal{C} as well. \square

We have a similar observation for analytic frameworks, when employing complexity results for Cred_σ^2 .

Theorem 10. For any argumentation semantics σ , with $\text{Cred}_\sigma^2 \in \mathcal{C}$ for a complexity class \mathcal{C} closed under conjunction, we have that deciding whether an AF is analytic for σ is in \mathcal{C} .

Proof. By definition an AF $F = (A, R)$ is σ -analytic if each pair $\{a, b\} \in A$ with neither $(a, b) \in R$ nor $(b, a) \in R$ is credulously accepted w.r.t. σ . Hence to check whether F is analytic for σ we simply conjoin all these tests (only polynomially many), each of which can be done in \mathcal{C} . \square

As P, NP and Σ_2^P are closed under conjunctions we obtain upper bounds for all semantics under our considerations.

In particular, we have the following results for naive semantics.

Corollary 11. The following problems are in P:

1. Given AF F , deciding whether $F \in \text{CAF}_{\text{nai}}$;
2. Given AF F , deciding whether $F \in \text{XAF}_{\text{nai}}$.

Towards our generic hardness result we introduce the concept of SCC-splittable⁵ semantics. Recall that we write $F|_S$ as shorthand for $(A_F \cap S, R_F \cap (S \times S))$.

Definition 5. A semantics σ is called SCC-splittable if there exists a function $\mathcal{GF}_\sigma : \mathfrak{F} \times 2^{\mathfrak{A}} \rightarrow 2^{\mathfrak{A}}$, with \mathfrak{F} being the set of all AFs over \mathfrak{A} , such that the following holds for every AF $F = (A, R) \in \mathfrak{F}$.

- $\mathcal{GF}_\sigma(F, A) = \sigma(F)$
- If $A = B \cup C$ and R does not contain attacks from C to B then

$$\sigma(F) = \bigcup_{E \in \mathcal{GF}_\sigma(F|_B, B)} \{E \cup E' \mid E' \in \mathcal{GF}_\sigma(F|_{C \setminus E_F^+}, U_E^C)\}$$

with $U_E^C = \{c \in C \setminus E_F^+ \mid \forall a \in B : (a, c) \in R \rightarrow a \in E_F^+\}$.

⁴A complexity class \mathcal{C} is closed under conjunctions iff for any problem $\Gamma \in \mathcal{C}$ the problem of deciding whether for a finite set of instances of Γ each of these instances is a yes-instance is also in \mathcal{C} .

⁵Here SCC refer to strongly connected component and reflects the fact that our notion of SCC-splittable is inspired by the notion of SCC-recursiveness [4].




Figure 11: The AF F' from the reduction in the proof of Theorem 13, for AF $F = (\{a, b, c, x\}, \{(a, b), (b, x), (c, x)\})$.

Observe that the second item implies that each strongly connected component of F is either included in B or C .

Splitting argumentation frameworks was studied in [6] where (among others) splittings for stable and preferred semantics are presented. Although the splitting theorem in [6] is not stated in terms of Definition 5 it immediately gives a function $\mathcal{G}\mathcal{F}_\sigma$ with the desired properties. We need one more definition.

Definition 6. A semantics σ is called *rational*, if for any AF F that is a clique (i.e. F is of the form $(A, \{(a, b) \mid a, b \in A, a \neq b\})$) it holds that $\sigma(F) = \{\{a\} \mid a \in A_F\}$.

Proposition 12. *Stable and preferred semantics are rational and SCC-splittable.*

Next we give the generic hardness results for semantics that are rational and SCC-splittable.

Theorem 13. *For any rational SCC-splittable argumentation semantics σ deciding whether an AF is compact for σ is as hard as Cred_σ when restricted to AFs without self-attacks.*

Proof. We reduce the problem Cred_σ to deciding whether an AF is compact for σ . That is given an instance $F = (A, R), x \in A$ of Cred_σ we build the following AF $F' = (A \cup A', R \cup R')$ with $A' = \{t_a \mid a \in A\}$ and

$$R' = \{(t_a, t_b) \mid a, b \in A, a \neq b\} \cup \{(t_a, b) \mid a, b \in A, a \neq x, b \neq a\}.$$

That is, we add a clique AF $C_A = (A', \{(t_a, t_b) \mid a, b \in A, a \neq b\})$ of size $|A|$ and link it to the original framework as follows: The argument t_x does not attack any of the original arguments. All the other arguments t_a attack all but one of the original arguments and thus, as we discuss below, enforces that this argument is credulously accepted. The construction is illustrated in Figure 11.

To prove the claim we have to show that x is credulously accepted in F iff F' is σ -compact. First observe that the new arguments in F' form a SCC and are not attacked by arguments from outside. As σ is SCC-splittable we can evaluate F' as follows:

1. Compute the extensions of the clique C_A .




Figure 12: The AF F' from the reduction in the proof of Theorem 14, for AF $F = (\{a, b, c\}, \{(a, b)\})$.

2. For each such extension E of C_A build the AF $F_{|A \setminus E_{F'}^+}$ by removing all arguments in $E_{F'}^+$ from F .
3. For each extension $E' \in \mathcal{GF}_\sigma(F_{|A \setminus E_{F'}^+}, U_E^A)$ return $E \cup E'$.

By assumption the extensions of C_A are the singletons $\{t_a\}$. First, consider $E_a = \{t_a\}$ with $a \neq x$, then $F_{|A \setminus E_{F'}^+} = (\{a\}, \{\})$ and $U_E^A = \{a\}$. We have $\mathcal{GF}_\sigma(F_{|A \setminus E_{F'}^+}, U_E^A) = \mathcal{GF}_\sigma((\{a\}, \{\}), \{a\}) = \sigma((\{a\}, \{\})) = \{a\}$ (the latter is since σ is rational). Thus for each $a \neq x$ the set $\{t_a, a\}$ is a σ -extension of F' . Second, consider $E_x = \{t_x\}$. Here $F_{|A \setminus E_{F'}^+} = F$ and $U_E^A = A$. Thus for each $E \subseteq A$ we have that $E \in \mathcal{GF}_\sigma(F, A) = \sigma(F)$ iff $\{t_x\} \cup E \in \sigma(F')$. Hence, x is credulously accepted (w.r.t. σ) in F iff x is credulously accepted (w.r.t. σ) in F' iff F' is σ -compact. \square

Theorem 14. *For any rational SCC-splittable argumentation semantics σ deciding whether an AF is analytic for σ is as hard as deciding whether an AF is compact for σ . The result even holds if one knows that the AF being tested for being analytic is already compact.*

Proof. We reduce the problem of deciding whether an AF F is compact to deciding whether F' is analytic. That is given an instance $F = (A, R)$ (we can assume that F has no self-attacks as otherwise it is an immediate no-instance) we build the following AF $F' = (A \cup A', R \cup R')$ with $A' = \{t\} \cup \{t_{\{a,b\}} \mid a, b \in A, (a, b) \notin R, (b, a) \notin R\}$ and

$$R' = \{(t_1, t_2) \mid t_1, t_2 \in A', t_1 \neq t_2\} \cup \{(t_{\{a,b\}}, c) \mid t_{\{a,b\}} \in A', c \in A, a \neq c, b \neq c\}.$$

That is we add a clique AF C of size at most $(|A|^2 + |A|)/2 + 1$ to F with a distinguished element t and link it to the original framework as follows: The argument t does not attack any of the original arguments. All the other arguments $t_{\{a,b\}}$ attack all original arguments in F except a and b (note that a and b are not necessarily distinct). The construction is illustrated in Figure 12.

To prove the claim we have to show that F is σ -compact iff F' is σ -analytic. First observe that the new arguments in F' form a strongly connected component (SCC) and are not attacked by arguments from outside. As σ is SCC-splittable we can evaluate F' as follows:

1. Compute the extensions of the clique C .

2. For each such extension E of C build the AF $F_{|A \setminus E_{F'}^+}$ by removing all arguments in $E_{F'}^+$ from F .
3. For each extension $E' \in \mathcal{GF}_\sigma(F_{|A \setminus E_{F'}^+}, U_E^A)$ return $E \cup E'$.

By assumption that σ is rational we have $\sigma(C) = \{\{t'\} \mid t' \in A'\}$. First consider an extension E of the form $\{t_{\{a,b\}}\}$ and recall that then we have $(a,b) \notin R$ and $(b,a) \notin R$. Then $F_{|A \setminus E_{F'}^+} = (\{a,b\}, R \cap \{a,b\} \times \{a,b\}) = (\{a,b\}, \{\})$. $U_E^A = \{a,b\}$. We have $\mathcal{GF}_\sigma(F_{|A \setminus E_{F'}^+}, U_E^A) = \mathcal{GF}_\sigma(\{\{a,b\}, \{\}\}, \{a,b\}) = \sigma(\{\{a,b\}, \{\}\}) = \{\{a,b\}\}$.⁶ Thus for each $a,b \in A$ such that $(a,b) \notin R$ and $(b,a) \notin R$, the set $\{t_{\{a,b\}}, a, b\}$ is a σ -extension of F' . This already shows that for any pair (x,y) of arguments in F' where x and y are different from the distinguished argument t in C , we have that x,y are jointly contained in at least one σ -extension iff there is no attack $x \rightarrow y$ or $y \rightarrow x$ in F' . Now, consider $E = \{t\}$. Here $F_{|A \setminus E_{F'}^+} = F$ and $U_E^A = A$. Thus for each $E \subseteq A$ we have that $E \in \mathcal{GF}_\sigma(F, A) = \sigma(F)$ iff $\{t\} \cup E \in \sigma(F')$. Recall that there is no attack between t and arguments in F . Now, F is σ -compact iff, for each $a \in A$, t occurs together with a in at least one σ -extension of F' . Together with our previous observation, we conclude that F is σ -compact iff F' is σ -analytic.

Finally, as for each $a \in A$ and $t_{\{a\}} = t_{\{a,a\}}$ the set $\{t_{\{a\}}, a\}$ is credulously accepted the AF F' is compact. \square

From the generic results above we immediately get the complexity characterization for stable and preferred semantics.

Corollary 15. *The following problems are NP-complete for $\sigma \in \{stb, prf\}$.*

1. Given AF F , deciding whether $F \in CAF_\sigma$;
2. Given AF F , deciding whether $F \in XAF_\sigma$; hardness already holds if the problem is restricted to AFs $F \in CAF_\sigma$.

Proof. Recall that $Cred_{stb}$ and $Cred_{prf}$ are NP-complete [17] and that NP is closed under conjunction. Membership thus follows from Theorems 9 and 10. Furthermore, stb and prf are SCC-splittable [6] and rational. Theorems 13 and 14 thus give the matching hardness results. \square

Theorems 13 and 14 do not apply to stage and semi-stable semantics (as they are not SCC-splittable). However we can extend the results to these semantics by carefully adapting the ideas from the proofs of Theorem 13 and 14. The main idea is still the same: we take the original AF F and add a gadget of arguments that attack certain arguments in F but whose arguments are not attacked by arguments of F . Such a gadget (replacing the clique) has to satisfy certain properties: (i) its evaluation is independent of F ; (ii) all arguments of the gadget are credulously accepted; (iii) there are certain arguments selecting a single argument, resp. a pair, of the original AF for acceptance, by attacking all the other arguments of the original AF; (iv) the gadget for testing $F \in CAF$ does not affect the acceptance of the argument under question and in the gadget for testing $F \in XAF$ there is an argument t that maintains all extensions E of F as extensions $\{t\} \cup E$.

⁶Notice that $\sigma(\{\{a,b\}, \{\}\}) = \{\{a,b\}\}$ for each rational SCC-splittable argumentation semantics σ .




Figure 13: The AF F' from Reduction 1, for AF $F = (\{a, b, c, x\}, \{(a, b), (b, x), (c, x)\})$.

Theorem 16. *Given AF F :*

1. *Deciding whether $F \in CAF_{stg}$ is Σ_2^P -complete.*
2. *Deciding whether $F \in CAF_{sem}$ is Σ_2^P -complete.*

We split the proof of Theorem 16 into several Lemmas. First, we have to show that both problems can be solved in Σ_2^P .

Lemma 17 (Membership in Σ_2^P). *Both deciding whether $F \in CAF_{stg}$ and deciding whether $F \in CAF_{sem}$ are in Σ_2^P .*

Proof. The membership in Σ_2^P follows from the memberships of $Cred_{stg}$ and $Cred_{sem}$ in Σ_2^P [27, 13] and Theorem 9. \square

For hardness we give a reduction that constructs an AF F' given an AF F and an argument x . Although, this reduction will be used for both hardness proofs we will apply it to different problems, i.e. $Cred_{stg}$ and $Cred_{sem}$, to show hardness for both stg and sem .

Reduction 1. *Given an AF $F = (A, R)$ and $x \in A$ we build the AF $F' = (A \cup A', R \cup R')$ with*

$$\begin{aligned}
 A' &= \{t_a \mid a \in A\} \cup \{t_y, t_z\} \cup \{h_1, h_2, h_3\} \\
 R' &= \{(t_a, t_b) \mid a, b \in A \cup \{y, z\}, a \neq b\} \cup \{(t_a, b) \mid a \in A \setminus \{x\}, b \in A \setminus \{a\}\} \cup \\
 &\quad \{(h_1, h_2), (h_2, h_3), (h_3, h_1)\} \cup \{(t_x, h_1), (t_y, h_2), (t_z, h_3)\}
 \end{aligned}$$

The construction is illustrated in Figure 13.

For both semantics we can assume that F has no self-attacks [27] and no stable extension. To achieve the second we can add an odd length cycle to F that is not connected to any other argument. This will guarantee that there is no stable extension and does not affect credulous acceptance w.r.t. semi-stable or stage semantics.

Lemma 18. *Given an AF F without self-attacks and stable extensions then $F' \in CAF_{stg}$ iff $(F, x) \in Cred_{stg}$.*

Proof. Below we will show that all arguments except x are always credulously accepted in F' and that x is credulously accepted in F' iff x is credulously accepted in F .

First, we show that each stage extension E contains at least one argument from $\{t_a \mid a \in A \cup \{y, z\}\}$. Suppose that not, then, as at most one of the arguments h_1, h_2, h_3 is contained in E , two of the sets $E \cup \{t_x\}$, $E \cup \{t_y\}$, $E \cup \{t_z\}$ are conflict-free. Hence we have a contradiction to the maximality of E . Further, as $\{t_a \mid a \in A \cup \{y, z\}\}$ forms a clique in F we get that each stage extension contains exactly one argument from the set.

Next we show that the ranges of naive sets E containing an argument from $\{t_a \mid a \in \{x, y, z\}\}$ cannot be contained in the ranges of conflict-free sets E' not containing any of these arguments.

- If $t_x \in E$ then h_1 is attacked by E and cannot be in E but is in the range of E . Now as h_2 gives the larger range than h_3 we can conclude that $h_2 \in E$ and $\{h_1, h_2, h_3\} \in E_{F'}^+$. By similar arguments we get that $\{h_1, h_2, h_3\} \subseteq E_{F'}^+$ if either $t_y \in E$ or $t_z \in E$.
- If $t_a \in E'$ with $a \notin \{x, y, z\}$ then only one of $\{h_1, h_2, h_3\}$ can be contained in E' and thus at most two of them are in the range of E' .

We will next consider these two kinds of extensions separately.

- First, consider the sets E containing an argument from $\{t_a \mid a \in \{x, y, z\}\}$. By the above we have that either $\{t_x, h_2\} \subseteq E$, $\{t_y, h_3\} \subseteq E$, or $\{t_z, h_1\} \subseteq E$. All of these three sets have the same attacks to the remaining arguments and thus we have that for each $E' \subseteq A$, $\{t_x, h_2\} \cup E' \in stg(F')$ iff $\{t_y, h_3\} \cup E' \in stg(F')$ iff $\{t_z, h_1\} \cup E' \in stg(F')$. As at least for $E' = \emptyset$ these sets are conflict-free this implies that the arguments $\{t_x, t_y, t_z, h_1, h_2, h_3\}$ are all credulously accepted in F' .

Moreover, the set $\{t_x, h_2\}$ does not attack any argument in A nor does A have any outgoing attacks. Thus

- (i) $\{t_x, h_2\} \cup E' \in cf(F')$ iff $E' \in cf(F)$ and
- (ii) as arguments in A do not attack arguments in A' we have that $(\{t_x, h_2\} \cup E')_{F'}^+ = \{t_x, h_2\}_{F'}^+ \cup E_{F'}^+$ and thus $(\{t_x, h_2\} \cup E')_{F'}^+$ is maximal when $E_{F'}^+$ is maximal.

Hence, $\{t_x, h_2\} \cup E' \in stg(F')$ iff $E' \in stg(F)$ and x is credulously accepted in F iff x is credulously accepted in F' .

As, by assumption, F has no stable extension there cannot be an extension E containing an argument from $\{t_a \mid a \in \{x, y, z\}\}$ and having all arguments A in its range.

- Second, consider the sets E containing an argument from $\{t_a \mid a \in A \setminus \{x\}\}$. Now it is easy to verify that the sets $\{t_a, a\}$ for $a \in A \setminus \{x\}$ are conflict-free sets of F' and have maximal range among the sets containing $\{t_a \mid a \in A \setminus \{x\}\}$. In particular A is in the range of each of these extensions, and as the extensions of the first type never have the whole set A in their

range (cf. first item) a set $\{t_a, a\}$ for $a \in A \setminus \{x\}$ is incomparable to these extensions, i.e. $\{t_a, a\}$ is a stage extension. Hence, we have that the arguments $\{a, t_a \mid a \in A \setminus \{x\}\}$ are credulously accepted. Moreover, no extensions E' with $\{t_a \mid a \in A \setminus \{x\}\} \cap E' \neq \emptyset$ can contain x .

Finally, combining the above results, we have that all arguments in A' except x are credulously accepted in F' and x is credulously accepted in F' iff x is credulously accepted in F iff F' is *stg*-compact. \square

Now, as Reduction 1 can be performed in polynomial-time and $Cred_{stg}$ is Σ_2^P -hard [27], Lemma 18 implies that deciding whether $F \in CAF_{stg}$ is Σ_2^P -hard.

Lemma 19. *Given an AF F without self-attacks and stable extensions then $F' \in CAF_{sem}$ iff $(F, x) \in Cred_{sem}$.*

The proof of the above lemma is very similar to the proof of Lemma 18 and thus omitted here, but provided in B. Now, as Reduction 1 can be performed in polynomial-time and $Cred_{sem}$ is Σ_2^P -hard [27], Lemma 19 implies that deciding whether $F \in CAF_{stg}$ is Σ_2^P -hard. This completes the proof of Theorem 16.

Next, starting from Theorem 16 we can show that also deciding whether an AF is analytic for stage or semi-stable is Σ_2^P -complete.

Theorem 20. *Given AF F :*

1. *Deciding whether $F \in XAF_{stg}$ is Σ_2^P -complete.*
2. *Deciding whether $F \in XAF_{sem}$ is Σ_2^P -complete.*

For both problems, hardness already holds if the problem is restricted to AFs $F \in CAF_\sigma$.

We split the proof of Theorem 20 into several Lemmas, starting with showing that both problems can be solved in Σ_2^P .

Lemma 21 (Membership in Σ_2^P). *Both deciding whether $F \in XAF_{stg}$ and deciding whether $F \in XAF_{sem}$ are in Σ_2^P .*

Proof. The membership in Σ_2^P follows from the memberships of $Cred_{stg}^2, Cred_{sem}^2$ in Σ_2^P [27, 13] and Theorem 10. \square

For hardness we give a reduction that constructs an AF F' given an AF F . Again this reduction will be used for both *stg* and *sem* but the hardness arguments will start from different problems, i.e. from testing whether an AF is in CAF_{stg} , in CAF_{sem} respectively.




Figure 14: The AF F' from Reduction 2, for AF $F = (\{a, b, c\}, \{(a, b)\})$.

Reduction 2. Given an AF $F = (A, R)$ and the AF⁷ $G = (A_G, R_G)$ from Figure 8 we build the AF $F' = (A \cup A_G \cup A', R \cup R_G \cup R')$ with

$$\begin{aligned} A' &= \{t_{a,b} \mid \{a, b\} \subseteq A, \{(a, b), (b, a)\} \cap R = \emptyset\} \cup \{t\} \\ R' &= \{(t_1, t_2) \mid t_1, t_2 \in A', t_1 \neq t_2\} \cup \{(t, x) \mid x \in A_G\} \cup \\ &\quad \{(t_{a,b}, c) \mid \{a, b\} \subseteq A, \{(a, b), (b, a)\} \cap R = \emptyset, c \in A \setminus \{a, b\}\} \end{aligned}$$

For both semantics we can assume that F has no self-attacks and no stable extension. This is by the fact that we made the same assumptions for the hardness proofs of Theorem 16 and the fact that Reduction 1 introduces neither self-attacks nor stable extensions.

Lemma 22. Given an AF F without self-attacks and stable extensions then $F' \in CAF_{stg}$, and $F \in CAF_{stg}$ iff $F' \in XAF_{stg}$.

Proof. First, we show that each stage extension E contains at least one argument from A' . Suppose that not, then $E \setminus A_G \cup \{t\}$ is a conflict-free set that has A_G in its range. Hence we have a contradiction to the range maximality of E . Further, as A' forms a clique in F we get that each extension contains exactly one argument from the set.

Next we show that the ranges of naive sets E containing argument t cannot be contained in the ranges of conflict-free sets E' containing an argument $t_{a,b}$ with $a, b \in A$.

- If $t \in E$ then all arguments in A_G are attacked by E and thus are in the range of E .
- If $t \notin E'$ at least one argument of A_G is not in the range of E' . Otherwise, $E' \cap A_G$ would be a stable extension of G , which contradicts $stb(G) = \emptyset$.

We will next consider these two kind of extensions separately.

- First, consider the sets E containing t . As t does not attack any argument in A nor does A have any outgoing attacks we have

$$(i) \quad \{t\} \cup E' \in cf(F') \text{ iff } E' \in cf(F) \text{ and}$$

⁷We can use here any AF G without self-attacks with $G \in XAF_{stg}$, $G \in XAF_{sem}$ and $stb(G) = \emptyset$.

(ii) as arguments in A do not attack arguments outside A we have that $(\{t\} \cup E')_{F'}^+ = \{t\}_{F'}^+ \cup E'_{F'}^+$ and thus $(\{t\} \cup E')_{F'}^+$ is maximal when $E'_{F'}^+$ is maximal.

Hence, $\{t\} \cup E' \in stg(F')$ iff $E' \in stg(F)$ and thus for each $a \in A$ we have that $\{t, a\}$ is credulously accepted in F' iff a is credulously accepted in F .

As, by assumption, F has no stable extension there cannot be an extension E containing t and having all arguments A in its range.

- Second, consider the sets E containing an argument $t_{a,b}$ with $a, b \in A$. Now it is easy to verify that the sets $\{t_{a,b}, a, b\}$ are conflict-free sets of F' and A is in the range of each of these extensions. As the extensions of the first type never have the whole set A in their range (cf. first item) a set $\{t_{a,b}, a, b\}$ is incomparable to these extensions.

As $t_{a,b}$ does not attack any argument in A_G nor does A_G have any outgoing attacks we have

(i) $\{t_{a,b}, a, b\} \cup E' \in cf(F')$ iff $E' \in cf(G)$ and

(ii) as arguments in A_G do not attack arguments outside A_G we have that $(\{t_{a,b}, a, b\} \cup E')_{F'}^+ = \{t_{a,b}, a, b\}_{F'}^+ \cup E'_{F'}^+$ and thus $(\{t_{a,b}, a, b\} \cup E')_{F'}^+$ is maximal when $E'_{F'}^+$ is maximal.

Thus, $\{t_{a,b}, a, b\} \cup E' \in stg(F')$ iff $E' \in stg(G)$. Now, as $G \in XAF_{stg}$ we have that for each $g, g' \in G$ with $(g, g'), (g', g) \notin R_G$ there is an $E' \in stg(G)$ with $g, g' \in E'$. Furthermore as G has no self-attacks it is also compact (cf. Proposition 5) and thus for each $g \in A_G$ there is an $E' \in stg(G)$ with $g \in E'$. From these stage extensions we obtain that:

- $\{t_{a,b}, a\}, \{t_{a,b}, b\}$ are credulously accepted in F' ;
- $\{t_{a,b}, g, g'\}$ is credulously accepted in F' , for $g, g' \in G$ with $(g, g'), (g', g) \notin R_G$;
- $\{t_{a,b}, g\}$ is credulously accepted in F' , for each $g \in G$;
- $\{a, g\}$ is credulously accepted in F' , for each $a \in A$ and $g \in G$;

Combining the above results, we have that all non-conflicting pairs of arguments in F' except $\{t, a\}$ with $a \in A$ are credulously accepted in F' . Thus F' is stg -analytic iff all the pairs $\{t, a\}$ with $a \in A$ are credulously accepted in F' iff each $a \in A$ is credulously accepted in F iff F is stg -compact.

Finally we show that $F' \in CAF_{stg}$ (independent of whether $F \in CAF_{stg}$). As (i) for each $a \in A$ the set $\{t_{a,a}, a\}$ is credulously accepted, and (ii) for each $g \in A_G$ and $a, b \in A$ with $(a, b), (b, a) \notin R$ the set $\{t_{a,b}, g\}$ is credulously accepted, the AF F' is stg -compact. \square

Now, as Reduction 2 can be performed in polynomial-time and CAF_{stg} is Σ_2^P -hard [Th. 16], Lemma 22 implies that deciding whether $F \in CAF_{stg}$ is Σ_2^P -hard.

Lemma 23. *Given an AF F without self-attacks and stable extensions then $F' \in CAF_{sem}$, and $F \in CAF_{sem}$ iff $F' \in XAF_{sem}$.*

Table 2: Complexity Results (\mathcal{C} -c denotes completeness for class \mathcal{C}).

	$F \in CAF_\sigma?$	$F \in XAF_\sigma?$
<i>nai</i>	in P	in P
<i>stb</i>	NP-c	NP-c
<i>prf</i>	NP-c	NP-c
<i>stg</i>	Σ_2^P -c	Σ_2^P -c
<i>sem</i>	Σ_2^P -c	Σ_2^P -c

The proof of the above lemma is very similar to the proof of Lemma 22 and thus omitted here, but provided in B. Now, as Reduction 2 can be performed in polynomial-time and CAF_{sem} is Σ_2^P -hard [Th. 16], Lemma 23 implies that deciding whether $F \in CAF_{stg}$ is Σ_2^P -hard. This completes the proof of Theorem 20. \square

In conclusion we have that for all the semantics under our considerations the complexity of testing whether an AF is compact or analytic is as hard as credulous acceptance. We summarize the results of this section in Table 2.

Complexity of further decision problems Similar to other subclasses, compact and analytic AFs decrease the complexity of certain decision problems. Let us first discuss the case of compact AFs. By definition for credulous acceptance (does an argument occur in at least one extension), this problem becomes trivial for this class. For skeptical acceptance (does an argument a occur in all extensions) in compact AFs the problem reduces to checking whether a is isolated. If yes, it is skeptically accepted; if no, a is connected to at least one further argument that has to be credulously accepted by the definition of compact AFs. But then, it is the case for any semantics that is based on conflict-free sets that a cannot be skeptically accepted, since it will not appear together with b in an extension. For analytic AFs we can distinguish between AFs with self-attacks and without. In the latter case the AFs are also compact (cf. Proposition 5) and thus credulous and skeptical acceptance can be solved as described above. In the former case, for credulous acceptance we only have to check whether the argument is self-attacking or not. For skeptical acceptance the behavior seems to diverge between different semantics. On the one hand, for deciding whether an argument is skeptically accepted w.r.t. stable semantics one can test if the argument is credulously accepted and all its attackers are not credulously accepted, which can be done in polynomial time. On the other hand side, for preferred and semi-stable semantics analytic AFs seem to have no computational benefits. Moreover, [7] showed that in compact AFs the verification problem (given AF F and a set of arguments E , is E a σ -extension of F ?) is still coNP-hard for stage, semi-stable and preferred semantics. These results can be extended to analytic AFs by the observation that the reductions used in the proofs of Theorems 14 and 20 are also valid reductions for the verification problem.

5 Explicit Conflict Conjecture

In this section we take another look at the issue of implicit conflicts and the possibility of making them explicit. In Section 3.2 we identified the classes of AFs where all conflicts are explicit w.r.t. a given semantics. Recall the notion of an analytic AF from Definition 4. In [7] the authors conjectured that, under stable semantics, every AF can be translated to an equivalent analytic AF (having the same set of arguments), i.e. that all implicit conflicts can be made explicit without changing the stable extensions. We will refute this conjecture and show that the claim also does not hold for preferred, semi-stable and stage semantics.

Definition 7. An AF F is called *quasi-analytic* for σ if there is an AF G such that $A_F = A_G$, $\sigma(F) = \sigma(G)$ and G is analytic for σ , i.e., it has only explicit conflicts for σ . On the other hand, F is called *non-analytic* for σ if it is not quasi-analytic for σ .

Example 4. Consider again the AF in Figure 2, which, as we have seen in Example 3, is not analytic for $\sigma \in \{stb, prf, sem, stg\}$. However, adding the attack $c \rightarrow d$ (or $d \rightarrow c$ or both) we obtain an equivalent (under σ) AF F' , where all conflicts are explicit. Thus F is quasi-analytic. \diamond

In other words, an AF is quasi-analytic for a given semantics σ if it can be translated to another AF that has the same arguments, has the same extensions under σ , and all conflicts are explicit. The conjecture from [7] says that every AF containing implicit conflicts for stable semantics is quasi-analytic, in the sense that all implicit conflicts can be made explicit without adding further arguments. We repeat the conjecture from [7], just parameterized by an arbitrary semantics. In line with the following definition, [7] claimed that ECC holds for stable semantics.

Definition 8. We say that the *Explicit Conflict Conjecture (ECC)* holds for semantics σ if every AF is quasi-analytic for σ .

First note that, as discussed in the introduction, ECC holds for naive semantics. Every pair of non-self-attacking arguments occurs together in a naive extension if and only if there is no attack between them. Hence a conflict can only be implicit for naive semantics if at least one of the arguments involved is self-attacking. But letting each self-attacking argument be attacked by all other arguments does not change the naive extensions (and obviously does not change the set of arguments), hence every AF is quasi-analytic.

In the remainder of this section we will refute ECC for all semantics in $\{stb, prf, sem, stg\}$ by providing non-analytic AFs.

Example 5. Take into account the AF $F = (A, R)$ depicted in Figure 15, which features an implicit conflict for stable semantics between a and b :

$$\begin{aligned} A &= \{a, b, c\} \cup \{u_i, v_i, x_i, y_i \mid i \in \{1, 2\}\} \\ R &= \{\langle a, c \rangle, \langle b, c \rangle\} \cup \{\langle \alpha_i, \beta_i \rangle \mid i \in \{1, 2\}, \alpha \in \{x, y\}, \beta \in \{u, v\}\} \\ &\quad \cup \{(u_i, a), (a, x_i), (v_i, b), (b, y_i), \langle u_i, v_i \rangle \mid i \in \{1, 2\}\} \end{aligned}$$




Figure 15: Illustration of the AF from Example 5.

In the following we refer to $A_1 = \{v_1\}$, $A_2 = \{u_1\}$, $A_3 = \{x_1, y_1\}$, and $B_1 = \{v_2\}$, $B_2 = \{u_2\}$, $B_3 = \{x_2, y_2\}$. The stable extensions of F can be separated into extensions containing c and others. For $i, j \in \{1, 2, 3\}$ the former are given as:

$$S_{ij} = \{c\} \cup A_i \cup B_j$$

If on the other hand $c \notin S$ one of a, b will be a member of S and thus:

$$\begin{array}{lll} S_1 = \{a, v_1, v_2\} & S_3 = \{a, v_1, y_2\} & S_5 = \{b, u_1, x_2\} \\ S_2 = \{b, u_1, u_2\} & S_4 = \{a, y_1, v_2\} & S_6 = \{b, x_1, u_2\} \end{array}$$

For $S \in \text{stb}(F)$ and wlog. $a \in S$ take into account that a is attacked by u_1 and the only possible defenders v_1 and y_1 are explicitly in conflict with b . Thus clearly a and b share an implicit conflict, as one cannot be defended without the other being attacked. However observe that all the other conflicts implicitly defined by the extension-set $\mathbb{S} = \{S_1, S_2, \dots, S_6\} \cup \{S_{ij} \mid i, j \in \{1, 2, 3\}\}$ are already given explicitly in F . Furthermore the remaining maximal conflict-free sets $S_a = \{a, y_1, y_2\}$ and $S_b = \{b, x_1, x_2\}$ do attack neither b nor a respectively and thus are not stable extensions of F . \diamond

We now proceed by showing that the AF depicted in Figure 15 and discussed in Example 5 serves as a counter-example for ECC for stable semantics.

Theorem 24. *There are non-analytic AFs for stable semantics.*

Proof. Consider the AF F from Example 5 and recall its set of stable extensions \mathbb{S} . We will show that there is no AF G with $A_G = A_F$, $\text{stb}(G) = \mathbb{S}$ and $(a, b) \in R_G$. (Observe that due to symmetry reasons we need not consider $(b, a) \in R_G$ and $(a, b) \notin R_G$.) For a contradiction take such an AF as given.

The extensions containing c ensure that there is no attack in G between arguments c and α_i for $\alpha \in \{x, u, v, y\}$ and $i \in \{1, 2\}$, or between α_1 and α_2 . By definition any stable extension $S \in \mathbb{S}$ attacks all outside arguments, $S \succ \alpha$ for $\alpha \in A_G \setminus S$. Hence from $S_3 = \{a, v_1, y_2\}$ being a stable extension we conclude $a \succ c$ and $\{a, y_2\} \succ \alpha_2$ for $\alpha \in \{x, u, v\}$. Similarly due to $S_4 = \{a, y_1, v_2\}$ we conclude that $\{a, y_1\} \succ \alpha_1$ for $\alpha \in \{x, u, v\}$. But now by assumption $a \succ b$




Figure 16: A non-analytic AF for prf as used in Example 6.

and thus for $S_a = \{a, y_1, y_2\}$ we acquire full range, $S_a \mapsto \alpha$ for any $\alpha \in A_G \setminus S_a$, i.e. S_a becomes an unwanted stable extension. Therefore F is non-analytic. \square

We observe that in this counter-example for ECC for stable semantics the stable extensions coincide with semi-stable, preferred and stage extensions. With the following lemma this leads to some straightforward generalizations.

Lemma 25. *Let F be an AF with $prf(F) = stb(F)$ (resp. $sem(F) = stb(F)$). If F is quasi-analytic for preferred (resp. semi-stable) semantics, then it is also quasi-analytic for stable semantics.*

Proof. By assumption, for $\sigma \in \{prf, sem\}$, there is a σ -analytic AF G such that $A_G = A_F$ and $\sigma(F) = \sigma(G)$. We want to show that $stb(G) = \sigma(G)$. Using the fact that for any AF F , $stb(F) \subseteq \sigma(F)$ holds, it remains to show that $\sigma(G) \subseteq stb(G)$. To this end observe that any attack of F still represents an explicit conflict in G . Now for $S \in stb(F)$ we know that for all $a \in A_F \setminus S$ we have $S \mapsto_F a$. Since by assumption also $S \in \sigma(F)$ this immediately implies an explicit conflict between S and a in G . Due to admissibility of σ -extensions this means that actually $S \mapsto_G a$ as otherwise S would not defend itself from a in G . Therefore we have $S \mapsto_G a$ for all $a \in A_G \setminus S$. Hence $S \in stb(G)$, resulting in $\sigma(G) = stb(G)$ and thus G being stb -analytic and also F being stb -quasi-analytic. \square

Using the AF F from Example 5 and the contraposition of Lemma 25 yields the following result, refuting ECC for preferred and semi-stable semantics.

Corollary 26. *There are non-analytic AFs for preferred and semi-stable semantics, respectively.*

The next example shows that some AFs prove to be non-analytic for preferred semantics while being quasi-analytic for all the other semantics under consideration.

Example 6. Take into account the AF $F = (A, R)$ as depicted in Figure 16 with

$$A = \{a_i, b_i, x_i, u_i \mid i \in \{1, 2, 3\}\}$$

$$R = \{(a_i, b_i), (b_i, x_i), (x_i, u_i) \mid i \in \{1, 2, 3\}\} \cup \{(x_1, x_2), (x_2, x_3), (x_3, x_1)\}$$




Figure 17: Analytic AF for stage semantics, cf. Example 5.

We have $\text{prf}(F) = \{S_a, S_b, A_1, A_2, A_3, B_1, B_2, B_3\}$ and

$$\begin{aligned}
 S_a &= \{a_1, a_2, a_3\} & S_b &= \{b_1, b_2, b_3, u_1, u_2, u_3\} \\
 A_1 &= \{a_2, a_3, b_1, x_2, u_1, u_3\} & B_1 &= \{a_1, b_2, b_3, x_1, u_2, u_3\} \\
 A_2 &= \{a_1, a_3, b_2, x_3, u_1, u_2\} & B_2 &= \{a_2, b_1, b_3, x_2, u_1, u_3\} \\
 A_3 &= \{a_1, a_2, b_3, x_1, u_2, u_3\} & B_3 &= \{a_3, b_1, b_2, x_3, u_1, u_2\}
 \end{aligned}$$

In the following we show that F is non-analytic for preferred semantics. For a contradiction we assume that there exists an analytic AF G with $A_G = A$ and $\text{prf}(F) = \text{prf}(G)$. We now investigate this hypothetical AF G . Observe that for $i, j \in \{1, 2, 3\}$ due to S_b there is no conflict between u_i and b_j , due to A_1, A_2, A_3 there is no conflict between u_i and a_j , and for $i \neq j$ there is no conflict between x_i and u_j ; in other words in G the u_i can be attacked only by the x_i . Furthermore we have an implicit conflict between a_1 and x_2 . Due to S_a being admissible and G being analytic now $S_a \rightsquigarrow_G x_2$. But then S_a defends u_2 and thus can not be a preferred extension in G . For symmetry reasons it follows that the implicit conflicts (a_i, x_j) of F cannot be made explicit for preferred semantics.

On the other hand for stable (or stage or semi-stable) semantics we observe that S_a is not an extension. Although the overall conflicts remain the same, this allows us to include conflicts (x_j, a_i) without any harm for the other extensions. Thus for stable, semi-stable and stage semantics this AF is quasi-analytic. \diamond

We still have not answered the question whether stage semantics possesses non-analytic AFs. A candidate for a non-analytic AF for stage semantics would be the AF F from Example 5, but it turns out to be quasi-analytic for stage semantics. In fact, the analytic AF G depicted in Figure 17 has the same stage extensions as F , $\text{stb}(F) = \text{stg}(F) = \text{stg}(G)$.

However, the following slightly more involved example yields a non-analytic AF for stage semantics.




Figure 18: Illustration of the AF from Example 7.

Example 7. Take into account the AF $F = (A, R)$ depicted in Figure 18 with:

$$\begin{aligned}
A &= \{a, b, c\} \cup \{u_i, v_i, x_i, y_i, r_i, s_i \mid i \in \{1, 2\}\} \\
R &= \{\langle a, c \rangle, \langle b, c \rangle\} \cup \{\langle r_i, x_i \rangle, \langle s_i, y_i \rangle \mid i \in \{1, 2\}\} \\
&\cup \{\langle \alpha_i, \beta_i \rangle \mid i \in \{1, 2\}, \alpha \in \{x, y\}, \beta \in \{u, v\}\} \\
&\cup \{\langle u_i, a \rangle, \langle a, x_i \rangle, \langle v_i, b \rangle, \langle b, y_i \rangle, \langle u_i, v_i \rangle \mid i \in \{1, 2\}\}
\end{aligned}$$

In the following we will refer to $M_{i1} = \{r_i, v_i, s_i\}$, $M_{i2} = \{r_i, u_i, s_i\}$, $M_{i3} = \{r_i, y_i\}$, $M_{i4} = \{x_i, s_i\}$, $M_{i5} = \{x_i, y_i\}$. The stable extensions of F can be separated into extensions containing c and others. For $i, j \in \{1 \dots 5\}$ the former are given as:

$$S_{ij} = \{c\} \cup M_{1i} \cup M_{2j}$$

If, on the other hand, $c \notin S$, one of a, b will be a member of S :

$$\begin{aligned}
S_1 &= \{a, r_1, r_2, v_1, v_2, s_1, s_2\} & S_4 &= \{a, r_1, r_2, y_1, v_2, s_2\} \\
S_2 &= \{b, r_1, r_2, u_1, u_2, s_1, s_2\} & S_5 &= \{b, r_1, u_1, x_2, s_1, s_2\} \\
S_3 &= \{a, r_1, r_2, v_1, y_2, s_1\} & S_6 &= \{b, r_2, x_1, u_2, s_1, s_2\}
\end{aligned}$$

Similarly to Example 5 we have that a and b share an implicit conflict for stable and thus stage semantics, as $stb(F) = stg(F) = \mathbb{S} = \{S_1 \dots S_6\} \cup \{S_{ij} \mid i, j \in \{1 \dots 5\}\}$. Again except for the implicit conflict between a and b all conflicts in F already are explicit, and the only other maximal conflict-free sets $S_a = \{a, r_1, r_2, y_1, y_2\}$ and $S_b = \{b, x_1, x_2, s_1, s_2\}$ are not stable extensions here.

◇

Theorem 27. *There are non-analytic AFs for stage semantics.*

Proof. Consider the AF $F = (A, R)$ from Example 7. We first show that F is non-analytic for stable semantics by assuming a contradicting analytic AF G of the same arguments and extensions. We will then use this observation to proceed similarly for stage semantics. As for any AF G with $stb(G) \neq \emptyset$ we have $stb(G) = stg(G)$, we will assume some AF G which is analytic for stage

semantics where $stb(G) = \emptyset$. In fact for both, stable and stage semantics, we show a slightly stronger result; for the given extension-set the conflict between a and b has to be implicit. For symmetry reasons, wlog. we assume $(a, b) \in R_G$. In what follows, we use the same naming schema for extensions as in Example 7.

For stable semantics we need $a \succ c$, since e.g. S_1 has to be a stable extension. From $S_{33} \in stb(G)$, $a \succ b$ by assumption and as observed $a \succ c$ we conclude $S_a \in stb(G)$, as $c \in S_{33}$ is allowed to attack only a and b . Thus if G is analytic for stable semantics then $stb(F) \neq stb(G)$.

We now turn to stage semantics and have the following observations:

- For $i \in \{1, 2\}$, due to maximal conflict-freeness and the given conflicts, we need explicit conflicts between s_i and y_i , r_i and x_i ($r_i, s_i \notin S_{55}$), between c and a , c and b ($a \notin S_{33}$, $b \notin S_{44}$), and between u_i and v_i ($v_i \notin S_{22}$). We will frequently make use of these necessities in the following.
- For the explicit conflict between s_1 and y_1 , we need $s_1 \succ y_1$ for otherwise $S_{55}^+ \subset S_{45}^+$. Similarly we conclude $s_2 \succ y_2$, $r_1 \succ x_1$ and $r_2 \succ x_2$;
- As the conflict between c and a is explicit, furthermore necessarily $c \succ a$ for otherwise (in case $a \succ c$ and $c \not\succ a$) $S_{11}^+ \subset S_1^+$;
- Now since u_i and v_i need to be in conflict we need $c \not\succ b$ for otherwise at least one of S_{ij} for $i, j \in \{1, 2\}$ becomes a stable extension. By conflict-implicitness hence $b \succ c$.
- From $c \succ a$, $r_1 \succ x_1$ and $s_1 \succ y_1$ we conclude $u_1 \succ v_1$ due to the danger of $S_{21}^+ \subset S_{11}^+$. Similarly $u_2 \succ v_2$.
- Since $c \succ a$ and $u_i \succ v_i$ furthermore we need $x_i \succ r_i$, $x_i \succ u_i$ and $x_i \succ v_i$, due to range comparison of M_{i4} and M_{i2} .
- By previous range observations we have to assume $b \not\succ a$ and $u_i \not\succ a$, for otherwise S_2 becomes a stable extension.
- But now $S_2^+ \subseteq S_b^+$, i.e. either we gain the unwanted extension S_b or we lose the desired extension S_2 .

□

To conclude this section we investigate the question of conditions such that ECC holds. We have mentioned earlier that every AF is quasi-analytic for naive semantics. This insight can be generalized as follows.

Proposition 28. *Let $\sigma \in \{stg, stb, sem, prf\}$. If for some AF F there exists an AF G such that $\sigma(F) = nai(G)$, then F is quasi-analytic for σ .*

Proof. Let F, G be AFs with $\sigma(F) = nai(G)$. We define the AF H with $A_H = A_F$ and $R_H = \{\langle a, b \rangle \mid (a, b) \in R_G, a, b \in \text{Args}_{\sigma(F)}\} \cup \{\langle a, x \rangle, \langle x, x \rangle \mid a \in A_F, x \notin \text{Args}_{\sigma(F)}\}$. As this AF H provides the same conflicts as the AF G for naive semantics, we deduce that also the maximal conflict-free sets are the same, $nai(H) = nai(G)$. By definition of H , for any $S \in nai(H)$ and $a \in A_F \setminus S$ we have $S \succrightarrow_H a$ and hence S is a stable extension of H . Finally observe that $stb(H) \subseteq \sigma(H) \subseteq nai(H)$ for any AF H , hence the result follows. \square

Another property that guarantees that ECC holds relies on the existence of what we call “identifying arguments”. We say that an AF F is *determined* for semantics σ if for every $S \in \sigma(F)$ there exists an $a \in S$ such that for $S' \in \sigma(F)$ we have that $a \in S'$ implies $S' = S$. In other words, every σ -extension contains an identifying argument in the sense that it does not occur in any other σ -extension. A simple necessary condition for an AF to be determined for σ is that the number of σ -extensions does not exceed the number of arguments.

Proposition 29. *Let $\sigma \in \{stb, prf, sem, stg\}$. Then, any AF F determined for σ is quasi-analytic for σ .*

Proof. Consider an AF F determined for σ and for each $S \in \sigma(F)$ let a_S be some fixed identifying argument. Now taking into account the sets $I = \{a_S \mid S \in \sigma(F)\}$ and $R_I = \{\langle a_S, a_{S'} \rangle \mid S, S' \in \sigma(F), S \neq S'\}$, clearly $\sigma((I, R_I)) = \{\{a_S\} \mid S \in \sigma(F)\}$. Furthermore let $O = A_F \setminus I$ be the remaining arguments of F and $R_O = \{\langle a, b \rangle \mid a, b \in O, (a, b) \notin \text{Pairs}_{\sigma(F)}\}$. We now define G as $A_G = A_F = O \cup I$ and $R_G = R_I \cup R_O \cup \{\langle a_S, b \rangle \mid S \in \sigma(F), b \in (O \setminus S)\}$. Observe that I forms a clique within G , a clique that is not attacked by arguments in O . For the SCC-splittable (cf. Definition 5) stable semantics we can evaluate G as follows:

1. Compute the extensions of the clique I .
2. For each such extension E of I build the AF $G_{|O \setminus E_G^+}$.
3. For each extension $E' \in \mathcal{GF}_{stb}(G_{|O \setminus E_G^+}, U_E^O)$ return $E \cup E'$.

Now the stable extensions of I are singletons $\{a_S\}$ for each $S \in stb(F)$. Moreover $G_{|O \setminus \{a_S\}_G^+} = (S \setminus \{a_S\}, \emptyset)$ and $U_{\{a_S\}}^O = S \setminus \{a_S\}$. We get $\mathcal{GF}_{stb}(G_{|O \setminus \{a_S\}_G^+}, U_{\{a_S\}}^O) = stb((S \setminus \{a_S\}, \emptyset)) = S \setminus \{a_S\}$. Hence $stb(G) = stb(F)$. The result for preferred semantics, which is also SCC-splittable, follows in the same way. For $\theta \in \{stg, sem\}$ we get $stb(G) = \theta(F)$ in the same way as above and since $\theta(F) \neq \emptyset$ it follows that $\theta(G) = stb(G) = \theta(F)$.

Finally observe that all conflicts in G for σ (among I , among O or between I and O) are explicit by definition. \square

6 Signatures

The last section dealt with the problem of making conflicts explicit without changing the set of arguments, or, in other words, finding an analytic AF with the same arguments that is equivalent

with respect to a given semantics. Abstaining from the condition that the set of arguments must be preserved, the focus is not on the given AF any more but on its sets of extensions. Given an extension-set \mathbb{S} and a semantics σ , the question is then whether the extension-set can be *analytically realized*, i.e. whether there is an analytic AF F having exactly $\sigma(F) = \mathbb{S}$, but imposing no restrictions on the arguments of F . We will deal with analytic realizability in Section 6.1. Likewise, Section 6.2 will be concerned with *compact realizability* where the AF realizing a given extension-set needs to be compact.

Prima facie this may seem similar to the concepts of analytic and compact argumentation frameworks studied in Section 3. However, relations between semantics from there do not carry over to realizability. For example we have seen in Theorem 2 that $CAF_{prf} \subset CAF_{nai}$, that is, every AF that is compact for preferred semantics is also compact for naive semantics and there exist AFs compact for naive but not compact for preferred semantics. In terms of compact realizability we will see that these semantics are related conversely, because compact realizability under naive semantics implies compact realizability under preferred semantics, but not vice versa (cf. Theorem 34).

Both analytic and compact realizability are restricted versions of the concept of (general) realizability studied in [23]. We first repeat the basic definitions and main results from there. Then we will for the analytic and compact scenario, respectively, first analyze the difference to general realizability and then deal with relations between the semantics under consideration.

Definition 9. An extension-set \mathbb{S} is called *realizable* under semantics σ if there is an AF F with $\sigma(F) = \mathbb{S}$. The *signature* of a semantics σ is defined as

$$\Sigma_\sigma = \{\sigma(F) \mid F \in AF_{\mathfrak{A}}\}.$$

The main results from [23] include $\Sigma_{nai} \subset \Sigma_{stg} = (\Sigma_{stb} \setminus \{\emptyset\}) \subset \Sigma_{prf} = \Sigma_{sem}$.

6.1 Analytic Signatures

In this section we deal with the restricted form of realizability, namely without the use of implicit conflicts.

Definition 10. An extension-set \mathbb{S} is called *analytically realizable* under semantics σ if there is some analytic AF $F \in XAF_\sigma$ with $\sigma(F) = \mathbb{S}$. The *analytic signature* (*x-signature*) Σ_σ^x of semantics σ consists of all extension-sets that are analytically realizable under σ :

$$\Sigma_\sigma^x = \{\sigma(F) \mid F \in XAF_\sigma\}$$

First of all note that every extension-set in the analytic signature of a semantics is also in the signature, i.e., $\Sigma_\sigma^x \subseteq \Sigma_\sigma$. In the following we will, for each semantics under consideration, either show that the relation is strict in the sense that certain extension-sets in Σ_σ are not analytically realizable or show that also $\Sigma_\sigma^x \supseteq \Sigma_\sigma$ holds, meaning that Σ_σ^x and Σ_σ coincide.

First we consider the relation between the signature Σ_{nai} and the analytic signature Σ_{nai}^x of naive semantics and formalize what we have already discussed in the introduction.

Theorem 30. *It holds that $\Sigma_{nai}^x = \Sigma_{nai}$.*

Proof. Consider some $\mathbb{S} \in \Sigma_{nai}$ with F being the AF realizing \mathbb{S} under naive semantics. It holds that a pair of arguments is contained in $Pairs_{\mathbb{S}}$ iff there is no attack between these arguments and none of them is self-attacking. Moreover, letting each self-attacking argument be attacked by all other arguments has no effect on the naive extensions. Hence the AF F' obtained from doing so has $nai(F') = \mathbb{S}$ and $F' \in XAF_{nai}$, therefore $\Sigma_{nai}^x = \Sigma_{nai}$. \square

Preferred and semi-stable semantics show strictly less expressiveness with respect to realizable extension-sets without implicit conflicts.

Theorem 31. *For $\sigma \in \{prf, sem\}$ it holds that $\Sigma_{\sigma}^x \subset \Sigma_{\sigma}$.*

Proof. Again consider the AF F in Figure 16 and let $\mathbb{S} = prf(F)$, which is given in Example 6. There we showed that there is no prf -analytic AF G having $\sigma(G) = \mathbb{S}$ and $A_G = A_F$. Here we can abstain from the last condition. So assume there is an AF $G' \in XAF_{prf}$ with $\sigma(G') = \mathbb{S}$. We know from Example 6 that there cannot be an attack between $S_a = \{a_1, a_2, a_3\} \in \mathbb{S}$ and u_2 and that in order for G' to be analytic $a_1 \succrightarrow_G x_2$. Moreover note that x_2 is the only possible attacker of u_2 among $Args_{\mathbb{S}}$. Finally, every additional argument $z \notin Args_{\mathbb{S}}$ in G' must be attacked by S_a since G' is prf -analytic and S_a must be admissible. This causes $S_a \cup \{u_2\}$ to be admissible in G' , hence S_a cannot be preferred in G' . Thus any AF realizing \mathbb{S} is non-analytic for preferred semantics or, in other words, $\mathbb{S} \in \Sigma_{prf} \setminus \Sigma_{prf}^x$.

Due to [26, 23] we know that $\Sigma_{sem} = \Sigma_{prf}$, hence there is an AF F' having $sem(F') = \mathbb{S}$. But when trying to analytically realize \mathbb{S} under sem , we make the same observations as above, meaning that $S_a \cup \{u_2\}$ is necessarily admissible, a contradiction to S_a being semi-stable. Hence also $\mathbb{S} \in \Sigma_{sem} \setminus \Sigma_{sem}^x$. \square

We now turn to stable and stage semantics. In contrast to preferred and semi-stable semantics, we will see that the use of additional arguments allows us to make each implicit conflict explicit. Therefore the analytic signature coincides with the signature for stable and stage semantics.

The following proposition shows that one additional argument allows, together with an appropriate modification of the attack relation, to make any single implicit conflict explicit.

Proposition 32. *For stable semantics and some AF F , if there is an implicit conflict between a and b , then there is an AF G with $|A_G| = |A_F| + 1$, $R_G \supseteq R_F$, $(a, b) \in R_G$ and $stb(G) = stb(F)$ and all implicit conflicts in G are implicit conflicts in F as well.*

Proof. Let F be an arbitrary AF with an implicit conflict between two arguments a and b . We define $R' = R_F \cup \{(a, b)\}$. Observe that $F' = (A_F, R')$ has the same and possibly more stable extensions as compared to F . By construction of F' , any $S \in stb(F') \setminus stb(F)$ has $a \in S$ and $S \not\rightarrow_F b$. We collect the arguments of these unwanted extensions in $A_a = Args_{(stb(F') \setminus stb(F))}$ and observe that $A_a \not\rightarrow_F b$. Now define the AF G with $A_G = A_F \cup \{x\}$ and

$$R_G = R' \cup \{(x, x)\} \cup \{(x, v) \mid v \in A_a\} \cup \{(u, x) \mid u \in A_F \setminus A_a\}.$$




Figure 19: A Venn-Diagram illustrating analytic signatures of stable, semi-stable, stage and preferred semantics.

First note that obviously $|A_G| = |A_F| + 1$, $R_G \supseteq R_F$, and $(a, b) \in R_G$. Moreover, since the new argument x attacks or is attacked by every other argument, G does not introduce any further implicit conflicts compared to F . It remains to show that $stb(G) = stb(F)$. Let $S' \in stb(F)$ and assume that $b \in S'$. As by assumption b and a do not occur together in any stable extension of F , we know that $b \succrightarrow_G x$ and thus $S' \in stb(G)$. On the other hand assume that $b \notin S'$. Then we have some $c \in S'$ with $c \succrightarrow_F b$. If $S' \notin stb(G)$, then only because $S' \not\precrightarrow_G x$, hence $S' \subseteq A_a$, a contradiction to $A_a \not\precrightarrow_F b$. Therefore $S' \in stb(G)$. Now assume there is some $S \in stb(G)$ with $S \notin stb(F)$. By the construction of G this S must be among $stb(F') \setminus stb(F)$. However, we then have $S \not\precrightarrow_G x$, a contradiction to $S \in stb(G)$, concluding the proof for $stb(F) = stb(G)$. \square

Now we can show that analytic and general signatures coincide for stable and stage semantics.

Theorem 33. For $\sigma \in \{stb, stg\}$ it holds that $\Sigma_\sigma^x = \Sigma_\sigma$.

Proof. We consider as special case $stb(F) = \emptyset$ or $stg(F) = \{\emptyset\}$ where by definition the AF $F = (\{x\}, \{(x, x)\})$ serves as analytic witness. Let $\mathbb{S} \in \Sigma_\sigma$, i.e., there is some AF F with $\sigma(F) = \mathbb{S}$. As by definition any AF F is finite we can have at most finitely many implicit conflicts for semantics $\sigma \in \{stb, stg\}$. Each of them can be removed by repeated application of Proposition 32 for $\sigma = stb$. Hence there is an analytic AF F' with $\sigma(F') = \mathbb{S}$, meaning that $\mathbb{S} \in \Sigma_{stb}^x$. For $\sigma = stg$ semantics we know from [26] that there is an AF G with $stb(G) = stg(G) = \mathbb{S}$. Now, again, we can remove all implicit conflicts and end up with the stg -analytic AF G' with $stg(G') = \mathbb{S}$. Hence $\mathbb{S} \in \Sigma_{stg}^x$. \square

So far we have compared general signatures and analytic signatures for the semantics under consideration. We have seen that preferred and semi-stable semantics can realize strictly more when allowing the use of implicit conflicts, while this is not the case for stable and stage semantics.

In the following we relate the analytic signatures of naive, stable, preferred, stage and semi-stable semantics to each other. For general signatures it was shown in [23] that $\Sigma_{nai} \subset \Sigma_{stg} = (\Sigma_{stb} \setminus \{\emptyset\}) \subset \Sigma_{sem} = \Sigma_{prf}$. In the analytic case preferred and semi-stable signatures do not coincide anymore.

Theorem 34. In accordance with Figure 19, it holds that:

1. $\Sigma_{nai}^x \subset \Sigma_\sigma^x$ for $\sigma \in \{stb, stg, sem, prf\}$;
2. $\Sigma_{stb}^x \setminus \{\emptyset\} = \Sigma_{stg}^x$;




Figure 20: The AF witnessing $\Sigma_{nai}^x \subset \Sigma_\sigma^x$ for $\sigma \in \{stb, sem, stg, prf\}$.

3. $\Sigma_{stg}^x \subset \Sigma_{sem}^x$;

4. $\Sigma_{sem}^x \subset \Sigma_{prf}^x$

Proof. (1) First recall from [23] that for a given $\mathbb{S} \in \Sigma_{nai}^x$, the canonic AF F where $A_F = \text{Args}_{\mathbb{S}}$ and $R_F = (A_F \times A_F) \setminus \text{Pairs}_{\mathbb{S}}$ gives $\mathbb{S} = nai(F) = \sigma(F)$, and F is analytic for σ , thus $\Sigma_{nai}^x \subseteq \Sigma_\sigma^x$.

Further consider the AF F where $A_F = \{x_1, x_2, x_3, y_1, y_2, y_3\}$ and $R_F = \{(x_i, x_j), (x_i, y_i) \mid i, j \in \{1, 2, 3\}\}$, cf. Figure 20, the AF featured in [26] to show that $\sigma(F) = \{\{x_1, y_2, y_3\}, \{x_2, y_1, y_3\}, \{x_3, y_1, y_2\}\}$ can not be realized under naive semantics. With the fact that this AF is analytic for σ we obtain $\Sigma_{nai}^x \not\subseteq \Sigma_\sigma^x$, hence $\Sigma_{nai}^x \subset \Sigma_\sigma^x$.

(2) Considering $\Sigma_{stb} \setminus \{\emptyset\} = \Sigma_{stg}$ [23] and Theorem 33 we obtain $\Sigma_{stb}^x \setminus \{\emptyset\} = \Sigma_{stg}^x$.

(3) For $\mathbb{S} \in \Sigma_{stg}^x$ with $\mathbb{S} \neq \{\emptyset\}$ we know from (2) that there is an analytic AF F with $stb(F) = \mathbb{S}$. Now as $\mathbb{S} \neq \{\emptyset\}$ also $sem(F) = \mathbb{S}$, hence $\mathbb{S} \in \Sigma_{sem}^x$. As obviously $\{\emptyset\} \in \Sigma_{sem}^x$ (witnessed by $(\{x\}, \{(x, x)\})$), we get $\Sigma_{stg}^x \subseteq \Sigma_{sem}^x$.

For properness take a look at the AF F from Figure 8, which, as discussed in the proof of Theorem 4, is analytic for semi-stable semantics. Now consider, for instance, $S = \{r_1, a_1, v_1, b_1\} \in sem(F)$. Observe that $c_i \notin S$ for $i \in \{1, 2, 3\}$ besides c_i not being in conflict with S . If there was an AF $F' \in XAF_{stg}$ with $stg(F') = sem(F)$, then there can not be any attack between S and c_i in F . But then $S \cup \{c_1\}$ is conflict-free in F' and its range is strictly larger than the range of S . Thus $sem(F) \notin \Sigma_{stg}^x$ and therefore $\Sigma_{stg}^x \subset \Sigma_{sem}^x$.

(4) For the last part of the theorem recall that the exact translation for $sem \rightarrow prf$ from [26] does not add any implicit conflicts between arguments from the original AF. In more detail for a given (analytic) AF F we add one self-attacking argument x_S for any unwanted preferred extension $S \in prf(F) \setminus sem(F)$, and further add attacks (x_S, a) for $a \in S$ and (b, x_S) for $b \in A_F \setminus S$. Thus the only implicit conflicts generated by this translation are conflicts between new and self-attacking arguments. However we can simply make such conflicts explicit by adding attacks between any self-attacking arguments, which does not affect preferred semantics, and hence $\Sigma_{sem}^x \subseteq \Sigma_{prf}^x$.

Now, for properness, consider the *prf*-analytic AF F from Figure 9. Define a cyclic successor functions with $s(1) = 2, s(2) = 3, s(3) = 1$ and $s(4) = 5, s(5) = 6, s(6) = 4$. We have as preferred extensions $prf(F) = \mathbb{S}_0 \cup \mathbb{S}_1 \cup \mathbb{S}_2$ with

$$\mathbb{S}_0 = \{\{x_i, y_j, z_{s(i)}, z_{s(j)}\} \mid i \in \{1, 2, 3\}, j \in \{4, 5, 6\} \text{ or } i \in \{4, 5, 6\}, j \in \{1, 2, 3\}\}$$

$$\mathbb{S}_1 = \{\{x_i, y_i, z_{s(i)}\} \mid i \in \{1, 2, 3, 4, 5, 6\}\}$$

$$\mathbb{S}_2 = \{\{x_i, y_{s(i)}, z_{s(s(i))}\}, \{x_{s(i)}, y_i, z_{s(s(i))}\} \mid i \in \{1, 2, 3, 4, 5, 6\}\}$$

Assume that there is some $G \in XAF_{sem}$ with $sem(G) = prf(F)$. We take a look at \mathbb{S}_1 and more specifically $\{x_1, y_1, z_2\} \in \mathbb{S}_1$. Now we need an explicit conflict between x_1 and x_4 , but in the selected set only x_1 can possibly defend against this attack, hence $(x_1, x_4) \in R_G$. The same argument works for x_1 and x_3 as well as z_2 and z_3 , meaning that also $(x_1, x_3), (z_2, z_3) \in R_G$. For symmetry reasons $\{(x_i, x_j), (x_j, x_i), (y_i, y_j), (y_j, y_i) \mid i \in \{1, 2, 3\}, j \in \{4, 5, 6\}\} \cup \{(x_{s(i)}, x_i), (z_i, z_{s(i)}) \mid i \in \{1, 2 \dots 6\}\} \subseteq R_G$.

We take a look at \mathbb{S}_2 and more specifically $\{x_1, y_2, z_3\} \in \mathbb{S}_2$. As there should be an explicit conflict between x_1 and x_2 with only x_1 possibly defending this extension against x_2 we need $(x_1, x_2) \in R_G$. Further as in this set only y_2 and z_3 can possibly attack z_2 we have the set $\{y_2, z_3\}$ attacking z_2 . For symmetry reasons $\{(x_i, x_{s(i)}), (y_i, y_{s(i)}) \mid i \in \{1, 2 \dots 6\}\} \subseteq R_G$ and each set $\{x_i, z_{s(i)}\}, \{y_i, z_{s(i)}\}$ for $i \in \{1, 2 \dots 6\}$ attacks z_i .

Finally we take a look at \mathbb{S}_0 and more specifically the set $S = \{x_1, y_4, z_2, z_5\} \in \mathbb{S}_0$. Since S necessarily is an admissible extension in an analytic AF we have that S attacks all rejected arguments. By the above observations we now have that S even attacks all arguments not being member of S in G , which means that S is a stable extension and stable semantics and semi-stable semantics thus coincide on G . But then, with $T = \{x_1, y_1, z_2\} \in \mathbb{S}_1$ not being in conflict with for instance z_4 we have that T can not be a stable or semi-stable extension in G . We finally conclude that indeed $prf(F) \notin \Sigma_{sem}^x$ and thus $\Sigma_{sem}^x \subset \Sigma_{prf}^x$. \square

6.2 Compact Signatures

We now turn to the issue of realizing extension-sets without the use of rejected arguments.

Definition 11. An extension-set \mathbb{S} is called *compactly realizable* under semantics σ if there is some compact AF $F \in CAF_\sigma$ with $\sigma(F) = \mathbb{S}$. The *compact signature* (*c-signature*) Σ_σ^c of semantics σ consists of all extension-sets that are compactly realizable under σ :

$$\Sigma_\sigma^c = \{\sigma(F) \mid F \in CAF_\sigma\}.$$

It is clear that $\Sigma_\sigma^c \subseteq \Sigma_\sigma$ holds for any semantics. The following theorem repeats the equality of compact and general signatures for naive semantics discussed in the introduction, and shows a \subseteq -relation for all other semantics.

Proposition 35. *It holds that*

1. $\Sigma_{nai}^c = \Sigma_{nai}$, and
2. $\Sigma_\sigma^c \subset \Sigma_\sigma$ for $\sigma \in \{stb, stg, sem, prf\}$.

Proof. nai: Consider some $\mathbb{S} \in \Sigma_{nai}$ with F being the AF realizing \mathbb{S} under naive semantics. It holds that an argument is contained in $Args_{\mathbb{S}}$ iff it is not self-attacking. Moreover removing any self-attacking argument together with its associated attacks has no effect on the naive extensions. Hence the AF F' obtained from removing all self-attacking arguments together with their associated attacks has $nai(F') = \mathbb{S}$ and $F' \in CAF_{nai}$, therefore $\Sigma_{nai}^c = \Sigma_{nai}$.




Figure 21: A Venn-Diagram illustrating compact signatures of stable, semi-stable, stage and preferred semantics.

By definition we have $\Sigma_\sigma^c \subseteq \Sigma_\sigma$. It remains to show that $\Sigma_\sigma^c \neq \Sigma_\sigma$ for $\sigma \in \{stb, stg, sem, prf\}$.

stb, stg: Consider the extension-set $\mathbb{S} = \{\{a, b, c\}, \{a, b, c'\}, \{a, b', c\}, \{a, b', c'\}, \{a', b, c\}, \{a', b, c'\}, \{a', b', c\}\}$ from the example in the introduction. We have seen that \mathbb{S} is realized under *stb* and *stg* by the AF F_1 from the introduction. Assume there is an AF $F = (Args_{\mathbb{S}}, R)$ realizing \mathbb{S} under *stb* or *stg*. Inspecting $Pairs_{\mathbb{S}}$ we infer that $R \subseteq \{(a, a'), (a', a), (b, b'), (b', b), (c, c'), (c', c)\}$. Note that, for any remaining choice of R , $stb(F) = stg(F)$. Now for $\{a, b, c\} \in stb(F)$ we need $(a, a'), (b, b'), (c, c') \in R$. On the other hand, for $\{a', b, c\}, \{a, b', c\}, \{a, b, c'\} \in stb(F)$ we need $(a', a), (b', b), (c', c) \in R$. But then also $\{a', b', c'\} \in stb(F)$. Hence $\mathbb{S} \notin \Sigma_{stb}^c$ and also $\mathbb{S} \notin \Sigma_{stg}^c$, witnessing $\Sigma_{stb}^c \subset \Sigma_{stb}$ and $\Sigma_{stg}^c \subset \Sigma_{stg}$.

prf, sem: Let $\sigma \in \{prf, sem\}$ and consider $\mathbb{S} = \{\{a, b\}, \{a, d, e\}, \{b, c, e\}\}$. $\mathbb{S} \in \Sigma_\sigma$ holds since Figure 3 shows an AF (with additional arguments) realizing \mathbb{S} as its semi-stable and preferred extensions. Now suppose there exists an AF $F = (Args_{\mathbb{S}}, R)$ such that $\sigma(F) = \mathbb{S}$. Since $\{a, d, e\}, \{b, c, e\} \in \mathbb{S}$, it is clear that R must not contain an edge involving e . But then, e is contained in each $E \in \sigma(F)$. It follows that $\sigma(F) \neq \mathbb{S}$ and therefore $\mathbb{S} \notin \Sigma_\sigma^c$. \square

In the following we relate the compact signatures of the semantics under consideration to each other. Recall that for general signatures it holds that $\Sigma_{nai} \subset \Sigma_{stg} = (\Sigma_{stb} \setminus \{\emptyset\}) \subset \Sigma_{sem} = \Sigma_{prf}$ [23]. Similarly, but not equivalently though, we have $\Sigma_{nai}^x \subset \Sigma_{stg}^x = (\Sigma_{stb}^x \setminus \{\emptyset\}) \subset \Sigma_{sem}^x \subset \Sigma_{prf}^x$ for analytic signatures (cf. Theorem 34). This picture changes when considering the relationships between compact signatures. Figure 21 depicts the relations between compact signatures which we will show in the next theorem. The dashed areas represent particular intersections for which the question of existence of extension-sets has to be left open. Also notice that stable semantics cannot realize the empty extension set within compact AFs.

Theorem 36. *In accordance with Figure 21, it holds that:*

1. $\Sigma_{nai}^c \subset \Sigma_\sigma^c$ for $\sigma \in \{stb, stg, sem, prf\}$;
2. $\Sigma_{stb}^c \subset \Sigma_\sigma^c$ for $\sigma \in \{stg, sem\}$;




Figure 22: AF showing $\Sigma_{prf}^c \setminus (\Sigma_{stb}^c \cup \Sigma_{sem}^c \cup \Sigma_{stg}^c) \neq \emptyset$.

3. $\Sigma_{prf}^c \setminus (\Sigma_{stb}^c \cup \Sigma_{sem}^c \cup \Sigma_{stg}^c) \neq \emptyset$;
4. $\Sigma_{stg}^c \setminus (\Sigma_{stb}^c \cup \Sigma_{prf}^c \cup \Sigma_{sem}^c) \neq \emptyset$;
5. $\Sigma_{stb}^c \setminus \Sigma_{prf}^c \neq \emptyset$;
6. $(\Sigma_{prf}^c \cap \Sigma_{sem}^c) \setminus (\Sigma_{stb}^c \cup \Sigma_{stg}^c) \neq \emptyset$;
7. $\Sigma_{sem}^c \setminus (\Sigma_{stb}^c \cup \Sigma_{prf}^c \cup \Sigma_{stg}^c) \neq \emptyset$.

Proof. (1) First recall that for a given $\mathbb{S} \in \Sigma_{nai}^c$, the canonic AF F where $A_F = \text{Args}_{\mathbb{S}}$ and $R_F = (A_F \times A_F) \setminus \text{Pairs}_{\mathbb{S}}$ gives $\mathbb{S} = \text{nai}(F) = \sigma(F)$, and F is compact for σ , thus $\Sigma_{nai}^c \subseteq \Sigma_{\sigma}^x$. Moreover, the AF depicted in Figure 20 is compact for $\sigma \in \{stb, stg, sem, prf\}$, but $\sigma(F)$ can not, as discussed in the proof of Theorem 34, be realized under the naive semantics. Hence $\Sigma_{nai}^c \subset \Sigma_{\sigma}^x$.

(2) $\Sigma_{stb}^c \subseteq \Sigma_{\sigma}^c$ for $\sigma \in \{stg, sem\}$, follows from the fact that $stg(F) = sem(F) = stb(F)$ for every $F \in \text{CAF}_{stb}$ [13]. Properness is by (4) and (7), to be shown in the remainder of this proof.

In the following we provide, as part of the proof, examples witnessing the remaining statements. The general procedure is as follows: Let $\sigma_1, \dots, \sigma_n$ and τ_1, \dots, τ_m be semantics. To show that $(\bigcap_{1 \leq i \leq n} \Sigma_{\sigma_i}^c) \setminus (\bigcup_{1 \leq j \leq m} \Sigma_{\tau_j}^c) \neq \emptyset$ holds, we fix some extension-set \mathbb{S} , provide an AF F with $\sigma_i(F) = \mathbb{S}$ for all $i \in \{1, \dots, n\}$, and show that \mathbb{S} is not compactly realizable under any of the semantics τ_1, \dots, τ_m .

We begin by showing (3) $\Sigma_{prf}^c \setminus (\Sigma_{stb}^c \cup \Sigma_{sem}^c \cup \Sigma_{stg}^c) \neq \emptyset$.

Example 8. Consider the extension-set $\mathbb{S} = \{\{a, b\}, \{a, x_i, s_i\}, \{b, y_i, s_i\}, \{x_i, y_i, s_i\} \mid 1 \leq i \leq 3\}$ and observe that the AF F depicted in Figure 22 has exactly $prf(F) = \mathbb{S}$. Since F is compact for prf we have $\mathbb{S} \in \Sigma_{prf}^c$. Let $\sigma \in \{stb, stg, sem\}$. We show that $\mathbb{S} \notin \Sigma_{\sigma}^c$. Towards a contradiction assume that there is an AF G with $A_G = \text{Args}_{\mathbb{S}}$ and $\sigma(G) = \mathbb{S}$. First observe that there cannot be any attack between a and b on the one hand and s_1, s_2 , and s_3 on the other. For $\sigma = stb$ we have a contradiction to $\sigma(G) = \mathbb{S}$ since $s_1, s_2, s_3 \notin \{a, b\}_G^+$. Also for $\sigma = stg$ we have a contradiction since for each i , $\{a, b, s_i\}$ is conflict-free and $\{a, b, s_i\}_G^+ \supset \{a, b\}_G^+$, hence $\{a, b\} \notin stg(G)$. Finally consider $\sigma = sem$. Let $S = \{a, x_1, s_1\}$, $T = \{x_1, y_1, s_1\}$. If there was no attack between a and y_1 then $S \cup T$ would be conflict-free and therefore $S, T \notin \sigma(G)$.




Figure 23: A directed cycle of nine arguments.




Figure 24: AF showing $\Sigma_{stb}^c \setminus \Sigma_{prf}^c \neq \emptyset$.

Since each of T and $\{a, b\}$ must defend itself, necessarily both $(y_1, a), (a, y_1) \in R_G$. By symmetry we get $\{\langle a, y_i \rangle, \langle b, x_i \rangle \mid 1 \leq i \leq 3\} \subseteq R_G$. Now in order to have $\{a, b\} \in \text{sem}(G)$, no s_i can be defended by $\{a, b\}$, hence each s_i must have an attacker that is not attacked by $\{a, b\}$ and s_i . Hence wlog. $\{(s_1, s_2), (s_2, s_3), (s_3, s_1)\} \subseteq R_G$. Now observe that S has to defend s_1 from s_3 , therefore $(x_1, s_3) \in R_G$. So far we have $S_G^+ \supseteq (\text{ArgS}_{\mathbb{S}} \setminus \{x_2, x_3\})$. S has to attack both x_2 and x_3 since otherwise either S would not defend itself or at least one of $S \cup \{x_2\}$ and $S \cup \{x_3\}$ would be admissible and have greater range than S . But now $S_G^+ = \text{ArgS}_{\mathbb{S}} \supset \{a, b\}_G^+$, a contradiction to $\{a, b\} \in \text{sem}(G)$. \diamond

We continue with (4) $\Sigma_{stg}^c \setminus (\Sigma_{stb}^c \cup \Sigma_{prf}^c \cup \Sigma_{sem}^c) \neq \emptyset$.

Example 9. Let \oplus such that $a \oplus b = (a + b) \bmod 9$. Consider the AF $F = (\{a_0, \dots, a_8\}, \{(a_i, a_{i \oplus 1}) \mid 0 \leq i < 9\})$, i.e. the directed cycle of nine arguments. We get $\text{stg}(F) = \{\{a_i, a_{i \oplus 2}, a_{i \oplus 4}, a_{i \oplus 6}\} \mid 0 \leq i < 9\}$. Now assume this extension-set is compactly realizable under stable, preferred or semi-stable semantics, i.e. there is some G with $\sigma(G) = \text{stg}(F)$ ($\sigma \in \{stb, prf, sem\}$) and $A_G = A_F$. Since a_i and a_j occur together in some stage extension of F for all i, j with $i \oplus 1 \neq j$ and $i \neq j \oplus 1$, the only possible attacks in G are (a_i, a_j) with $i \oplus 1 = j$ or $i = j \oplus 1$. Now let $S_i = \{a_i, a_{i \oplus 2}, a_{i \oplus 4}, a_{i \oplus 6}\}$. In order to have $S_i \in \sigma(G)$, a_i has to attack $a_{i \oplus 8}$ and $a_{i \oplus 6}$ has to attack $a_{i \oplus 7}$, first for S_i to be maximal and second to be defended. Hence $R_G = \{\langle a_i, a_{i \oplus 1} \rangle \mid 0 \leq i < 9\}$ and $\sigma(G) = \text{stg}(F) \cup \{a_i, a_{i \oplus 3}, a_{i \oplus 6} \mid 0 \leq i < 3\}$, showing that there is no compact AF realizing $\text{stg}(F)$ under σ . \diamond

The following example witnesses that (5) $\Sigma_{stb}^c \setminus \Sigma_{prf}^c \neq \emptyset$.

Example 10. Consider stable semantics for the AF F depicted in Figure 24 and let $\mathbb{S} = \text{stb}(F)$ be its extension-set. Observe that neither $\{a, b, c\}$ nor any superset is a stable extension.




Figure 25: AF showing $(\Sigma_{prf}^c \cap \Sigma_{sem}^c) \setminus (\Sigma_{stb}^c \cup \Sigma_{stg}^c) \neq \emptyset$.

Assume there exists some AF G compactly realizing \mathbb{S} under preferred semantics, i.e. $prf(G) = \mathbb{S}$ and $A_G = Arg_{\mathbb{S}}$. One can check that F is analytic for stable semantics, i.e. for the AF G there can only be attacks between arguments being linked in Figure 24.

Consider the extension $S = \{b, c, x_1, s_1\} \in \mathbb{S}$. For $S \in prf(G)$ there are two possible reasons for $a \notin S$. Either a is in conflict with S or a is not defended by S . Assume a not to be defended by S . Then $x_2 \succ_G a$ and $x_1 \not\succeq_G x_2$ and $s_1 \not\succeq_G x_2$. But then $x_2 \notin S$ defends itself, hence S cannot be a maximal admissible set in G . It follows that a is in conflict with S , the only possibility being a conflict with x_1 , hence $x_1 \succ_G a$ ($a \succ_G x_1$ is not sufficient since no other argument in S can defend x_1 against a). Considering $\{a, y_1, z_1, s_2\} \in \mathbb{S}$, none of y_1 , z_1 , and s_2 can defend a against x_1 , hence also $a \succ_G x_1$.

Similarly, one can justify the existence of symmetric attacks between a and x_2 , b and y_i , and c and z_i ($i \in \{1, 2\}$). Therefore the set $\{a, b, c\}$ is admissible in G , hence there must be some $S' \in prf(G)$ with $S' \supseteq \{a, b, c\}$, a contradiction to \mathbb{S} being compactly realizable under the preferred semantics. \diamond

We proceed with an example showing that (6) $(\Sigma_{prf}^c \cap \Sigma_{sem}^c) \setminus (\Sigma_{stb}^c \cup \Sigma_{stg}^c) \neq \emptyset$.

Example 11. Consider the AF F from Figure 25. We have $\mathbb{S} = sem(F) = prf(F) = \{\{v_i, y_j, r_i, s_j\}, \{w_i, x_j, t_i, s_j\}, \{v_i, w_j, r_i, t_j\} \mid 1 \leq i, j \leq 3\}$. For $\sigma = stg$ or $\sigma = stb$, assume there is an AF G with $\sigma(G) = \mathbb{S}$ and $A_G = Arg_{\mathbb{S}}$. First note that for all $i, j \in \{1, 2, 3\}$ each pair $\{v_i, s_j\}, \{w_i, s_j\}, \{r_i, s_j\}, \{t_i, s_j\}$ is contained in some element of \mathbb{S} , hence there cannot be an attack between any of these pairs in G . Now let $S = \{v_i, w_j, r_i, t_j\}$ for some $i, j \in \{1, \dots, 3\}$. We have $S_G^+ \subseteq A_G \setminus \{s_1, s_2, s_3\}$, hence S cannot be a stable extension of G . Moreover, since G must be self-loop-free, $S \cup \{s_k\}$ with $1 \leq k \leq 3$ is conflict-free and obviously has a bigger range than S . Therefore S cannot be a stage extension in G . \diamond

For (7) we will make use of the following lemma, which might be of interest on its own.

Lemma 37. Let $\sigma, \tau \in \{stb, prf, sem, stg\}$ and F, G be τ -compact AFs such that $\tau(F) \notin \Sigma_{\sigma}^c$ and $A_F \cap A_G = \emptyset$. It holds that $\tau(F \cup G) \notin \Sigma_{\sigma}^c$.

Proof. Assume there is some compact AF H such that $\sigma(H) = \tau(F \cup G)$. Since $A_F \cap A_G = \emptyset$, it follows that $\tau(F \cup G) = \tau(F) \times \tau(G)$. Due to compactness every argument $a \in A_F$ occurs together with every argument $b \in A_G$ in some τ -extension of $F \cup G$, meaning that H cannot contain any attack between a and b . Hence $\sigma(H) = \sigma(H_1) \times \sigma(H_2)$ with $A_{H_1} = A_F$ and $A_{H_2} = A_G$. Therefore it must hold that $\sigma(H_1) = \tau(F)$, a contradiction to the assumption that $\tau(F) \notin \Sigma_G^c$. \square

Now we get (7) $\Sigma_{sem}^c \setminus (\Sigma_{stb}^c \cup \Sigma_{prf}^c \cup \Sigma_{stg}^c) \neq \emptyset$ as follows: Let $F = F_1 \cup F_2$ where F_1 is the AF in Figure 24 and F_2 is the AF in Figure 25 (observe that for $A_{F_1} \cap A_{F_2} = \emptyset$ some renaming is necessary). From $sem(F_1) \notin \Sigma_{prf}^c$ (see Example 10) we get $sem(F) = (sem(F_1) \times sem(F_2)) \notin \Sigma_{prf}^c$ by Lemma 37. In the same way $sem(F) \notin \Sigma_{stb}^c \cup \Sigma_{stg}^c$ follows from $sem(F_2) \notin \Sigma_{stb}^c \cup \Sigma_{stg}^c$ (see Example 11).

This concludes the proof of Theorem 36. \square

Comparing the insights obtained from Theorem 36 with the results on expressiveness of semantics in [23] we observe notable differences depending on whether rejected arguments are allowed or not. When allowing rejected arguments (as utilized in [23]), preferred and semi-stable semantics are equally expressive and at the same time strictly more expressive than stable and stage semantics. As we have seen, this does not carry over to the compact setting where, with the exception of $\Sigma_{stb}^c \subset \Sigma_{sem}^c$ and $\Sigma_{stb}^c \subset \Sigma_{stg}^c$, signatures become incomparable.

What remains an open issue is the existence of extension-sets lying in the intersection between Σ_{prf}^c (resp. Σ_{sem}^c) and Σ_{stg}^c but outside of Σ_{stb}^c (see Venn-diagram in Figure 21). We approach this issue in the remainder of this section.

Lemma 38. *In self-attack free AFs every stage extension that is admissible is also stable.*

Proof. Take some AF $F = (A, R)$, and some admissible stage extension S , $S \in stg(F)$, $S \in adm(F)$ as given. Suppose there is some argument that is not in the range of S , i.e. $a \in A \setminus S_F^+$. Then by admissibility a cannot attack S , by assumption S does not attack a . Consider that any stage extension is maximal conflict-free, thus for $a \notin S$ we in fact would need $(a, a) \in R$. It follows that there is no such argument a and thus $S_F^+ = A$. Hence $S \in stb(F)$. \square

Proposition 39. *Let $\sigma \in \{sem, prf\}$ and F, G be σ -compact AFs such that $stg(F) = \sigma(G)$. If $stg(F) \notin \Sigma_{stb}^c$ then*

1. $F \neq G$, and
2. G is non-analytic for σ .

Proof. Assume that $F = G$. But then, as by assumption $stg(F) = \sigma(F)$, by Lemma 38 also $\sigma(F) = stb(F)$, a contradiction to the assumption that $stg(F) \notin \Sigma_{stb}^c$. Therefore $F \neq G$.

For a contradiction, wlog. assume G to be σ -analytic (for any quasi-analytic H there is some corresponding analytic G). Observe that for stage extensions $S \in stg(F)$ and any argument $a \in A \setminus S$ it holds that either there is an explicit conflict between S and a in F , or a is self-attacking in F , for otherwise S_F^+ would not be maximal. With $stg(F) = \sigma(G)$ and G being analytic for the admissibility based semantics σ this means that $S \succrightarrow_G a$, i.e. $S_G^+ = A$. With all σ -extensions becoming stb -extensions and the fact that $stb(F) \subseteq \sigma(F)$ for any F , we derive a contradiction to the initial statement: $stb(G) = stg(F)$. \square

Assume that for $\sigma \in \{prf, stg\}$ there exists an extension-set $\mathbb{S} \in (\Sigma_\sigma^c \cap \Sigma_{stg}^c) \setminus \Sigma_{stb}^c$. Now Proposition 39 says that \mathbb{S} is compactly realized by different AFs under σ and stg , i.e. $stg(F) = \mathbb{S}$ and $\sigma(G) = \mathbb{S}$ with $F \neq G$. Moreover, G is non-analytic. Recent investigations encourage us to conjecture that such extension-sets do not exist (we already know that $\Sigma_{stb}^c \subseteq \Sigma_{sem}^c \cap \Sigma_{stg}^c$ (cf. Theorem 36.2) and $\Sigma_{stb}^c \setminus \Sigma_{prf}^c \neq \emptyset$ (cf. Theorem 36.5) hold):


Conjecture. It holds that $\Sigma_{prf}^c \cap \Sigma_{stg}^c \subset \Sigma_{stb}^c$ and $\Sigma_{sem}^c \cap \Sigma_{stg}^c = \Sigma_{stb}^c$.

6.3 Numbers of Extensions in Compact Frameworks

In the previous section we have related the semantics under consideration with respect to their capabilities in compactly realizing extension-sets. The concrete problem which was tackled in [23] is deciding, given an extension-set \mathbb{S} , whether \mathbb{S} is realizable. For compact realizability this is, in general, a hard problem; that is, we have no reason to believe that we can do any better than guessing a compact AF and checking whether its extension-set coincides with \mathbb{S} . Nevertheless, in this section we provide a number of shortcuts to detect non-compactness. By “shortcut”, we mean a property of the given extension-set \mathbb{S} that is easily computable (preferably in polynomial time) and lets us (sometimes) give a definitive answer to the decision problem. These shortcuts are related to numerical aspects of argumentation frameworks, some of which have been studied in graph theory.

Among the most basic properties that are necessary for compact realizability, we find numerical aspects like possible cardinalities of σ -extension-sets.

Example 12. Consider the following AF F_2 :



Let us determine the stable extensions of F_2 . Clearly, taking one a_i , one b_i and one c_i yields a conflict-free set that is also stable as long as it attacks z . Thus from the $3 \cdot 2 \cdot 3 = 18$ combinations, only one (the set $\{a_1, b_1, c_2\}$) is not stable, whence F_2 has $18 - 1 = 17$ stable extensions. We note that this AF is not compact since z occurs in none of the extensions. Is there an equivalent stb -compact AF? The results of this section will provide us with a negative answer. \diamond

(author?) ((year?)) have shown that there is a correspondence between the maximal number of stable extensions in argumentation frameworks and the maximal number of maximal independent sets in undirected graphs [33]. Recently, the result was generalized to further semantics [23,

Proposition 11 and Theorem 5].⁸ To set the scene for the subsequent results building upon it, we recall the result below (Theorem 40). For any natural number n we define:⁹

$$\sigma_{\max}(n) = \max \{ |\sigma(F)| \mid F \in AF_{\mathfrak{A}}, |A_F| = n \}$$

$\sigma_{\max}(n)$ returns the maximal number of σ -extensions among all AFs with n arguments. Surprisingly, there is a closed expression for σ_{\max} .

Theorem 40. *The function $\sigma_{\max}(n) : \mathbb{N} \rightarrow \mathbb{N}$ is given by*

$$\sigma_{\max}(n) = \begin{cases} 1, & \text{if } n = 0 \text{ or } n = 1, \\ 3^s, & \text{if } n \geq 2 \text{ and } n = 3s, \\ 4 \cdot 3^{s-1}, & \text{if } n \geq 2 \text{ and } n = 3s + 1, \\ 2 \cdot 3^s, & \text{if } n \geq 2 \text{ and } n = 3s + 2. \end{cases}$$

What about the maximal number of σ -extensions on weakly connected¹⁰ graphs? Does this number coincide with $\sigma_{\max}(n)$? The next theorem provides a negative answer to this question and thus gives space for impossibility results as we will see. For a natural number n define

$$\sigma_{\max}^{\text{con}}(n) = \max \{ |\sigma(F)| \mid F \in AF_{\mathfrak{A}}, |A_F| = n, F \text{ connected} \}$$

$\sigma_{\max}^{\text{con}}(n)$ returns the maximal number of σ -extensions among all connected AFs with n arguments. Again, a closed expression exists.

Theorem 41. *The function $\sigma_{\max}^{\text{con}}(n) : \mathbb{N} \rightarrow \mathbb{N}$ is given by*

$$\sigma_{\max}^{\text{con}}(n) = \begin{cases} n, & \text{if } n \leq 5, \\ 2 \cdot 3^{s-1} + 2^{s-1}, & \text{if } n \geq 6 \text{ and } n = 3s, \\ 3^s + 2^{s-1}, & \text{if } n \geq 6 \text{ and } n = 3s + 1, \\ 4 \cdot 3^{s-1} + 3 \cdot 2^{s-2}, & \text{if } n \geq 6 \text{ and } n = 3s + 2. \end{cases}$$

Proof. First some notations: for an AF $F = (A, R)$, denote its irreflexive version by

$$\text{irr}(F) = (A, R \setminus \{(a, a) \mid a \in A\});$$

denote its symmetric version by

$$\text{sym}(F) = (A, R \cup \{(b, a) \mid (a, b) \in R\});$$

⁸We mention that it is not the case that for all semantics σ the so-called *diversity function* $\Delta_{\sigma}(n)$ introduced in [23] coincides with $\sigma_{\max}(n)$ as introduced in [9] and defined below. This can be seen by considering complete semantics [23, 10].

⁹In this section, unless stated otherwise we use σ as a placeholder for stable, semi-stable, preferred, stage and naive semantics.

¹⁰In the following we simply write connected and take it to mean weakly connected.

and its associated undirected graph by

$$\text{und}(F) = (A, \{\{a, b\} \mid (a, b) \in R\}).$$

Furthermore, for a simple and undirected graph $G = (V, E)$ we use $MIS(G)$ for the set of maximal independent sets of G . Remember, a set $S \subseteq V$ is called *independent* if no edge $e \in E$ has both its endpoints in S . Moreover, an independent set S is called *maximal independent* if it is \subseteq -maximal among the independent sets of G . Finally, we denote its associated symmetric AF by

$$\text{dir}(G) = (V, \{(a, b), (b, a) \mid \{a, b\} \in E\}).$$

Now for the proof. We start with showing that the number of naive extensions does not exceed the claimed value range of $\sigma_{\max}^{\text{con}}(n)$. Given a connected AF F . Observe that the deletion of self-loops does not reduce the number of naive extensions, i.e. $|nai(F)| \leq |nai(\text{irr}(F))|$. This can be seen as follows. First, for any $E \in nai(F)$ exists an $E' \in nai(\text{irr}(F))$, such that $E \subseteq E'$ and second, for each two $E_1, E_2 \in nai(F)$ there is no $E' \in nai(\text{irr}(F))$, such that $E_1 \subseteq E'$ and $E_2 \subseteq E'$ simultaneously. Furthermore, it is easy to see that for any irreflexive AF G , $nai(G) = MIS(\text{und}(G))$. Roughly speaking, this is due to the fact that first, both concepts call for \subseteq -maximal sets and second, naive semantics does not distinguish between the presence of an attack (a, b) or the presence of (b, a) or the presence of both of them. Consequently, $|nai(F)| \leq |MIS(\text{und}(\text{irr}(F)))|$. Fortunately, due to Theorem 2 in [30] the maximal number of maximal independent sets in connected n -graphs are exactly given by the claimed value range of $\sigma_{\max}^{\text{con}}(n)$.


We proceed with arguing that the maximal number of stable extensions within the class of connected AFs is at least as large as the claimed value range of $\sigma_{\max}^{\text{con}}(n)$. In Figure 1 in [30] graphs realizing the maximal number of maximal independent sets for connected n -graphs are presented. These so-called *extremal graphs* can be used to derive AFs where former maximal independent sets become stable extensions. This can be done by replacing undirected edges by symmetric directed attacks. This construction is justified by the fact that for any simple graph G , $|MIS(G)| = |nai(\text{dir}(G))|$ and furthermore, as shown in [14, Propositions 4 and 5] naive and stable semantics coincide on the class of irreflexive and symmetric AFs. Example 13 below provides an illustration.

In order to conclude the proof we use well-known subset-relations between the considered semantics (compare Section 2). Since $stb(F) \subseteq stg(F) \subseteq nai(F)$ for any AF F , we derive that $|stb(F)| \leq |stg(F)| \leq |nai(F)|$. Furthermore, we have already shown that first, $\sigma_{\max}^{\text{con}}(n)$ does not exceed the claimed value range in case of naive semantics and second, $\sigma_{\max}^{\text{con}}(n)$ is at least as great as the claimed value range in case of stable semantics. Consequently, the stated equality provides us with a tight upper bound for stable, stage and naive semantics. What about semi-stable and preferred semantics? Since the result is already shown for stable semantics and in consideration of $stb(F) \subseteq sem(F) \subseteq prf(F)$ for any AF F , it suffices to prove that $\sigma_{\max}^{\text{con}}(n)$ does not exceed the claimed value range in case of preferred semantics. This can be seen as follows. First, one may easily verify that for any AF F we have, $|prf(F)| \leq |prf(\text{irr}(F))|$ as well as $prf(F) \subseteq prf(\text{sym}(F))$. Hence, $|prf(F)| \leq |prf(\text{sym}(\text{irr}(F)))|$. In [3, Corollary 1] it was already shown that preferred and stable semantics agree on irreflexive and symmetric AFs, i.e. for any AF F , $prf(\text{sym}(\text{irr}(F))) =$

$stb(sym(irr(F)))$). In summary, for any AF F we have, $|prf(F)| \leq |stb(sym(irr(F)))|$. Assuming the existence of an AF F possessing more preferred extension than the claimed value range of $\sigma_{\max}^{\text{con}}(n)$ implies the existence of an witnessing AF, namely $sym(irr(F))$, possessing more stable extension than the claimed value range of $\sigma_{\max}^{\text{con}}(n)$ in contrast to the already shown upper bound. Hence, the stated value range of $\sigma_{\max}^{\text{con}}(n)$ serves as a tight upper bound for semi-stable and preferred semantics too. \square

The following illustration provides an example how connected AFs having the maximal number of σ -extensions look like.

Example 13. Consider the following AF G :



The AF G is connected and possesses 22 σ -extensions. More precisely:

$$\sigma(G) = \{\{a_i, b_j, c_k\} \mid 1 \leq i, j, k \leq 3\} \setminus \{\{a_i, b_j, c_k\} \mid i = j = 3 \vee j = k = 3\}$$

This justifies $|\sigma(G)| = 27 - 5 = 22$. Furthermore, G consists of 9 arguments. Applying Theorem 41 we obtain $\sigma_{\max}^{\text{con}}(n) = 2 \cdot 3^{3-1} + 2^{3-1} = 2 \cdot 3^2 + 2^2 = 18 + 4 = 22$. This means, G is an extremal AF within the class of connected graphs. As an aside, in case of arbitrary frameworks, the maximal number of stable extensions given n arguments can be realized by deleting the mutual attacks between a_3 and b_3 as well as b_3 and c_3 (cf. Theorem 40). Restoring mutual attacks between one pair only yields the second largest number, which will be proven in Theorem 42. \diamond

A further interesting question concerning arbitrary AFs is whether all natural numbers less than $\sigma_{\max}(n)$ are realizable by AFs possessing n arguments.¹¹ The following theorem shows that there is a serious gap between the maximal and second largest number. For any positive natural n define

$$\sigma_{2\text{nd}}(n) = \max(\{|\sigma(F)| \mid F \in \text{AF}_{\text{alt}}, |A_F| = n\} \setminus \{\sigma_{\max}(n)\})$$

$\sigma_{2\text{nd}}(n)$ returns the second largest number of σ -extensions among all AFs with n arguments. Graph theory provides us with an expression.

Theorem 42. Function $\sigma_{2\text{nd}}(n) : \mathbb{N} \setminus \{0\} \rightarrow \mathbb{N}$ is given by

$$\sigma_{2\text{nd}}(n) = \begin{cases} \sigma_{\max}(n) - 1, & \text{if } 1 \leq n \leq 7, \\ \sigma_{\max}(n) \cdot \frac{11}{12}, & \text{if } n \geq 8 \text{ and } n = 3s + 1, \\ \sigma_{\max}(n) \cdot \frac{8}{9}, & \text{otherwise.} \end{cases}$$

¹¹We sometimes speak about realizing a natural number k and mean finding an AF having exactly k extensions.

Proof. At first we argue that the second largest number of σ -extensions is at least as large as the claimed value range of $\sigma_{2\text{nd}}(n)$. For this it suffices to present witnessing AFs. In [31, Theorem 2.4] graphs realizing the second largest number of maximal independent sets for n -graphs are given. These simple graphs can be used to derive AFs where former maximal independent sets become σ -extensions. Replacing undirected edges by symmetric directed attacks accounts for this. This can be seen as follows. First, for any simple graph G , $|MIS(G)| = |nai(dir(G))|$. Second, for any irreflexive and symmetric AF F we have, $stb(F) = nai(F)$ [14, Propositions 4 and 5] and finally, applying well-known subset-relations, namely $stb(F) \subseteq sem(F) \subseteq prf(F)$ and $stb(F) \subseteq stg(F)$ (for any AF F) justifies the claim for all considered semantics.

We show now that the second largest number of σ -extensions does not exceed the claimed value range of $\sigma_{2\text{nd}}(n)$. Given an AF F where $|A_F| = n$. Observe that we have nothing to show if $n \leq 7$ since $\sigma_{2\text{nd}}(n)$ is given as the maximal number minus one. Let $n \geq 8$ and suppose, towards a contradiction, that $l \cdot \sigma_{\max}(n) < \sigma_{2\text{nd}}(n) = |\sigma(F)| < \sigma_{\max}(n)$ where l depends on the remainder of n on division by 3 ($l \in \{\frac{11}{12}, \frac{8}{9}\}$). Similar to the proof of Theorem 41 one may easily show that for all considered semantics σ , $|\sigma(F)| \leq |\sigma(sym(irr(F)))|$ as well as that for any symmetric and irreflexive G , $\sigma(F) = MIS(und(G))$. This means, $l \cdot \sigma_{\max}(n) < |\sigma(F)| \leq |MIS(und(sym(irr(F))))| \leq \sigma_{\max}(n)$. We further conclude that $|MIS(und(sym(irr(F))))| = \sigma_{\max}(n)$. This equality has to hold because the term $l \cdot \sigma_{\max}(n)$ as well as the value range of $\sigma_{\max}(n)$ precisely coincide with the second largest or maximal number of maximal independent sets in simple graphs [31, 33]. This means, $l \cdot \sigma_{\max}(n) < |MIS(und(sym(irr(F))))| < \sigma_{\max}(n)$ would contradict the second largest number of maximal independent sets. Note that up to isomorphisms the extremal graphs are uniquely determined (cf. Theorem 1 in [30]). In the following we use K_n to denote a complete graph on n vertices. Depending on the remainder of n on division by 3 we have K_3 s for $n \equiv 0$, either one K_4 or two K_2 s and the rest are K_3 s in case of $n \equiv 1$ and one K_2 plus K_3 s for $n \equiv 2$. Remember that we have $|\sigma(F)| < |\sigma(sym(irr(F)))| = \sigma_{\max}(n)$. In particular, this means $F \neq sym(irr(F))$. Consequently, depending on the remainder we may thus estimate $|\sigma(F)| \leq k \cdot \sigma_{\max}(n)$ where $k \in \{\frac{2}{3}, \frac{3}{4}, \frac{1}{2}\}$. This can be seen as follows: First, computing the σ -extensions of an AF can be reduced to computing the σ -extensions of each of its component (see Lemma 46) and second, the minimal factors decreasing the number of σ -extension (through adding self-loops or deleting attacks) within a component where 3, 4 or 2 arguments attack each other are $\frac{2}{3}$, $\frac{3}{4}$ or $\frac{1}{2}$, respectively. We finally state $l \cdot \sigma_{\max}(n) < |\sigma(F)| \leq k \cdot \sigma_{\max}(n)$ where $l \in \{\frac{11}{12}, \frac{8}{9}\}$ and $k \in \{\frac{2}{3}, \frac{3}{4}, \frac{1}{2}\}$. This is a clear contradiction concluding the proof. \square

The attentive reader might have noticed that all maximal number functions introduced in this subsection refer to arbitrary argumentation frameworks instead of compact ones. The following theorem shows that this is not incidental since the compact versions of these functions return the same values (provided that they are defined). We first introduce the corresponding functions:

$$\begin{aligned}\sigma_{\max}^c(n) &= \max \{|\sigma(F)| \mid F \in CAF_\sigma, |A_F| = n\} \\ \sigma_{\max}^{c,\text{con}}(n) &= \max \{|\sigma(F)| \mid F \in CAF_\sigma, |A_F| = n, F \text{ connected}\} \\ \sigma_{2\text{nd}}^c(n) &= \max (\{|\sigma(F)| \mid F \in CAF_\sigma, |A_F| = n\} \setminus \{\sigma_{\max}^c(n)\})\end{aligned}$$

Theorem 43. For any $n \in \mathbb{N}$, $\sigma_{\max}^c(n) = \sigma_{\max}(n)$ and $\sigma_{\max}^{c, \text{con}}(n) = \sigma_{\max}^{\text{con}}(n)$. Furthermore, for any $n \in \mathbb{N} \setminus \{0, 1\}$, $\sigma_{2\text{nd}}^c(n) = \sigma_{2\text{nd}}(n)$.

Proof. (sketch) Given $n \in \mathbb{N}$ as claimed above.

(\leq) Obviously, $\sigma_{\max}^c(n) \leq \sigma_{\max}(n)$, $\sigma_{\max}^{c, \text{con}}(n) \leq \sigma_{\max}^{\text{con}}(n)$ and $\sigma_{2\text{nd}}^c(n) \leq \sigma_{2\text{nd}}(n)$ since $CAF_\sigma \subseteq AF_{\mathfrak{A}}$.
(\geq) Inspecting the proofs of Theorems 40, 41 and 42 (respective the proofs of the mentioned references therein) reveals that the witnessing examples, i.e. the AFs realizing a certain maximal number are already compact. \square

From now on we implicitly presuppose that the introduced maximal number functions restricted to compact frameworks coincide with their unrestricted versions.

Example 14. Recall that the (non-compact) AF we discussed in Example 12 had the extension-set \mathbb{S} with $|\mathbb{S}| = 17$ and $|\text{Args}_{\mathbb{S}}| = 8$. Is there a stable-compact AF with the same extensions? Firstly, nothing definitive can be said by Theorem 40 since $17 \leq 18 = \sigma_{\max}(8)$. Furthermore, in accordance with Theorem 41 the set \mathbb{S} cannot be compactly σ -realized by a connected AF since $17 > 15 = \sigma_{\max}^{\text{con}}(8)$. Finally, using Theorem 42 we infer that the set \mathbb{S} is not compactly σ -realizable because $\sigma_{2\text{nd}}(8) = 16 < 17 < 18 = \sigma_{\max}(8)$. \diamond

The compactness property is instrumental here, since Theorem 42 has no counterpart in non-compact AFs. More generally, allowing additional arguments as long as they do not occur in extensions enables us to realize any number of stable extensions up to the maximal one.

Proposition 44. Let n be a natural number. For each $k \leq \sigma_{\max}(n)$, there is an AF F with $|\text{Args}_{\text{stb}(F)}| = n$ and $|\text{stb}(F)| = k$.

Proof. To realize k stable extensions with n arguments, we start with the construction for the maximal number from Theorem 40. We then subtract extensions as follows: We choose $\sigma_{\max}(n) - k$ arbitrary distinct stable extensions of the AF realizing the maximal number. To exclude them, we use the construction of Definition 9 in [23]. \square

Corollary 45. Let n be a natural number and σ among preferred, semi-stable and stage semantics. For each $k \leq \sigma_{\max}(n)$, there is an AF F with $|\text{Args}_{\sigma(F)}| = n$ and $\sigma(F) = k$.

Proof. Follows from Lemmata 2.2 and 4.2 in [23]. \square

Now we are prepared to provide possible short cuts when deciding realizability of a given extension-set by initially simply counting the extensions. First some formal definitions.

Definition 12. Given an AF $F = (A, R)$, the component-structure $\mathcal{K}(F) = \{K_1, \dots, K_n\}$ of F is the set of sets of arguments, where each K_i coincides with the arguments of a weakly connected component of the underlying graph; $\mathcal{K}_{\geq 2}(F) = \{K \in \mathcal{K}(F) \mid |K| \geq 2\}$.

Example 15. The AF $F = (\{a, b, c\}, \{(a, b)\})$ has component-structure $\mathcal{K}(F) = \{\{a, b\}, \{c\}\}$. \diamond

The component-structure $\mathcal{K}(F)$ gives information about the number n of components of F as well as the size $|K_i|$ of each component. Knowing the components of an AF, computing the σ -extensions can be reduced to computing the σ -extensions of each component and building the cross-product. The AF resulting from restricting F to component K_i is given by $F \downarrow_{K_i} = (K_i, R_F \cap K_i \times K_i)$.

Lemma 46. *Given an AF F with component-structure $\mathcal{K}(F) = \{K_1, \dots, K_n\}$ it holds that the extensions in $\sigma(F)$ and the tuples in $\sigma(F \downarrow_{K_1}) \times \dots \times \sigma(F \downarrow_{K_n})$ are in one-to-one correspondence.*

Proof. By induction on n ; the base case $n = 1$ is trivial. For the induction step let $\mathcal{K}(F) = \{K_1, \dots, K_n, K_{n+1}\}$.

“ \subseteq ”: Let $S \in \sigma(F)$. Define $D_{n+1} = S \cap K_{n+1}$. By induction hypothesis, there are sets D_1, \dots, D_n such that each D_i is a σ -extension of $F \downarrow_{K_i}$ and $S \setminus K_{n+1} = D_1 \cup \dots \cup D_n$. We have to show that D_{n+1} is a σ -extension of $F \downarrow_{K_{n+1}}$. $\sigma = stb$: Clearly D_{n+1} is conflict-free, and any $a \in K_{n+1} \setminus D_{n+1}$ is attacked since S is stable and the attacks must come from D_{n+1} due to connectivity. $\sigma \in \{nai, prf\}$: If there is a conflict-free/admissible superset of D_{n+1} , then S is not naive/preferred for F . $\sigma \in \{stg, sem\}$: If there is a superset of D_{n+1} with greater range, then S is not stage/semi-stable for F .

“ \supseteq ”: Let D_1, \dots, D_n, D_{n+1} such that each D_i is a σ -extension of $F \downarrow_{K_i}$. Define $S = D_1 \cup \dots \cup D_n \cup D_{n+1}$; we show that $S \in \sigma(F)$. By induction hypothesis, $D_1 \cup \dots \cup D_n \in \sigma(F \downarrow_{K_1, \dots, K_n})$. $\sigma = stb$: Clearly S is conflict-free since all D_i are conflict-free; since D_{n+1} is stable for $F \downarrow_{K_{n+1}}$ it attacks all $a \in K_{n+1} \setminus D_{n+1}$ and thus S is stable for F . $\sigma \in \{nai, prf\}$: If S is not naive/preferred for F , there is a conflict-free/admissible superset of S in F . There is at least one additional argument, that is either in $D_1 \cup \dots \cup D_n$ or in D_{n+1} . But the first is impossible due to induction hypothesis, and the second due to presumption. $\sigma \in \{stg, sem\}$: If S is not stage/semi-stable for F , there is a conflict-free/admissible set S' with greater range. The range difference must manifest itself in $D_1 \cup \dots \cup D_n$ or D_{n+1} , which leads to a contradiction with the induction hypothesis and the presumption that D_{n+1} is stage/semi-stable for $F \downarrow_{K_{n+1}}$. \square

Given an extension-set \mathbb{S} we want to decide whether \mathbb{S} is realizable by a compact AF under semantics σ . For an AF $F = (A, R)$ with $\sigma(F) = \mathbb{S}$ we know that there cannot be a conflict between any pair of arguments in $Pairs_{\mathbb{S}}$, hence $R \subseteq \overline{Pairs_{\mathbb{S}}} = (A \times A) \setminus Pairs_{\mathbb{S}}$. The next proposition implicitly shows that for argument-pairs $(a, b) \notin Pairs_{\mathbb{S}}$, although there is not necessarily a direct conflict between a and b , they are definitely in the same component. In other words, this shows that implicit conflicts cannot arise across (weakly connected) components but only within them.

Proposition 47. *Let \mathbb{S} be an extension-set. (1) The transitive closure of $\overline{Pairs_{\mathbb{S}}}$, the set $(\overline{Pairs_{\mathbb{S}}})^*$, is an equivalence relation, that is, it is reflexive, symmetric, and transitive. (2) For each AF $F \in CAF_{\sigma}$ that compactly realizes \mathbb{S} under semantics σ (that is, $\sigma(F) = \mathbb{S}$), the component structure $\mathcal{K}(F)$ of F is given by the equivalence classes of $(\overline{Pairs_{\mathbb{S}}})^*$, that is, $\mathcal{K}(F)$ is the quotient set of $Args_{\mathbb{S}}$ by $(\overline{Pairs_{\mathbb{S}}})^*$.*

Proof. Consider some extension-set \mathbb{S} together with an AF $F \in CAF_{\sigma}$ with $\sigma(F) = \mathbb{S}$. We have to show that for any pair of arguments $a, b \in Args_{\mathbb{S}}$ it holds that $(a, b) \in (\overline{Pairs_{\mathbb{S}}})^*$ iff a and b are connected in the graph underlying F .

If a and b are connected in F , this means that there is a sequence s_1, \dots, s_n such that $a = s_1$, $b = s_n$, and $(s_1, s_2), \dots, (s_{n-1}, s_n) \notin \text{Pairs}_{\mathbb{S}}$, hence $(a, b) \in (\overline{\text{Pairs}_{\mathbb{S}}})^*$.

If $(a, b) \in (\overline{\text{Pairs}_{\mathbb{S}}})^*$ then also there is a sequence s_1, \dots, s_n such that $a = s_1$, $b = s_n$, and $(s_1, s_2), \dots, (s_{n-1}, s_n) \in \overline{\text{Pairs}_{\mathbb{S}}}$. Consider some $(s_i, s_{i+1}) \in \overline{\text{Pairs}_{\mathbb{S}}}$ and assume, towards a contradiction, that s_i occurs in another component of F than s_{i+1} . Recall that $F \in \text{CAF}_{\sigma}$, so each of s_i and s_{i+1} occur in some extension and $\sigma(F) \neq \emptyset$. Hence, by Lemma 46, there is some σ -extension $E \supseteq \{s_i, s_{i+1}\}$ of F , meaning that $(s_i, s_{i+1}) \in \text{Pairs}_{\mathbb{S}}$, a contradiction. Hence all s_i and s_{i+1} for $1 \leq i < n$ occur in the same component of F , proving that also a and b do so. \square

It is particularly nice to note that the only conditions we used in the proof were compactness and conflict-freeness, which indeed shows the Proposition for all five semantics considered here.

We will denote the component-structure induced by an extension-set \mathbb{S} as $\mathcal{K}(\mathbb{S})$, i.e. $\mathcal{K}(\mathbb{S})$ is the quotient set of $\text{Args}_{\mathbb{S}}$ by $(\overline{\text{Pairs}_{\mathbb{S}}})^*$. Note that, by Proposition 47, $\mathcal{K}(\mathbb{S})$ is equivalent to $\mathcal{K}(F)$ for every $F \in \text{CAF}_{\sigma}$ with $\sigma(F) = \mathbb{S}$. Given \mathbb{S} , the computation of $\mathcal{K}(\mathbb{S})$ can be done in polynomial time. With this we can use results from graph theory together with number-theoretical considerations in order to get impossibility results for compact realizability.

Recall that for a single connected component with n arguments the maximal number of σ -extensions is denoted by $\sigma_{\max}^{\text{con}}(n)$ and its values are given by Theorem 41. In the compact setting it further holds for a connected AF F with at least 2 arguments that $|\sigma(F)| \geq 2$.

Proposition 48. *Given an extension-set \mathbb{S} where $|\mathbb{S}|$ is odd, it holds that if $\exists K \in \mathcal{K}(\mathbb{S}) : |K| = 2$ then \mathbb{S} is not compactly realizable under semantics σ .*

Proof. Assume to the contrary that there is an $F \in \text{CAF}_{\sigma}$ with $\sigma(F) = \mathbb{S}$. We know that $\mathcal{K}(F) = \mathcal{K}(\mathbb{S})$. By assumption there is a $K \in \mathcal{K}(\mathbb{S})$ with $|K| = 2$, whence $|\sigma(K)| = 2$. Thus by Lemma 46 the total number of σ -extensions is even. Contradiction. \square

Example 16. Consider the extension-set $\mathbb{S} = \{\{a, b, c\}, \{a, b', c'\}, \{a', b, c'\}, \{a', b', c\}, \{a, b, c'\}, \{a', b, c\}, \{a, b', c\}\} = \text{stb}(F_1)$ where F_1 is the non-compact AF from the proof of Proposition 35. There, it took us some effort to argue that \mathbb{S} is not compactly *stb*-realizable. Proposition 48 now gives an easier justification: $\text{Pairs}_{\mathbb{S}}$ yields $\mathcal{K}(\mathbb{S}) = \{\{a, a'\}, \{b, b'\}, \{c, c'\}\}$. Thus \mathbb{S} with $|\mathbb{S}| = 7$ cannot be realized. \diamond

We denote the set of possible numbers of σ -extensions of a compact AF with n arguments as $\mathcal{P}(n)$; likewise we denote the set of possible numbers of σ -extensions of a compact and *connected* AF with n arguments as $\mathcal{P}^c(n)$. Although we know that $p \in \mathcal{P}(n)$ implies $p \leq \sigma_{\max}(n)$, there may be $q \leq \sigma_{\max}(n)$ that are not realizable by a compact AF under σ ; likewise for $q \in \mathcal{P}^c(n)$.

Clearly, any $p \leq n$ is possible by building an undirected graph with p arguments where every argument attacks all other arguments, a K_p , and filling up with k isolated arguments (k distinct copies of K_1) such that $k + p = n$. This construction obviously breaks down if we want to realize more extensions than we have arguments, that is, $p > n$. In this case, we have to use Lemma 46 and further graph-theoretic gadgets for addition and even a limited form of subtraction. Let us show how for $n = 7$ any number of extensions up to the maximal number 12 is realizable. For $12 = 3 \cdot 4$, Theorem 40 yields the realization, a disjoint union of a K_3 and a K_4 ($\triangle \boxtimes$). For

the remaining numbers, we have that $8 = 2 \cdot 4 \cdot 1$ and so we can combine a K_2 , a K_4 and a K_1 ($\dashv \dashv \boxtimes \cdot$). Likewise, $9 = 3 \cdot 3 \cdot 1$ ($\triangle \triangle \cdot$); $10 = 3 \cdot 3 + 1$ ($\triangle \triangle \triangle$) and finally $11 = 3 \cdot 4 - 1$ ($\triangle \boxtimes$). These small examples already show that \mathcal{P} and \mathcal{P}^c are closely intertwined and let us deduce some general corollaries: Firstly, $\mathcal{P}^c(n) \subseteq \mathcal{P}(n)$ since connected AFs are a subclass of AFs. Next, $\mathcal{P}(n) \subseteq \mathcal{P}(n+1)$ as in the step from $\triangle \triangle$ to $\triangle \triangle \cdot$. We even know that $\mathcal{P}(n) \subset \mathcal{P}(n+1)$ since $\sigma_{\max}(n+1) \in \mathcal{P}(n+1) \setminus \mathcal{P}(n)$. Furthermore, whenever $p \in \mathcal{P}(n)$, then $p+1 \in \mathcal{P}^c(n+1)$, as in the step from $\triangle \triangle$ to $\triangle \triangle \triangle$. The construction that goes from 12 to 11 above obviously only works if there are two weakly connected components overall, which underlines the importance of the component structure of the realizing AF. Indeed, multiplication of extension numbers of single components is our only chance to achieve overall numbers that are substantially larger than the number of arguments. This is what we will turn to next.

Having to leave the exact contents of $\mathcal{P}(n)$ and $\mathcal{P}^c(n)$ open, we can still state the following result:

Proposition 49. *Let \mathbb{S} be an extension-set that is compactly realizable under semantics σ where $\mathcal{K}_{\geq 2}(\mathbb{S}) = \{K_1, \dots, K_n\}$. Then for each $1 \leq i \leq n$ there is a $p_i \in \mathcal{P}^c(|K_i|)$ such that $|\mathbb{S}| = \prod_{i=1}^n p_i$.*

Example 17. Consider the extension-set $\mathbb{S}' = \{\{a, b, c\}, \{a, b', c'\}, \{a', b, c'\}, \{a', b', c\}\}$. (In fact there exists a (non-compact) AF F with $stb(F) = \mathbb{S}'$). We have the same component-structure $\mathcal{K}(\mathbb{S}') = \mathcal{K}(\mathbb{S})$ as in Example 16, but since now $|\mathbb{S}'| = 4$ we cannot use Proposition 48 to show impossibility of realization in terms of a compact AF. But with Proposition 49 at hand we can argue in the following way: $\mathcal{P}^c(2) = \{2\}$ and since $\forall K \in \mathcal{K}(\mathbb{S}') : |K| = 2$ it must hold that $|\mathbb{S}'| = 2 \cdot 2 \cdot 2 = 8$, which is obviously not the case. \diamond

In particular, we have a straightforward non-realizability criterion whenever $|\mathbb{S}|$ is a prime number: the AF (if any) must have at most one weakly connected component of size greater than two. Theorem 41 gives us the maximal number of σ -extensions in a single weakly connected component. Thus whenever the number of desired extensions is larger than that number and prime, it cannot be realized.

Corollary 50. *Let extension-set \mathbb{S} with $|\text{Args}_{\mathbb{S}}| = n$ be compactly realizable under σ . If $|\mathbb{S}|$ is a prime number, then $|\mathbb{S}| \leq \sigma_{\max}^{\text{con}}(n)$.*

Example 18. Let \mathbb{S} be an extension-set with $|\text{Args}_{\mathbb{S}}| = 9$ and $|\mathbb{S}| = 23$. We find that $\sigma_{\max}^{\text{con}}(9) = 2 \cdot 3^2 + 2^2 = 22 < 23 = |\mathbb{S}|$ and thus \mathbb{S} is not compactly realizable under semantics σ . \diamond

We can also make use of the derived component structure of an extension-set \mathbb{S} . Since the total number of extensions of an AF is the product of these numbers for its weakly connected components (Lemma 46), each non-trivial component contributes a non-trivial amount to the total. Hence if there are more components than the factorization of $|\mathbb{S}|$ has primes in it, then \mathbb{S} cannot be realized.

Corollary 51. *Let extension-set \mathbb{S} be compactly realizable under σ and let $f_1^{z_1} \cdot \dots \cdot f_m^{z_m}$ be the integer factorization of $|\mathbb{S}|$, where f_1, \dots, f_m are prime numbers. Then,*

$$z_1 + \dots + z_m \geq |\mathcal{K}_{\geq 2}(\mathbb{S})|.$$

Example 19. Consider an extension-set \mathbb{S} containing 21 extensions and $|\mathcal{K}_{\geq 2}(\mathbb{S})| = 3$. Since $21 = 3^1 * 7^1$ and further $1 + 1 < 3$, \mathbb{S} is not compactly realizable under semantics σ . \diamond

We conclude this section with a partial recipe for determining compact (non-)realizability. Given an extension-set \mathbb{S} , compute:

- the number of extensions $k = |\mathbb{S}|$,
- the number of arguments $n = |\text{Args}_{\mathbb{S}}|$,
- the component-structure $\mathcal{K}(\mathbb{S})$, in particular the number of non-trivial components $c = |\mathcal{K}_{\geq 2}(\mathbb{S})|$,
- the integer factorization of $k = f_1^{z_1} \cdot \dots \cdot f_m^{z_m}$

Towards deciding compact realizability, we can use the results of this section in the following way:

1. If $\sigma_{\max}(n) < k$ then \mathbb{S} is not compactly realizable.
2. If $\sigma_{2\text{nd}}(n) < k < \sigma_{\max}(n)$ then \mathbb{S} is not compactly realizable.
3. If $c = 1$ and $\sigma_{\max}^{\text{con}}(n) < k$ then \mathbb{S} is not compactly realizable.
4. If k is prime and $\sigma_{\max}^{\text{con}}(n) < k$ then \mathbb{S} is not compactly realizable.
5. If k is odd and there is a $K \in \mathcal{K}(\mathbb{S})$ with $|K| = 2$ then \mathbb{S} is not compactly realizable.
6. If $z_1 + \dots + z_m < c$ then \mathbb{S} is not compactly realizable.

7 Discussion

In this work we studied several aspects concerning the fundamental concepts of *rejected arguments* and *implicit conflicts* in abstract argumentation frameworks. For that, we focused on naive, stable, preferred, semi-stable and stage semantics, all of which satisfy the principle of I-maximality [4]. We omitted prominent basic semantics like complete, admissible and conflict-free sets from our studies as they do not align in this respect. However, for the sake of completeness we provide a complete account of results for the issues tackled in this paper for these semantics in C.

The idea of avoiding rejected arguments or implicit conflicts, lead us to introduce and study the novel classes of *compact and analytic argumentation frameworks*, each parameterized by a particular semantics. For both classes we obtained a similar picture for the relationship between semantics.

Concerning *computational issues*, we have analyzed the complexity of deciding whether a given AF is compact (resp. analytic) for a semantics σ . Our results range from tractability for naive semantics, over NP-completeness for stable and preferred semantics, up to Σ_2^P -completeness for semi-stable and stage semantics. We also have argued that the problem of credulous acceptance becomes trivial and skeptical acceptance is often polynomial time computable when we restrict ourselves to the subclasses under consideration, while the verification problems remain as hard as in the general case. The overall picture is now as follows: On the one hand we have illustrated that the classes of compact and analytic AFs provide computational benefits both in practice and in terms of theoretical worst-case analysis. On the other hand testing for membership in one of the classes is, for most of the semantics, of rather high complexity and thus these classes cannot be directly used to improve systems. However, for future work, we plan to take the rather negative complexity results for deciding membership in the subclass of compact (resp. analytic) AFs into account and seek for efficiently checkable (but not necessarily complete) criteria in order to decide whether a given AF (a) is compact (resp. analytic); and (b) whether it can be easily transformed into a compact (resp. analytic) one. The ultimate goal remains to design pre-processing procedures that identify rejected arguments that can be removed and implicit conflicts that can be made explicit; in other words, simplifications of the given AF into a semantically equivalent AF with better computational properties.

One of our main results was the *refutation of the Explicit Conflict Conjecture*, originally proposed in [7] for stable semantics. In fact, for each semantics σ among stable, preferred, semi-stable, and stage, we provided AFs where it is not possible to find an equivalent (under σ) AF where all conflicts become explicit. As a consequence, this result shows that in order to express a certain set \mathbb{S} of extensions via an AF, one cannot just draw attacks between any pair of arguments that do not occur jointly in any extension $E \in \mathbb{S}$. We believe that this not only gives a new insight into the fundamental properties of argumentation semantics, but also is important to be taken into account in research about the dynamics and evolvement of AFs. In particular these results indicate that techniques similar to Conflict-Driven Clause Learning [35] in SAT-solvers cannot be directly applied in the argumentation setting, as it can happen that a solver identifies a conflict between two arguments but it is impossible to add this conflict to the AF without changing the outcome. This is in particular interesting as most of the leading abstract argumentation systems are built on top of SAT-solvers [37, 38].

Finally, we addressed the question of *signatures and realizability*. We studied the relationship between signatures of compact and resp. analytic AFs. Our results complement the analysis from [23] and give a more fine-grained landscape about the expressive power of different semantics when the shape of AFs is restricted. Building on initial research from [9], we also analyzed possible numbers of extensions AFs can yield under a semantics at hand. Extending these considerations to admissible and complete semantics will be part of future work (cf. [10] for a conjecture regarding the maximal number of complete extensions). Results of the latter kind can also be beneficial for argumentation systems, since they may allow AF solvers to navigate more efficiently through the search space of possible extensions.

References

- [1] Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. An introduction to argumentation semantics. *Knowledge Eng. Review*, 26(4):365–410, 2011.
- [2] Pietro Baroni, Paul E. Dunne, and Massimiliano Giacomin. On the resolution-based family of abstract argumentation semantics and its grounded instance. *Artificial Intelligence*, 175(3-4):791–813, 2011.
- [3] Pietro Baroni and Massimiliano Giacomin. Characterizing defeat graphs where argumentation semantics agree. In Guillermo R. Simari and Paolo Torroni, editors, *Proc. ArgNMR*, pages 33–48, 2007.
- [4] Pietro Baroni and Massimiliano Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence*, 171(10-15):675–700, 2007.
- [5] Pietro Baroni and Massimiliano Giacomin. A systematic classification of argumentation frameworks where semantics agree. In Philippe Besnard, Sylvie Doutre, and Anthony Hunter, editors, *Proc. COMMA*, volume 172 of *Frontiers in Artificial Intelligence and Applications*, pages 37–48. IOS Press, 2008.
- [6] Ringo Baumann. Splitting an argumentation framework. In James P. Delgrande and Wolfgang Faber, editors, *Proc. LPNMR*, volume 6645 of *Lecture Notes in Computer Science*, pages 40–53. Springer, 2011.
- [7] Ringo Baumann, Wolfgang Dvořák, Thomas Linsbichler, Hannes Strass, and Stefan Woltran. Compact argumentation frameworks. In Torsten Schaub, Gerhard Friedrich, and Barry O’Sullivan, editors, *Proc. ECAI*, volume 263 of *Frontiers in Artificial Intelligence and Applications*, pages 69–74. IOS Press, 2014.
- [8] Ringo Baumann and Christof Spanring. Infinite Argumentation Frameworks – On the Existence and Uniqueness of Extensions. In Thomas Eiter, Hannes Strass, Mirosław Truszczyński, and Stefan Woltran, editors, *Advances in Knowledge Representation, Logic Programming, and Abstract Argumentation - Essays Dedicated to Gerhard Brewka on the Occasion of His 60th Birthday*, volume 9060 of *Lecture Notes in Computer Science*, pages 281–295. Springer, 2015.
- [9] Ringo Baumann and Hannes Strass. On the Maximal and Average Numbers of Stable Extensions. In Elizabeth Black, Sanjay Modgil, and Nir Oren, editors, *Proc. TAFE 2013*, volume 8306 of *Lecture Notes in Computer Science*, pages 111–126. Springer, 2014.
- [10] Ringo Baumann and Hannes Strass. Open Problems in Abstract Argumentation. In Thomas Eiter, Hannes Strass, Mirosław Truszczyński, and Stefan Woltran, editors, *Advances in Knowledge Representation, Logic Programming, and Abstract Argumentation - Essays Dedicated to Gerhard Brewka on the Occasion of His 60th Birthday*, volume 9060 of *Lecture Notes in Computer Science*, pages 325–339. Springer, 2015.

- [11] Trevor J. M. Bench-Capon and Paul E. Dunne. Argumentation in artificial intelligence. *Artificial Intelligence*, 171(10-15):619–641, 2007.
- [12] Martin Caminada and Leila Amgoud. On the evaluation of argumentation formalisms. *Artificial Intelligence*, 171(5-6):286–310, 2007.
- [13] Martin Caminada, Walter A. Carnielli, and Paul E. Dunne. Semi-stable semantics. *J. Log. Comput.*, 22(5):1207–1254, 2012.
- [14] Sylvie Coste-Marquis, Caroline Devred, and Pierre Marquis. Symmetric argumentation frameworks. In Lluís Godo, editor, *Proc. ECSQARU*, volume 3571 of *Lecture Notes in Computer Science*, pages 317–328. Springer, 2005.
- [15] Sylvie Coste-Marquis, Sébastien Konieczny, Jean-Guy Mailly, and Pierre Marquis. On the revision of argumentation systems: Minimal change of arguments statuses. In Chitta Baral, Giuseppe De Giacomo, and Thomas Eiter, editors, *Proc. KR*, pages 52–61. AAAI Press, 2014.
- [16] Martin Diller, Adrian Haret, Thomas Linsbichler, Stefan Rümmele, and Stefan Woltran. An extension-based approach to belief revision in abstract argumentation. In Qiang Yang and Michael Wooldridge, editors, *Proc. IJCAI*, pages 2926–2932. AAAI Press, 2015.
- [17] Yannis Dimopoulos and Alberto Torres. Graph theoretical structures in logic programs and default theories. *Theoretical Computer Science*, 170(1-2):209–244, 1996.
- [18] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, 1995.
- [19] Paul E. Dunne. Computational properties of argument systems satisfying graph-theoretic constraints. *Artificial Intelligence*, 171(10–15):701–729, 2007.
- [20] Paul E. Dunne. The computational complexity of ideal semantics. *Artificial Intelligence*, 173(18):1559–1591, 2009.
- [21] Paul E. Dunne and Trevor J. M. Bench-Capon. Coherence in finite argument systems. *Artificial Intelligence*, 141(1/2):187–203, 2002.
- [22] Paul E. Dunne, Wolfgang Dvořák, Thomas Linsbichler, and Stefan Woltran. Characteristics of multiple viewpoints in abstract argumentation. In Christoph Beierle and Gabriele Kern-Isberner, editors, *Proc. DKB*, pages 16–30, 2013.
- [23] Paul E. Dunne, Wolfgang Dvořák, Thomas Linsbichler, and Stefan Woltran. Characteristics of multiple viewpoints in abstract argumentation. *Artificial Intelligence*, 228:153–178, 2015.
- [24] Wolfgang Dvořák. *Computational Aspects of Abstract Argumentation*. PhD thesis, Vienna University of Technology, 2012.

- [25] Wolfgang Dvořák, Matti Järvisalo, Johannes Peter Wallner, and Stefan Woltran. Complexity-sensitive decision procedures for abstract argumentation. *Artificial Intelligence*, 206:53–78, 2014.
- [26] Wolfgang Dvořák and Christof Spanring. Comparing the expressiveness of argumentation semantics. *J. Log. Comput.*, 2016.
- [27] Wolfgang Dvořák and Stefan Woltran. Complexity of semi-stable and stage semantics in argumentation frameworks. *Inf. Process. Lett.*, 110(11):425–430, 2010.
- [28] Sjur K. Dyrkolbotn. How to argue for anything: Enforcing arbitrary sets of labellings using AFs. In Chitta Baral, Giuseppe De Giacomo, and Thomas Eiter, editors, *Proc. KR*, pages 626–629. AAAI Press, 2014.
- [29] Uwe Egly, Sarah Alice Gaggl, and Stefan Woltran. Answer-set programming encodings for argumentation frameworks. *Argument & Computation*, 1(2):147–177, 2010.
- [30] Jerrold R. Griggs, Charles M. Grinstead, and David R. Guichard. The number of maximal independent sets in a connected graph. *Discrete Mathematics*, 68(23):211–220, 1988.
- [31] Zemin Jin and Xueliang Li. Graphs with the second largest number of maximal independent sets. *Discrete Mathematics*, 308(23):5864–5870, 2008.
- [32] Thomas Linsbichler, Christof Spanring, and Stefan Woltran. The hidden power of abstract argumentation semantics. In Elizabeth Black, Sanjay Modgil, and Nir Oren, editors, *Proc. TAFE*, volume 9524 of *Lecture Notes in Computer Science*, pages 146–162. Springer, 2015.
- [33] John W. Moon and Leo Moser. On cliques in graphs. *Israel Journal of Mathematics*, 3(1):23–28, 1965.
- [34] Iyad Rahwan and Guillermo R. Simari, editors. *Argumentation in Artificial Intelligence*. Springer, 2009.
- [35] João P. Marques Silva, Inês Lynce, and Sharad Malik. Conflict-driven clause learning SAT solvers. In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*, pages 131–153. IOS Press, 2009.
- [36] Hannes Strass. The relative expressiveness of abstract argumentation and logic programming. In Sven Koenig and Blai Bonet, editors, *Proc. AAAI*, pages 1625–1631. AAAI Press, 2015.
- [37] Matthias Thimm and Serena Villata. System descriptions of the first international competition on computational models of argumentation (ICCA’15). *CoRR*, abs/1510.05373, 2015.
- [38] Matthias Thimm, Serena Villata, Federico Cerutti, Nir Oren, Hannes Strass, and Mauro Vallati. Summary report of the first international competition on computational models of argumentation. *AI Magazine*, 37(1):102, 2016.

- [39] Bart Verheij. Two approaches to dialectical argumentation: admissible sets and argumentation stages. In John-Jules C. Meyer and Linda C. van der Gaag, editors, *Proc. NAIC*, pages 357–368, 1996.

A Implicit Conflicts and Runtime

This appendix reports the results of experiments carried out in order to understand the impact of implicit conflicts on the efficiency of solvers. The underlying assumption is that the explication of implicit conflicts leads to a more succinct representation of an AF, which in turn supports the solver by providing more information to work with. The results for extension enumeration under stable and preferred semantics confirm this hypothesis.

The First International Competition on Computational Models of Argumentation (ICCMA, see <http://argumentationcompetition.org/2015/>) featured three types of benchmark graphs, each following a certain graph model. The one which turned out to be the hardest for most solvers is based on a construction that aims at having many stable (and preferred) extensions.¹² As the set containing the largest instances of this type even had to be removed from the competition due to its difficulty and therefore seems most interesting to analyze, we focus on this group of instances.¹³ For stable and preferred semantics we analyze the correlation between the number of implicit conflicts and the runtime of solving. We focus on the particular task of extension enumeration and employ the winning solvers of the competition, that is ASPARTIX [29] for stable semantics and CEGARTIX [25] for preferred semantics. We report statistics for the instances of the competition as well as for a larger set of instances constructed with the same generator to get results which are statistically more significant.

Instances of the considered type usually feature many rejected arguments – the ICCMA-instances have an average of 78% for stable and 76% for preferred semantics. Therefore we only take into account implicit conflicts that occur between arguments that are not rejected. Hence, in this appendix, given an AF F and semantics σ , there is an implicit conflict between arguments $a, b \in A_F$ if $(a, b) \notin Pairs_{\sigma(F)}$, $(a, b), (b, a) \notin R_F$ and $a, b \in Args_{\sigma(F)}$. The σ -explication of F is then $(A_F, R_F \cup \{(a, b), (b, a) \mid \text{implicit conflict between } a \text{ and } b\})$. Note that, in general, the explicated AF is not equivalent w.r.t. σ to the original one. However, it is not a simplification in terms of extensions, as the number of extensions of the explicated AF is always greater or equal than in the original AF.

The tests have been run on a machine with two AMD Opteron 6308 processors (3.5GHz) having 2 physical cores each; each of these cores puts at disposal 2 logical cores (per hyperthreading), 192 GB RAM (12 x 16GB) and a timeout of 600 seconds for each instance.

The experiments show a clear tendency to shorter runtime when decreasing the number of implicit conflicts. Figure 26 shows the experimental results for the ICCMA instances and Figure 27 for the generated instances. The correlation between implicit conflicts per non-rejected argument and runtime of solving turns out to be slightly greater for preferred semantics. Explication of implicit conflicts leads to a speedup in solving for both semantics, more significantly for stable though. Table 3 shows a few more statistics for the generated instances. While the solvers have

¹²Generators can be found at <https://sourceforge.net/projects/probo/>.

¹³Besides, the other two types are of limited interest also because of other reasons: (i) Instances of the first type have only one (large) extension under all standard semantics and are therefore relatively easy to solve and, more importantly, have no implicit conflicts among non-rejected arguments; (ii) The other instance type has a rich structure of strongly connected components leading to very few (or no) extensions and only a small amount of implicit conflicts.




Figure 26: Experimental results for implicit conflicts on ICCMA 2015 instances. Out of the 48 instances there are 30 (63%) for *stb* and 33 (69%) for *prf* which contain at least one implicit conflict. Among those, the number of implicit conflicts divided by the number of non-rejected arguments is 16.0 for *stb* and 15.6 for *prf* on average. The scatter plots at the top relate the number of implicit conflicts per non-rejected argument to the runtime of extension enumeration. We get a correlation coefficient of 0.16 for *stb* and 0.26 for *prf*. The diagrams at the bottom show the cumulative runtime for extension enumeration for the original set of AFs on the one hand and the corresponding explicated AFs on the other.




Figure 27: Experimental results for implicit conflicts on an additional 1061 instances obtained from the same generator as the most difficult instances of ICCMA 2015. This set of AFs contains 533 (50%) for *stb* and 567 (53%) for *prf* with at least one implicit conflict. Among those, the number of implicit conflicts divided by the number of non-rejected arguments is 17.1 for *stb* and 17.0 for *prf* on average. Again, the scatter plots at the top of the figure relate the number of implicit conflicts per non-rejected argument to the runtime of extension enumeration. Now we get a correlation coefficient of 0.24 for *stb* and 0.30 for *prf*. The cumulative runtimes for extension enumeration indicating a significant speedup for the explicated AFs are shown at the bottom of the figure.

Table 3: Statistics for the generated instances. The total number of instances is 533 for stable and 567 for preferred semantics. Every instance which led to a timeout for the explicated instance also did so for the original one.

Semantics	stable		preferred	
	original	explicit	original	explicit
solved instances	512	528	544	561
timeouts	21	5	23	6
faster	115	416	257	304
uniquely solved	0	16	0	17
average runtime (s)	49.7	11.6	60.7	23.9
maximal runtime (s)	600.0	600.0	600.0	600.0

about four times as many timeouts for the original instances than for the explicated instances for both semantics, the number of instances which are solved faster is less significantly different for preferred semantics. The explanation is that, for preferred semantics, explication only pays off if the original AF is sufficiently difficult to solve.

In this appendix we have only considered implicit conflicts among non-rejected arguments. Dropping this constraint would have the following effects on the experiments: due to the already mentioned high number of rejected arguments (i) the size of explicated AFs significantly increases, and (ii) the number of implicit conflicts mostly depends on the number of rejected arguments. The effect on the runtime for extension enumeration on these explicated instances strongly differs between stable and preferred semantics. We get runtime on a constantly low level (between one and two seconds) for stable semantics and runtimes significantly higher than on the original instances for preferred instances. While the former observation is somehow in accordance with [7] as discussed at the end of Section 4, we attribute the latter to the almost quadratic increase in instance size.

B Proofs of Section 4

Proof of Lemma 19. Below we will show that all arguments except x are always credulously accepted in F' and that x is credulously accepted in F' iff x is credulously accepted in F .

First we show that each semi-stable extension E contains at least one argument from $\{t_a \mid a \in A \cup \{y, z\}\}$. Suppose that not, then $E \cup \{t_x\}$ is a conflict-free set and, as t_x defends itself against all its attackers, the set $E \cup \{t_x\}$ is also admissible. Hence we have a contradiction to the maximality of E . Further, as $\{t_a \mid a \in A \cup \{y, z\}\}$ forms a clique in F we get that each extension contains exactly one argument from the set.

Next we show that the ranges of preferred extensions E containing an argument from $\{t_a \mid a \in \{x, y, z\}\}$ cannot be contained in the ranges of admissible sets E' not containing any of these arguments.

- If $t_x \in E$ then h_2 is defended against all its attackers, that are t_y and h_1 , and thus also

$h_2 \in E$. As h_1 is attacked by t_x and h_3 is attacked by h_2 we have $\{h_1, h_2, h_3\} \subseteq E_{F'}^+$. By similar arguments we get that $\{h_1, h_2, h_3\} \subseteq E_{F'}^+$ if either $t_y \in E$ or $t_z \in E$.

- If $t_a \in E'$ with $a \notin \{x, y, z\}$ then none of the h_1, h_2, h_3 can be in the range, as they form an odd length cycle and all attacking arguments from outside are counter attacked by E' .

We will next consider these two kind of extensions separately.

- First, consider the sets E containing an argument from $\{t_a \mid a \in \{x, y, z\}\}$. By the above we have that either $\{t_x, h_2\} \subseteq E$, $\{t_y, h_3\} \subseteq E$, or $\{t_z, h_1\} \subseteq E$. All of these three sets have the same attacks to the remaining arguments and thus we have that for each $E' \subseteq A$, $\{t_x, h_2\} \cup E' \in \text{sem}(F')$ iff $\{t_y, h_3\} \cup E' \in \text{sem}(F')$ iff $\{t_z, h_1\} \cup E' \in \text{sem}(F')$. As at least for $E' = \emptyset$ these sets are also admissible this implies that the arguments $\{t_x, t_y, t_z, h_1, h_2, h_3\}$ are credulously accepted in F' .

Moreover, $\{t_x, h_2\}$ defends the arguments A against all attack from arguments in A' and does not attack any argument in A . Thus

- (i) $\{t_x, h_2\} \cup E' \in \text{adm}(F')$ iff $E' \in \text{adm}(F)$ and
- (ii) as arguments in A do not attack arguments in A' we have that $(\{t_x, h_2\} \cup E')_{F'}^+ = \{t_x, h_2\}_{F'}^+ \cup E_{F'}^+$ and thus $(\{t_x, h_2\} \cup E')_{F'}^+$ is maximal when $E_{F'}^+$ is maximal.

Hence, $\{t_x, h_2\} \cup E' \in \text{sem}(F')$ iff $E' \in \text{sem}(F)$ and x is credulously accepted in F iff x is credulously accepted in F' .

As, by assumption, F has no stable extension there cannot be an extension E containing an argument from $\{t_a \mid a \in \{x, y, z\}\}$ and having all arguments A in its range.

- Second, consider the sets E containing an argument from $\{t_a \mid a \in A \setminus \{x\}\}$. Now it is easy to verify that the sets $\{t_a, a\}$ for $a \in A \setminus \{x\}$ are admissible sets of F' and have maximal range among the extensions containing $\{t_{a,b} \mid a, b \in A\}$. In particular A is in the range of each of these extensions, and thus they are incomparable with the extension of the first type, i.e. they are semi-stable. Hence, we have that the arguments $\{a, t_a \mid a \in A \setminus \{x\}\}$ are credulously accepted. Moreover, no extensions E' with $\{t_a \mid a \in A \setminus \{x\}\} \cap E' \neq \emptyset$ can contain x .

Finally, combining the above results, we have that all arguments in A' except x are credulously accepted in F' and x is credulously accepted in F iff x is credulously accepted in F' iff F' is *stg*-compact. \square

Proof of Lemma 23. First, we show that each semi-stable extension E contains at least one argument from A' . Suppose that not, then $E \setminus A_G \cup \{t\}$ is an admissible set that has A_G in its range. Hence we have a contradiction to the range maximality of E . Further, as A' forms a clique in F we get that each semi-stable extension contains exactly one argument from the set.

Next we show that the ranges of preferred extensions E containing argument t cannot be contained in the ranges of admissible sets E' containing an argument $t_{a,b}$ with $a, b \in A$.

- If $t \in E$ then all arguments in A_G are attacked by E and thus are in the range of E .
- If $t \notin E'$ at least one argument of A_G is not in the range of E' . Otherwise, $E' \cap A_G$ would be a stable extension of G , which contradicts $stb(G) = \emptyset$.

We will next consider these two kind of extensions separately.

- First, consider the sets E containing t . As t does not attack any argument in A nor does A have any outgoing attacks we have

- (i) $\{t\} \cup E' \in adm(F')$ iff $E' \in adm(F)$ and
- (ii) as arguments in A do not attack arguments outside A we have that $(\{t\} \cup E')_{F'}^+ = \{t\}_{F'}^+ \cup E'_{F'}^+$ and thus $(\{t\} \cup E')_{F'}^+$ is maximal when $E'_{F'}^+$ is maximal.

Hence, $\{t\} \cup E' \in sem(F')$ iff $E' \in sem(F)$ and $\{t, a\}$ is credulously accepted in F' iff a is credulously accepted in F .

As, by assumption, F has no stable extension there cannot be a semi-stable extension E containing t and having all arguments A in its range.

- Second, consider the extensions E containing an argument $t_{a,b}$ with $a, b \in A$. Now it is easy to verify that the sets $\{t_{a,b}, a, b\}$ are admissible sets of F' and A is in the range of each of these extensions. Thus they are incomparable with the extension of the first type.

As $t_{a,b}$ does not attack any argument in A_G nor does A_G have any outgoing attacks we have

- (i) $\{t_{a,b}, a, b\} \cup E' \in adm(F')$ iff $E' \in adm(G)$ and
- (ii) as arguments in A_G do not attack arguments outside A_g we have that $(\{t_{a,b}, a, b\} \cup E')_{F'}^+ = \{t_{a,b}, a, b\}_{F'}^+ \cup E'_{F'}^+$ and thus $(\{t_{a,b}, a, b\} \cup E')_{F'}^+$ is maximal when $E'_{F'}^+$ is maximal.

Thus, $\{t_{a,b}, a, b\} \cup E' \in sem(F')$ iff $E' \in sem(G)$. Now, as $G \in XAF_{sem}$ we have that for each $g, g' \in G$ with $(g, g'), (g', g) \notin R_G$ there is an $E' \in sem(G)$ with $g, g' \in E'$. Furthermore as G has no self-attacks it is also compact (cf. Proposition 5) and thus for each $g \in A_G$ there is an $E' \in sem(G)$ with $g \in E'$. From these stage extensions we obtain that:

- $\{t_{a,b}, a\}, \{t_{a,b}, b\}$ are credulously accepted in F' ;
- $\{t_{a,b}, g, g'\}$ is credulously accepted in F' , for $g, g' \in G$ with $(g, g'), (g', g) \notin R_G$;
- $\{t_{a,b}, g\}$ is credulously accepted in F' , for each $g \in G$;
- $\{a, g\}$ is credulously accepted in F' , for each $a \in A$ and $g \in G$;

Combining the above results, we have that all non-conflicting pairs of arguments in F' except $\{t, a\}$ with $a \in A$ are credulously accepted in F' . Thus F' is *stg-analytic* iff all the pairs $\{t, a\}$ with $a \in A$ are credulously accepted in F' iff each $a \in A$ is credulously accepted in F iff F is *sem-compact*.

Finally we show that $F' \in CAF_{sem}$ (independent of whether $F \in CAF_{sem}$). As (i) for each $a \in A$ the set $\{t_{a,a}, a\}$ is credulously accepted, and (ii) for each $g \in A_G$ and $a, b \in A$ with $(a, b), (b, a) \notin R$ the set $\{t_{a,b}, g\}$ is credulously accepted, the AF F' is *stg-compact*. \square

C Basic Semantics

In the following show how the concepts considered in the paper carry over to conflict-free, admissible and complete semantics. For that we have to introduce complete semantics first [18]:

Definition 13. Given $F \in AF_{\exists!}$, $S \in com(F)$ if $S \in adm(F)$ and for each $a \in A$ that is defended by S it holds that $a \in S$.

Proposition 52. *It holds that $CAF_{cf} = CAF_{nai}$ and $CAF_{adm} = CAF_{com} = CAF_{prf}$.*

Proof. Since naive (resp. preferred) extensions of any given AF F are exactly the \subseteq -maximal conflict-free (resp. admissible and complete) extensions it holds that $Args_{nai(F)} = Args_{cf(F)}$ and $Args_{prf(F)} = Args_{adm(F)} = Args_{com(F)}$. Therefore $F \in CAF_{nai}$ iff $F \in CAF_{cf}$ and $F \in CAF_{prf}$ iff $F \in CAF_{adm}$ iff $F \in CAF_{com}$. \square

Proposition 53. *It holds that $XAF_{cf} = XAF_{nai}$ and $XAF_{adm} = XAF_{com} = XAF_{prf}$.*

Proof. Since naive (resp. preferred) extensions of any given AF F are exactly the \subseteq -maximal conflict-free (resp. admissible and complete) extensions it holds that $Pairs_{nai(F)} = Pairs_{cf(F)}$ and $Pairs_{prf(F)} = Pairs_{adm(F)} = Pairs_{com(F)}$. Therefore $F \in XAF_{nai}$ iff $F \in XAF_{cf}$ and $F \in XAF_{prf}$ iff $F \in XAF_{adm}$ iff $F \in XAF_{com}$. \square

Corollary 54. *The following problems are in P:*

1. *Given AF F , deciding whether $F \in CAF_{cf}$;*
2. *Given AF F , deciding whether $F \in XAF_{cf}$.*

Proof. Follows from Corollary 11 and Propositions 52 and 53. \square

Corollary 55. *The following problems are NP-complete for $\sigma \in \{adm, com\}$.*

1. *Given AF F , deciding whether $F \in CAF_{\sigma}$;*
2. *Given AF F , deciding whether $F \in XAF_{\sigma}$; hardness already holds if the problem is restricted to AFS $F \in CAF_{\sigma}$.*

Proof. Follows from Corollary 15 and Propositions 52 and 53. \square

Proposition 56. *ECC holds for cf.*

Proof. Consider an arbitrary AF F , and observe that in order to have an implicit conflict between arguments a and b , at least one of these arguments has to be self-attacking. Hence we obtain an analytic AF F' having $cf(F') = cf(F)$ by letting each self-attacking argument be attacked by all other arguments. \square

Proposition 57. *ECC does not hold for adm and com.*




Figure 28: AF F compactly realizing an extension-set $\mathbb{S} \notin \Sigma_{adm}^c \cup \Sigma_{cf}^c$ under com .

Proof. Assume ECC holds for adm (resp. com) and let F be an AF which is non-analytic for prf . By assumption there is an AF F' with $A_{F'} = A_F$ and $adm(F') = adm(F)$ (resp. $com(F') = com(F)$). But then also $prf(F') = prf(F)$, a contradiction to F being non-analytic for prf . \square

Proposition 58. *It holds that*

1. $\Sigma_{cf}^c = \Sigma_{cf}$ and
2. $\Sigma_{\sigma}^c \subset \Sigma_{\sigma}$ for $\sigma \in \{adm, com\}$.

Proof. (1) Follows directly from Proposition 35.1.

(2) Let $\sigma \in \{adm, com\}$, $\mathbb{S} = \{\emptyset, \{a, b\}\}$ and assume there is an AF $F \in CAF_{\sigma}$ with $\sigma(F) = \mathbb{S}$. Since a and b are free of conflict it must hold that $F = (\{a, b\}, \emptyset)$. But then we get $adm(F) = \{\emptyset, \{a\}, \{b\}, \{a, b\}\}$ and $com(F) = \{\{a, b\}\}$, hence $\mathbb{S} \notin \Sigma_{\sigma}^c$. On the other hand there is the non-compact AF $F' = (\{a, b, c, d\}, \{(a, c), (c, c), (c, b), (b, d), (d, d), (d, a)\})$ having $\sigma(F') = \mathbb{S}$, hence $\mathbb{S} \in \Sigma_{\sigma}$. \square

70

Proposition 59. *It holds that*

1. $\Sigma_{cf}^c \subset \Sigma_{adm}^c$ and
2. $\Sigma_{com}^c \not\subseteq \Sigma_{\sigma}^c$ and $\Sigma_{\sigma}^c \not\subseteq \Sigma_{com}^c$ for $\sigma \in \{cf, adm\}$.

Proof. (1) Given an arbitrary AF F it holds that $cf(F) = adm(sym(F))$, where $sym(F)$ is the AF obtained from making all attacks of F symmetric, hence $\Sigma_{cf}^c \subseteq \Sigma_{adm}^c$. Properness is by the AF $G = (\{a, b, c, d\}, \{(a, b), (b, c), (c, d), (d, a)\})$, that is the directed cycle of four arguments, having $adm(G) = \{\emptyset, \{a, c\}, \{b, d\}\}$, which is an extension-set not realizable under cf . This is by the fact that if $\{a, c\}$ is conflict-free in some AF then clearly also $\{a\}$ and $\{c\}$ must be conflict-free. Hence $\Sigma_{cf}^c \subset \Sigma_{adm}^c$.

(2) $\Sigma_{com}^c \not\subseteq \Sigma_{\sigma}^c$: Any extension-set \mathbb{S} containing exactly one non-empty set of arguments S is compactly realizable under com by the AF (S, \emptyset) , but not under cf and adm since \emptyset is not contained in \mathbb{S} . The following example shows that these trivial cases are not the only AFs in $\Sigma_{com}^c \setminus \Sigma_{\sigma}^c$. To this end consider the AF F depicted in Figure 28. We have $com(F) = \{\emptyset, \{a_1\}, \{a_2\}, \{b_1\}, \{b_2\}, \{a_1, b_2\}, \{a_2, b_1\}, \{a_1, a_2, c\}, \{b_1, b_2, d\}\}$. On the one hand it is easy to see that F is compact for complete semantics, on the other hand observe that both $\{a_1\}, \{a_2\} \in com(F)$, $(a_1, a_2) \in Pairs_{com(F)}$, but $\{a_1, a_2\} \notin com(F)$. So $com(F)$ violates a necessary condition

for admissible and conflict-free extension-sets (cf. [23]). Hence $com(F) \notin \Sigma_\sigma$ and therefore by Proposition 58 also $com(F) \notin \Sigma_\sigma^c$.

$\Sigma_\sigma^c \not\subseteq \Sigma_{com}^c$: Let $F = (\{a, b, c\}, \{\langle a, b \rangle, \langle b, c \rangle\})$ and observe that $cf(F) = adm(F) = \{\emptyset, \{a\}, \{b\}, \{c\}, \{a, c\}\}$. Now assume there is an AF $G \in CAF_{com}$ with $com(G) = cf(F)$. Clearly $A_G = \{a, b, c\}$ and $R_G \subseteq \{(a, b), (b, a), (b, c), (c, b)\}$. Now for $\emptyset \in com(G)$ each argument must be attacked and, moreover, the singletons $\{a\}$, $\{b\}$ and $\{c\}$ must defend themselves. Hence it must be that $G = F$ which means $com(G) = \{\emptyset, \{a, c\}, \{b\}\}$, a contradiction. \square

Proposition 60. *It holds that*

1. $\Sigma_{cf}^x = \Sigma_{cf}$ and
2. $\Sigma_\sigma^x \subset \Sigma_\sigma$ for $\sigma \in \{adm, com\}$.

Proof. (1) Follows directly from Theorem 30.

(2) Consider the AF depicted in Figure 16 which was discussed in Example 6. We show in Theorem 31 that $prf(F) \notin \Sigma_{prf}^x$. Observe that F has the same implicit conflicts (namely between a_1 and x_2 , a_2 and x_3 , and a_3 and x_1) for preferred, admissible and complete semantics. Now assuming that $com(F) \in \Sigma_{com}^x$ (resp. $adm(F) \in \Sigma_{adm}^x$) means that there is some AF F' which is analytic for com (resp. adm) and has $com(F') = com(F)$ (resp. $adm(F') = adm(F)$). But then F' is also analytic for prf and has $prf(F') = prf(F)$, a contradiction to $prf(F) \notin \Sigma_{prf}^x$. Hence $com(F) \notin \Sigma_{com}^x$ and $adm(F) \notin \Sigma_{adm}^x$. \square

Proposition 61. *It holds that*

1. $\Sigma_{cf}^x \subset \Sigma_{adm}^x$ and
2. $\Sigma_{com}^x \not\subseteq \Sigma_\sigma^x$ and $\Sigma_\sigma^x \not\subseteq \Sigma_{com}^x$ for $\sigma \in \{cf, adm\}$.

Proof. (1) The argument from in the proof of Proposition 59 applies here as well.

(2) $\Sigma_{com}^x \not\subseteq \Sigma_\sigma^x$: Consider the AF F from the proof of Proposition 59 (resp. Figure 28) and extend it by a symmetric attack between arguments c and d as follows $F' = (A_F, R_F \cup \{\langle c, d \rangle\})$. Now $com(F) = com(F')$ and it is easy to verify that F' is analytic for complete semantics, but as discussed before $com(F) = com(F') \notin \Sigma_\sigma$. Hence, we have a witness for $\Sigma_{com}^x \not\subseteq \Sigma_\sigma^x$.

$\Sigma_\sigma^x \not\subseteq \Sigma_{com}^x$: Again consider the AF $F = (\{a, b, c\}, \{\langle a, b \rangle, \langle b, c \rangle\})$ and recall that $cf(F) = adm(F) = \{\emptyset, \{a\}, \{b\}, \{c\}, \{a, c\}\}$. Assume there is an AF $G \in XAF_x$ with $com(G) = cf(F)$. Clearly $\{a, b, c\} \subseteq A_G$ and $R_G \cap (A_G \times A_G) \subseteq \{(a, b), (b, a), (b, c), (c, b)\}$. Consider arguments in A_G that are different from a, b, c . As such arguments do not appear in any extensions they have to be self-attacking and in conflict with all the other arguments. From $\{a\}, \{b\}, \{c\} \in cf(F) = adm(F)$ we obtain that a, b, c attack all arguments in $A_G \setminus \{a, b, c\}$. Now as $\{b\} \in com(G)$ we have that a and c must be attacked by some argument not attacked by b . Thus $(b, a) \in R_G$ and $(b, c) \in R_G$ and as $\{a\}, \{c\} \in com(G)$ and have to defend themselves also $(a, b) \in R_G$ and $(c, b) \in R_G$. But then we have $com(G) = \{\emptyset, \{a, c\}, \{b\}\}$, a contradiction. \square