dbai

TECHNICAL REPORT

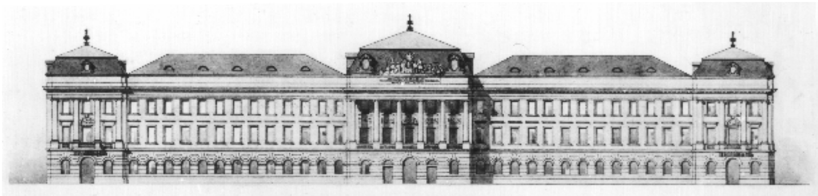

INSTITUT FÜR LOGIC AND COMPUTATION

ABTEILUNG DATENBANKEN UND ARTIFICIAL INTELLIGENCE

Graph-Classes of Argumentation Frameworks with Collective Attacks

DBAI-TR-2021-120

Wolfgang Dvořák **Matthias König** **Stefan Woltran**

DBAI TECHNICAL REPORT

2021

Institut für Logic and Computation
Abteilung Datenbanken und
Artificial Intelligence
Technische Universität Wien
Favoritenstr. 9
A-1040 Vienna, Austria
Tel: +43-1-58801-18403
Fax: +43-1-58801-918403
sek@dbai.tuwien.ac.at
www.dbai.tuwien.ac.at

TECHNISCHE UNIVERSITÄT WIEN
Vienna University of Technology

DBAI TECHNICAL REPORT
DBAI TECHNICAL REPORT DBAI-TR-2021-120, 2021

Graph-Classes of Argumentation Frameworks with Collective Attacks

Wolfgang Dvořák [1] **Matthias König** [2] **Stefan Woltran** [3]

Abstract. Argumentation frameworks with collective attacks (SETAFs) have gained increasing attention in recent years as they provide a natural extension of the well-known abstract argumentation frameworks (AFs) due to Dung. Concerning complexity, it is known that for the standard reasoning tasks in abstract argumentation, SETAFs show the same behavior as AFs, i.e. they are mainly located on the first or second level of the polynomial hierarchy. However, while for AFs there is a rich literature on easier fragments, complexity analyses in this direction are still missing for SETAFs. In particular, the well-known graph-classes of acyclic AFs, even-cycle-free AFs, symmetric AFs, and bipartite AFs have been shown tractable. In this paper, we aim to extend these results to the more general notion of SETAFs. In particular, we provide various syntactic notions on SETAFs that naturally generalize the graph properties for directed hypergraphs, and perform a complexity analysis of the prominent credulous and skeptical acceptance problems for several different widely used semantics.

[1]Institute of Logic and Computation, TU Wien, Austria. E-mail: dvorak@dbai.tuwien.ac.at
[2]Institute of Logic and Computation, TU Wien, Austria. E-mail: mkoenig@dbai.tuwien.ac.at
[3]Institute of Logic and Computation, TU Wien, Austria. E-mail: woltran@dbai.tuwien.ac.at

1 Introduction

Formal argumentation provides formalisms to resolve conflicts in potentially inconsistent or incomplete knowledge, which is essential to draw conclusions of any kind in such a setting. In this context, argumentation frameworks (AFs), introduced in the influential paper by Dung [5], turned out to be a versatile system for reasoning tasks in an intuitive setting. In AFs we view arguments just as abstract entities, represented by nodes in a directed graph, independent from their internal structure. Conflicts are modeled in form of attacks between these arguments, constituting the edges of said graph representation. Different semantics have been defined for AFs and deliver sets of arguments that are jointly acceptable given the topology of attacks in the AF at hand. However, by their limited syntax it is hard to formalize certain naturally occurring statements in AFs, which is why various generalizations of the standard formalism have been proposed, see, e.g. [1]. One such generalization extends the syntax by *collective attacks*, i.e. a construction where a set T of arguments attacks an argument h, but no proper subset of T does; the resulting class of frameworks is often referred to as *SETAFs*. The underlying structure of SETAFs then is a *directed hypergraph*. When they introduced SETAFs [22], Nielsen and Parsons argued that collective attacks naturally appear in various contexts, e.g. when languages are not closed under conjunction. In fact, in certain settings standard AFs require artificial additional arguments and attacks, while the same setting can be natively represented in SETAFs. These observations have been backed up by recent practically driven investigations [25]. Moreover, SETAFs have been proven to be strictly more expressive than AFs, as shown in [11] by means of signatures. In spite of these advantages, there has not yet been much work on computational aspects of SETAFs. The general complexity of the most common reasoning tasks has been investigated in [12], where also an implementation of a solver for SETAFs with answer-set programming has been introduced. Moreover, algorithmic approaches for SETAFs have been studied in [15, 21].

The main aim of this paper is to deepen the complexity analysis of [12] which has shown that the complexity of SETAFs coincides with the results for classical AFs in general. In particular, this means that reasoning in many popular semantics is on the first or second level of the polynomial hierarchy. To still achieve manageable runtimes with large instances, the approach we shall take in this paper is to restrict the syntax of SETAFs. We propose certain constraints on the hypergraph structure such that the induced class of frameworks is easy to reason on (i.e. the problems in question are computable in *polynomial time*). On AFs this approach turned out to be fruitful: we say an AF is acyclic, symmetric, or bipartite, if its attack relation is, respectively. The thereby obtained graph classes are *tractable fragments* of AFs [2, 6, 7, 10]. Even though there exist translations from SETAFs to AFs [23, 19], it is not at all clear whether tractability results for AFs carry over to SETAFs. This is due to the fact that these translations can lead to an exponential blowup in the number of arguments; moreover certain structural properties are lost in the translation.

In what follows, we thus focus on defining graph properties for SETAFs "from scratch" - these can then be checked and exploited without a detour via AFs. Our main contributions can be summarized as follows:

- Novel definitions for graph classes of directed hypergraphs: these notions are conservative generalizations (i.e. in the special case of AFs they coincide with the respective classical notions) of well known properties of directed graphs such as acyclicity, symmetry, bipartiteness and 2-colorability. As a byproduct of the detailed analysis we state certain syntactical and semantical properties of SETAFs within these classes.

- We pinpoint the complexity of credulous and skeptical reasoning in the respective graph classes w.r.t. seven widely used argumentation semantics, that is admissible, grounded, complete, preferred, stable, stage, and semi-stable [22, 12, 19]. We provide (efficient) algorithms to reason on these computationally easy frameworks, and give negative results by providing hardness results for classes that yield no computational speedup.

- We establish the status of *tractable fragments* for the classes acyclicity, even-cycle-freeness, primal-bipartiteness, and self-attack-free full-symmetry. In fact, we not only show that these classes are easy to reason in, but the respective properties can also be recognized efficiently. This result allows one to perform such a check as a subroutine of a general-purpose SETAF-solver such that the overall asymptotic runtime is polynomial in case the input framework belongs to such a class.

Note that some proofs are not given in full length, they are carried out in detail in Appendix A.

2 Preliminaries

2.1 Argumentation Frameworks

Throughout the paper, we assume a countably infinite domain $\mathfrak{A}$ of possible arguments.

Definition 1. *A SETAF is a pair $SF = (A, R)$ where $A \subseteq \mathfrak{A}$ is finite, and $R \subseteq (2^A \setminus \{\emptyset\}) \times A$ is the attack relation. For an attack $(T, h) \in R$ we call T the* tail *and h the* head *of the attack. SETAFs (A, R), where for all $(T, h) \in R$ it holds that $|T| = 1$, amount to (standard Dung) AFs. In that case, we usually write (t, h) to denote the set-attack $(\{t\}, h)$.*

Given a SETAF (A, R), we write $S \mapsto_R a$ if there is a set $T \subseteq S$ with $(T, a) \in R$. Moreover, we write $S' \mapsto_R S$ if $S' \mapsto_R a$ for some $a \in S$. We drop subscript R in $\mapsto_R$ if there is no ambiguity. For $S \subseteq A$, we use S^+_R to denote the set $\{a \mid S \mapsto_R a\}$ and define the range *of S (w.r.t. R), denoted $S^\oplus_R$, as the set $S \cup S^+_R$.*

Example 1. *Consider the SETAF $SF = (A, R)$ with $A = \{a, b, c, d\}$ and $R = \{(\{a, b\}, c), (\{a, c\}, b), (\{c\}, d)\}$. For an illustration see Figure 1a - the dashed attacks are collective attacks.*

We will now define special 'kinds' of attacks and fix the notions of redundancy-free and self-attack-free SETAFs.

Definition 2. *Given a SETAF $SF = (A, R)$, an attack $(T, h) \in R$ is* redundant *if there is an attack $(T', h) \in R$ with $T' \subset T$. A SETAF without redundant attacks is* redundancy-free*. An attack $(T, h) \in R$ is a* self-attack *if $h \in T$. A SETAF without self-attacks attacks is* self-attack-free.

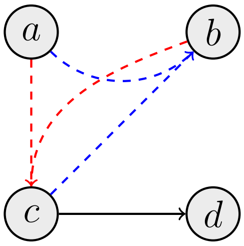

(a) SETAF SF

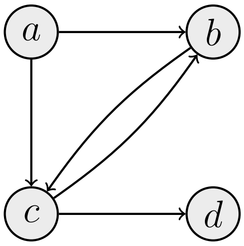

(b) $primal(SF)$

Figure 1: An example SETAF and its primal graph.

Redundant attacks can be efficiently detected and then be omitted without changing the standard semantics [16, 23]. In the following we always assume redundancy-freeness for all SETAFs, unless stated otherwise. The well-known notions of conflict and defense from classical Dung-style-AFs naturally generalize to SETAFs.

Definition 3. *Given a SETAF $SF = (A, R)$, a set $S \subseteq A$ is* conflicting *in SF if $S \mapsto_R a$ for some $a \in S$. A set $S \subseteq A$ is* conflict-free *in SF, if S is not conflicting in SF, i.e. if $T \cup \{h\} \not\subseteq S$ for each $(T, h) \in R$. cf(SF) denotes the set of all conflict-free sets in SF.*

Definition 4. *Given a SETAF $SF = (A, R)$, an argument $a \in A$ is* defended *(in SF) by a set $S \subseteq A$ if for each $B \subseteq A$, such that $B \mapsto_R a$, also $S \mapsto_R B$. A set $T \subseteq A$ is defended (in SF) by S if each $a \in T$ is defended by S (in SF).*

The semantics we study in this work are the grounded, admissible, complete, preferred, stable, stage and semi-stable semantics, which we will abbreviate by *grd*, *adm*, *com*, *pref*, *stb*, *stage* and *sem* respectively [12, 19, 22].

Definition 5. *Given a SETAF $SF = (A, R)$ and a conflict-free set $S \in cf(SF)$. Then,*

- *$S \in adm(SF)$, if S defends itself in SF,*
- *$S \in com(SF)$, if $S \in adm(SF)$ and $a \in S$ for all $a \in A$ defended by S,*
- *$S \in grd(SF)$, if $S = \bigcap_{T \in com(SF)} T$,*
- *$S \in pref(SF)$, if $S \in adm(SF)$ and there is no $T \in adm(SF)$ s.t. $T \supset S$,*
- *$S \in stb(SF)$, if $S \mapsto a$ for all $a \in A \setminus S$,*
- *$S \in stage(SF)$, if $\nexists T \in cf(SF)$ with $T_R^\oplus \supset S_R^\oplus$, and*
- *$S \in sem(SF)$, if $S \in adm(SF)$ and $\nexists T \in adm(SF)$ s.t. $T_R^\oplus \supset S_R^\oplus$.*

Table 1: Extensions of the example SETAF SF from Example 1.

σ	$\sigma(SF)$
cf	$\{\emptyset, \{a\}, \{b\}, \{c\}, \{d\}, \{a,b\}, \{a,c\}, \{a,d\}, \{b,c\}, \{b,d\}, \{a,b,d\}\}$
adm	$\{\emptyset, \{a\}, \{a,b\}, \{a,c\}, \{a,b,d\}\}$
com	$\{\{a\}, \{a,c\}, \{a,b,d\}\}$
grd	$\{\{a\}\}$
pref/*stb*/*stage*/*sem*	$\{\{a,c\}, \{a,b,d\}\}$

For an example of the extensions of a SETAF see Table 1. The relationship between the semantics has been clarified in [12, 19, 22] and matches with the relations between the semantics for Dung AFs, i.e. for any SETAF SF:

$$stb(SF) \subseteq sem(SF) \subseteq pref(SF) \subseteq com(SF) \subseteq adm(SF) \subseteq cf(SF) \tag{1}$$

$$stb(SF) \subseteq stage(SF) \subseteq cf(SF). \tag{2}$$

The following property also carries over from Dung AFs: For any SETAF SF, if $stb(SF) \neq \emptyset$ then $stb(SF) = sem(SF) = stage(SF)$.

2.2 Complexity

We assume the reader to have basic knowledge in computational complexity theory[1], in particular we make use of the complexity classes L (logarithmic space), P (polynomial time), NP (non-deterministic polynomial time), coNP, Σ_2^{P} and Π_2^{P}. For a given SETAF $SF = (A, R)$ and an argument $a \in A$, we consider the standard reasoning problems (under semantics σ) in formal argumentation:

- Credulous acceptance $Cred_\sigma$: Is the argument a contained in at least one σ extension of SF?, and
- Skeptical acceptance $Skept_\sigma$: Is the argument a contained in all σ extensions of SF?

The complexity landscape of SETAFs coincides with that of Dung AFs and is depicted in Table 2. As SETAFs generalize Dung AFs the hardness results for Dung AFs [2, 4, 8, 9, 17, 18] (for a survey see [10]) carry over to SETAFs. Also the same upper bounds hold for SETAFs [12]. However, while the complexity results for AFs can be interpreted as complexity w.r.t. the number of arguments $|A|$, the complexity results for SETAFs should be understood as complexity w.r.t. $|A| + |R|$ (as $|R|$ might be exponentially larger than $|A|$).

[1] For a gentle introduction to complexity theory in the context of formal argumentation, see [10].

Table 2: Complexity for AFs and SETAFs (C-c denotes completeness for C).

	grd	*adm*	*com*	*pref*	*stb*	*stage*	*sem*
$Cred_\sigma$	P-c	NP-c	NP-c	NP-c	NP-c	Σ_2^P-c	Σ_2^P-c
$Skept_\sigma$	P-c	trivial	P-c	Π_2^P-c	coNP-c	Π_2^P-c	Π_2^P-c

3 Graph Classes

The directed hypergraph-structure of SETAFs is rather specific and to the best of the authors' knowledge the hypergraph literature does not provide generalizations of common graph classes to this kind of directed hypergraphs. Thus we first identify such generalizations for SETAFs for the graph classes of interest. Then, we show the tractability of acyclicity and even-cycle-freeness (the latter does not hold for stage semantics) in SETAFs, and that odd-cycle-freeness lowers the complexity to the first level of the polynomial hierarchy as for AFs. Then, we adapt the notion of symmetry in different natural ways, only one of which will turn out to lower the complexity of reasoning as with symmetric AFs. Finally, we will adapt and analyze the notions of bipartiteness and 2-colorability. Again we will see a drop in complexity only for a particular definition of this property on hypergraphs. All of the classes generalize classical properties of directed graphs in a way for SETAFs such that in the special case of AFs (i.e. for SETAFs where for each attack (T, h) the tail T consists of exactly one argument) they coincide with said classical notions, respectively. Finally, we will argue that these classes are not only efficient to reason on, but are also efficiently recognizable. Hence, we can call them *tractable fragments of argumentation frameworks with collective attacks*.

When defining these classes we will use the notion of the *primal graph*, an implementation of the hypergraph structure of a SETAF into a directed graph. An illustration is given in Figure 1.

Definition 6. *Given a SETAF $SF = (A, R)$. Then its* primal graph *is defined as $primal(SF) = (A', R')$, where $A' = A$, and $R' = \{(t, h) \mid (T, h) \in R, t \in T\}$.*

3.1 Acyclicity

Akin to cycles in AFs, we define cycles on SETAFs as a sequence of arguments such that there is an attack between each consecutive argument.

Definition 7. *A* cycle *C of length $|C| = n$ is a sequence of pairwise distinct arguments $C = (a_1, a_2, \ldots, a_n, a_1)$ such that for each a_i there is an attack (A_i, a_{i+1}) with $a_i \in A_i$, and there is an attack (A_n, a_1) with $a_n \in A_n$. A SETAF is* cyclic *if it contains a cycle (otherwise it is* acyclic*),* even-cycle-free *if it contains no cycles of even length, and* odd-cycle-free *if it contains no cycles of odd length.*

Note that a SETAF SF is acyclic if and only if its primal graph $primal(SF)$ is acyclic. It can easily be seen that acyclic SETAFs are well founded [22], i.e. there is no infinite sequence of sets $B_1, B_2, \ldots$, such that for all i, B_i is the tail of an attack towards an argument in B_{i-1}. As shown in [22], this means grounded, complete, preferred, and stable semantics coincide. Moreover, as therefore there always is at least one stable extension, stable, semi-stable and stage semantics coincide as well, and the lower complexity of $Cred_{grd}$ and $Skept_{grd}$ carries over to the other semantics. Together with the hardness from AFs, we immediately obtain our first result.

Theorem 1. *For acyclic SETAFs the problems $Cred_\sigma$ and $Skept_\sigma$ for $\sigma \in \{grd, com, pref, stb, stage, sem\}$ are* P*-complete. Moreover $Cred_{adm}$ is* P*-complete.*

For AFs we have that the absence of even-length cycles forms a tractable fragment for all semantics under our consideration but stage. The key lemma is that every AF with more than one complete extension has to have a cycle of even length [9]. This property also holds for SETAFs, which in turn means even-cycle-free SETAFs have exactly one complete extension, namely the grounded extension, which is then also the only preferred and semi-stable extension. Our proof of this property follows along the lines of the respective known proof for AFs. Moreover, the grounded extension is the only candidate for a stable extension, and thus for reasoning with stable semantics it suffices to check whether the grounded extension is stable. Finally, note that the hardness of $Cred_{stage}$ and $Skept_{stage}$ carries over from AFs (cf. [10]) to SETAFs.

Theorem 2. *For even-cycle-free SETAFs the problems $Cred_\sigma$ and $Skept_\sigma$ for $\sigma \in \{com, pref, stb, sem\}$ are* P*-complete. Moreover the problem $Cred_{adm}$ is* P*-complete, the problem $Cred_{stage}$ is Σ_2^{P}-complete, and the problem $Skept_{stage}$ is Π_2^{P}-complete.*

For odd-cycle free SETAFs the situation is just like with odd-cycle-free AFs [8]. If there is a sequence of arguments $(a_1, a_2, \dots)$, we say a_1 *indirectly attacks* the arguments a_{2*i-1} and *indirectly defends* the arguments a_{2*i} for $i \geq 1$ (cf. [22]). As odd-cycle-free SETAFs are *limited controversial* [22], i.e. there is no infinite sequence of arguments such that each argument indirectly attacks and defends the next, they are coherent, i.e. stable and preferred semantics coincide, and therefore we experience a drop of the complexity to the first level of the polynomial hierarchy.

Theorem 3. *For odd-cycle-free SETAFs the problems $Cred_\sigma$ for $\sigma \in \{adm, stb, pref, com, stage, sem\}$ are* NP*-complete, problems $Skept_\sigma$ for $\sigma \in \{stb, pref, stage, sem\}$ are* coNP*-complete, and the problems $Cred_{grd}$, $Skept_{grd}$, and $Skept_{com}$ are* P*-complete.*

3.2 Symmetry

In the following we provide two generalizations of symmetry[2] for SETAFs. The first definition via the primal graph is inspired by the notion of counter-attacks: an AF $F = (A, R)$ is symmetric if for every attack $(a, b) \in R$ there is a counter-attack $(b, a) \in R$. As we will show, the corresponding definition for SETAFs is not sufficiently restrictive to lower the complexity of the reasoning problems in questions, except for a fast way to decide whether an argument is in the grounded extension or not. For an illustration of the following definitions see Figure 2.

Definition 8. *A SETAF $SF = (A, R)$ is* primal-symmetric *iff for every attack $(T, h) \in R$ and $t \in T$ there is an attack $(H, t) \in R$ with $h \in H$.*

[2]Further symmetry-notions for SETAFs have been investigated in [20].

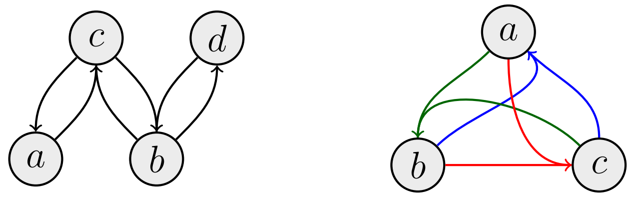

(a) Primal-symmetry (b) Full-symmetry

Figure 2: Different notions of symmetry.

As expected, a SETAF is primal-symmetric iff its primal graph is symmetric. Notice that the notion of primal-symmetry coincides with the definition of symmetry of Abstract Dialectical Frameworks in [3]. The next notion intuitively captures the "omnidirectionality" of symmetric attacks: for every attack all involved arguments have to attack each other. In the definition of fully-symmetry we distinguish between self-attacks and attacks which are not self-attacks.

Definition 9. *A SETAF* $SF = (A, R)$ *is* fully-symmetric *iff for every attack* $(T, h) \in R$ *we either have*

- *if* $h \in T$*, then* $\forall x \in T$ *it holds* $(T, x) \in R$*, or*
- *if* $h \notin T$*, then* $\forall x \in S$ *it holds* $(S \setminus \{x\}, x) \in R$ *with* $S = T \cup \{h\}$*.*

We have that every fully-symmetric SETAF is primal-symmetric, the converse does not hold. In symmetric AFs every argument defends itself against all incoming attacks, hence, admissible sets coincide with conflict-free sets, and it becomes computationally easy to reason on admissible, complete, and preferred extensions. However, this is not the case with our notions of symmetry for SETAFs. Consider the fully-symmetric (and thus also primal-symmetric) SETAF from Figure 2b: we have that for example the singleton set $\{a\}$ is conflict-free, but $\{a\}$ cannot defend itself against the attacks towards a. That is, the argument for tractability from AFs does not transfer to SETAFs. This corresponds to the the fact that we will obtain full hardness for the admissibility-based semantics in question, when making no further restrictions on the graph structure.

For both notions of symmetry we have that an argument is in the grounded extension iff it is not in the head of any attack, which can easily be checked in logarithmic space. This is by the characterization of the grounded extension as least fixed point of the *characteristic function*[22], i.e. the grounded extension can be computed by starting from the empty set and iteratively adding all defended arguments. For primal-symmetric SETAFs with and without self-attacks, as well as fully-symmetric SETAFs (allowing self-attacks) this is the only computational speedup we can get, the remaining semantics maintain their full complexity.

In order to show the hardness for primal-symmetric SETAFs we provide a translation that transforms each SETAF $SF = (A, R)$ in a primal-symmetric SETAF SF': we construct SF' from SF by adding, for each attack $r = (T, h)$ and $t \in T$, mutually attacking arguments $a^1_{r,t}, a^2_{r,t}$, the (ineffective) counter-attack $(\{a^1_{r,t}, a^2_{r,t}, h\}, t)$, and attacks $(t, a^1_{r,t}), (t, a^2_{r,t})$. It can be verified that the resulting SETAF SF' is primal-symmetric, does not introduce self-attacks and preserves the acceptance status of the original arguments.

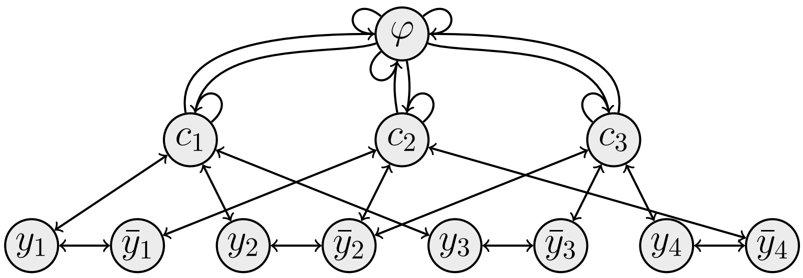

Figure 3: Illustration of SF^1_φ for a formula φ with atoms $Y = \{y_1, y_2, y_3, y_4\}$, and clauses $C = \{\{y_1, y_2, y_3\}, \{\bar{y}_1, \bar{y}_2, \bar{y}_4)\}, \{\bar{y}_2, \bar{y}_3, y_4\}\}$.

Theorem 4. *For primal-symmetric SETAFs (with or without self-attacks) the problems $Cred_{grd}$, $Skept_{grd}$ and $Skept_{com}$ are in* L*, the complexity of the other problems under our consideration coincides with the complexity for the general problems (see Table 2).*

We will see the same hardness results for fully-symmetric SETAFs, but here the hardness relies on the use of self-attacks. Stable, stage, and semi-stable semantics have already their full complexity in symmetric AFs allowing self-attacks [10]. For the admissible, complete and preferred semantics, hardness can be shown with adjustments to the standard reductions. That is, we substitute some of the occurring directed attacks (a, b) by classical symmetric attacks $(a, b), (b, a)$, and others by symmetric self-attacks $(\{a, b\}, a), (\{a, b\}, b)$. For instance, for admissible semantics, given a CNF-formula φ with clauses C over atoms Y we define $SF^1_\varphi = (A', R')$ (cf. Figure 3), with $A' = \{\varphi\} \cup C \cup Y \cup \bar{Y}$ and R' given by (a) the usual attacks $\{(y, \bar{y}), (\bar{y}, y) \mid y \in Y\}$, (b) symmetric attacks from literals to clauses $\{(y, c), (c, y) \mid y \in c, c \in C\} \cup \{(\bar{y}, c), (c, \bar{y}) \mid \bar{y} \in c, c \in C\}$, and (c) the symmetric self-attacks $\{(\{c, \varphi\}, \varphi), (\{c, \varphi\}, c) \mid c \in C\}$. The attacks (c) ensure that all c have to be attacked in order to accept φ and that all c are unacceptable.

Theorem 5. *For fully-symmetric SETAFs (allowing self-attacks) the problems $Cred_{grd}$, $Skept_{grd}$ and $Skept_{com}$ are in* L*, the complexity of credulous and skeptical acceptance for the other semantics under our consideration coincides with the complexity for the general problems (see Table 2).*

Investigations on symmetric AFs often distinguish between AFs with and without self-attacks [10]. Indeed, also for *self-attack-free* fully-symmetric SETAFs we have that all naive extensions (i.e. $\subseteq$-maximal conflict-free sets) are stable, hence, one can construct a stable extension containing an arbitrary argument a by starting with the conflict-free set $\{a\}$ and expanding it to a maximal conflict-free set. As stable extensions are admissible, complete, preferred, stage, and semi-stable, an argument is trivially credulously accepted w.r.t. these semantics. Similarly, it is easy to decide whether an argument is in all extensions.

Theorem 6. *For self-attack-free fully-symmetric SETAFs the problems $Cred_\sigma$ are trivially true for $\sigma \in \{adm, com, pref, stb, stage, sem\}$. The problems $Skept_\sigma$ are in* L *for $\sigma \in \{grd, com, pref, stb, stage, sem\}$. Moreover, $Cred_{grd}$ is in* L.

Figure 4: Different notions of bipartiteness.

3.3 Bipartiteness

In the following we will provide two generalizations of bipartiteness; the first - primal-bipartiteness - extends the idea of partitioning for directed hypergraphs, the second is a generalization of the notion of 2-colorability. In directed graphs bipartiteness and 2-colorability coincide. However, this is not the case in SETAFs with their directed hypergraph-structure. As it will turn out, 2-colorability is not a sufficient condition for tractable reasoning, whereas primal-bipartiteness makes credulous and skeptical reasoning P-easy. For an illustration of the respective definitions see Figure 4.

Definition 10. *Let $SF = (A, R)$ be a SETAF. Then SF is primal-bipartite iff its primal graph $primal(SF)$ is bipartite, i.e. iff there is a partitioning of A into two sets (Y, Z), such that*

- *$Y \cup Z = A$, $Y \cap Z = \emptyset$, and*
- *for every $(T, h) \in R$ either $h \in Y$ and $T \subseteq Z$, or $h \in Z$ and $T \subseteq Y$.*

For bipartite AFs, Dunne provided an algorithm to enumerate the arguments that appear in admissible sets [6]; this algorithm can be adapted for SETAFs (see Algorithm 1). Intuitively, the algorithm considers the two sets of the partition separately. For each partition it iteratively removes arguments that cannot be defended, and eventually ends up with an admissible set. The union of the two admissible sets then forms a superset of every admissible set in the SETAF. As primal-bipartite SETAFs are odd-cycle-free, they are coherent [22], which means preferred and stable extensions coincide. This necessarily implies the existence of stable extensions, which means they also coincide with stage and semi-stable extensions. These results suffice to pin down the complexity of credulous and skeptical reasoning for the semantics under our consideration.

Theorem 7. *For primal-bipartite SETAFs the problems $Cred_\sigma$ and $Skept_\sigma$ for $\sigma \in \{com, pref, stb, stage, sem\}$ are* P*-complete. Moreover the problem $Cred_{adm}$ is* P*-complete.*

It is noteworthy that the complexity of deciding whether a set S of arguments is *jointly* credulously accepted w.r.t. preferred semantics in primal-bipartite SETAFs was already shown to be NP-complete for bipartite AFs (and, hence, for SETAFs) in [6]; however, this only holds if the arguments in question distribute over both partitions - for arguments that are all within one partition this problem is in P, which directly follows from the fact that Algorithm 1 returns the set Y_i of credulously accepted arguments - which is itself an admissible set.

Algorithm 1: Compute the set of credulously accepted arguments w.r.t. *pref* semantics

Input : A primal-bipartite SETAF $SF = (A, R)$ with a partitioning (Y, Z)
Output: The admissible set Y_i of credulously accepted arguments in Y

$i := 0$
$Y_0 := Y$
$R_0 := R$
repeat
  $i := i + 1$
  $Y_i := Y_{i-1} \setminus \{y \mid y \in Y_{i-1},$ there is some $(Z', y) \in R_{i-1}$ with $Z' \subseteq Z$ such that $\forall z \in Z' \; |\{(Y', z) \mid (Y', z) \in R_{i-1}\}| = 0\}$
  $R_i := R_{i-1} \setminus \{(Y', z) \mid Y' \subseteq Y, z \in Z, Y' \not\subseteq Y_i\}$
until $Y_i = Y_{i-1}$;

It is natural to ask whether the more general notion of 2-colorability also yields a computational speedup. We capture this property for SETAFs by the following definition:

Definition 11. *Let $SF = (A, R)$ be a SETAF. Then SF is 2-colorable iff there is a partitioning of A into two sets (Y, Z), such that*

- *$Y \cup Z = A$, $Y \cap Z = \emptyset$, and*
- *for every attack $(T, h) \in R$ we have $(T \cup \{h\}) \cap Y \neq \emptyset$ and $(T \cup \{h\}) \cap Z \neq \emptyset$.*

Note that both primal-bipartiteness and 2-colorability do not allow self-loops (a, a) with a single argument in the tail, but 2-colorable SETAFs may contain self-attacks (T, h) with $|T| \geq 2$.

For admissibility-based semantics that preserve the grounded extension (such as *grd*, *com*, *pref*, *stb*, *sem*) it is easy to see that the problems remain hard in 2-colorable SETAFs: intuitively, one can add two fresh arguments to any SETAF and add them to the tail T of every attack (T, h) - they will be in each extension of the semantics in question, and other than that the extensions will coincide with the original SETAF (this translation is *faithful*, cf. [18]). To establish hardness for stage semantics we can adapt the existing reductions by replacing self-attacking arguments by a construction with additional arguments such that 2-colorability is ensured, and replace certain classical AF-attacks by collective attacks.

Theorem 8. *For 2-colorable SETAFs the complexity of $Cred_\sigma$ and $Skept_\sigma$ for all semantics under our consideration coincides with the complexity of the general problem (see Table 2).*

3.4 Tractable Fragments

The (relatively speaking) low complexity of reasoning in SETAFs with the above described features on its own is convenient, but to be able to fully exploit this fact we also show that these classes are easily *recognizable*. As mentioned in [13], the respective AF-classes can be efficiently decided by graph algorithms. As for acyclicity, even-cycle-freeness, and primal-bipartiteness it suffices

Table 3: Tractable fragments in SETAFs.

		grd	*adm*	*com*	*pref*	*stb*	*stage*	*sem*
General	$Cred_\sigma$	P-c	NP-c	NP-c	NP-c	NP-c	Σ_2^{P}-c	Σ_2^{P}-c
	$Skept_\sigma$	P-c	trivial	P-c	Π_2^{P}-c	coNP-c	Π_2^{P}-c	Π_2^{P}-c
Acyclicty	$Cred_\sigma$	P-c	P-c	P-c	P-c	P-c	P-c	P-c
	$Skept_\sigma$	P-c	trivial	P-c	P-c	P-c	P-c	P-c
Even-cycle-freeness	$Cred_\sigma$	P-c	P-c	P-c	P-c	P-c	Σ_2^{P}-c	P-c
	$Skept_\sigma$	P-c	trivial	P-c	P-c	P-c	Π_2^{P}-c	P-c
self-attack-free	$Cred_\sigma$	in L	trivial	trivial	trivial	trivial	trivial	trivial
full-symmetry	$Skept_\sigma$	in L	trivial	in L	in L	in L	in L	in L
Primal-bipartiteness	$Cred_\sigma$	P-c	P-c	P-c	P-c	P-c	P-c	P-c
	$Skept_\sigma$	P-c	trivial	P-c	P-c	P-c	P-c	P-c

to analyze the primal graph, these results carry over to SETAFs. Moreover, for primal-bipartite SETAFs we can efficiently compute a partitioning, which is needed as input for Algorithm 1. Finally, we can test for full-symmetry efficiently as well: one (naive) approach is to just loop over all attacks and check whether there are corresponding attacks towards each involved argument. Likewise, a test for self-attack-freeness can be performed efficiently. Summarizing the results of this work, we get the following theorem.

Theorem 9. *Acyclicity, even-cycle-freeness, self-attack-free full-symmetry, and primal-bipartiteness are* tractable fragments *for SETAFs.*

In particular, for credulous and skeptical reasoning in the semantics under our consideration the complexity landscape including tractable fragments in SETAFs is depicted in Table 3.

4 Conclusion

In this work, we introduced and analyzed various different syntactic classes for SETAFs. These new notions are conservative generalizations of properties of directed graphs, namely acyclicity, even/odd-cycle-freeness, symmetry, and bipartiteness, which have been shown to lower the complexity for acceptance problems of AFs. The starting point for our definitions is the *primal graph* of the SETAF, a structural embedding of directed hypergraph into a directed graph. Other than establishing basic properties, we performed a complete complexity analysis for credulous and skeptical reasoning in classes of SETAFs with these generalized properties.

For the notions regarding cycles, we established the same properties for acyclicity, even-cycle-freeness, and odd-cycle-freeness for SETAFs that also hold for AFs. This includes the fact that the same upper and lower bounds on the complexity holds as in AFs, namely reasoning in acyclicity becomes tractable for all semantics under our consideration, even-cycle-freeness becomes tractable for all semantics but stage, and in odd-cycle-free SETAFs the complexity drops to the first level of the polynomial hierarchy. The symmetry notions we introduced generalize the concept of counter-attacks. We have established that a symmetric primal graph is not a sufficient condition for a SETAF to lower the complexity. The more restricting notion of full-symmetry yields a drop in complexity, but only if one also requires the SETAFs to be self-attack-free. Allowing self-attacks, even this notion does not yield a drop in the complexity for the semantics in question, which is the case for admissible, preferred, and complete semantics in AFs. We also investigated notions of bipartiteness. While in directed graphs bipartiteness and 2-colorability coincide, this in not the case in directed hypergraphs. We provided an algorithm that allows one to reason efficiently on primal-bipartite SETAFs, a result that does not apply for the more general notion of 2-colorable SETAFs. Finally, we argued that these classes can also be efficiently recognized, which is a crucial condition if one wants to implement the more efficient algorithms as a sub-routine of a general SETAF-solver.

In the future, tractability for SETAFs could be established by performing parametrized complexity analysis, as it has been done for AFs [10, 14]. In particular, we understand these results as a starting point for investigations in terms of backdoors (i.e. measuring and exploiting a bounded *distance* of a given SETAF to a certain tractable class), along the lines of similar investigations for AFs [13]. Moreover, it is important to analyze whether SETAFs that occur in applications belong to any of the graph-classes introduced in this work. For example, it can be checked that the frameworks generated for a particular application in [25]—even though they do not belong to one of our tractable fragments—enjoy a (weak) symmetry-property, which allows one to reason in L on the grounded extension. This can be shown using the same proof as for our primal-symmetry result. Finally, as the purpose of the algorithms featured in this work was solely to illustrate the membership to the respective complexity classes, undoubtedly they yield a potential for improvement and optimization.

Acknowledgments

This research has been supported by the Vienna Science and Technology Fund (WWTF) through project ICT19-065, and by the Austrian Science Fund (FWF) through projects P30168 and P32830.

A Proof Details

In order to carry out the missing proofs in detail, we introduce some additional notation. *Translations* (cf. [18]) are a special kind of reduction, they allow us to easily reduce decision problems on SETAFs to other decision problems on SETAFs.

A (SETAF-)translation Tr is a function that takes a SETAF SF and outputs another SETAF $Tr(SF)$ with specific properties; if every SETAF in the range of the translation shares a syntactic property and certain semantic restrictions are preserved (i.e. certain relationships between the extensions apply), we obtain a hardness-result for the class that is defined by the syntactic property. First we will adapt the syntactic notions for translations that are originally defined for AFs (see [18]) for SETAFs.

Definition 12. *A* (SETAF-)translation Tr *is a function which maps SETAFs to SETAFs. Let* $SF = (A, R)$ *be a SETAF and* $Tr(SF) = (A', R')$ *its translated image. Moreover let* σ, σ' *be two semantics. A translation* Tr *is called*

- efficient *if for every* SF *its* $Tr(SF)$ *can be computed efficiently,*
- embedding *if for every* $SF = (A, R)$ *we have* $A \subseteq A'$ *and* $R = R' \cap ((2^A \setminus \emptyset) \times A)$.
- exact *for* $\sigma \Rightarrow \sigma'$ *if for every SETAF* SF *we have* $\sigma(SF) = \sigma'(Tr(SF))$,
- faithful *for* $\sigma \Rightarrow \sigma'$ *if for every SETAF* SF *we have* $\sigma(SF) = \{E \cap A \mid E \in \sigma'(Tr(SF))\}$ *and* $|\sigma(SF)| = |\sigma'(Tr(SF))|$,
- acceptance-preserving *for* $\sigma \Rightarrow \sigma'$ *if for every SETAF* SF *we have* $\sigma(SF) = \{E \cap A \mid E \in \sigma'(Tr(SF))\}$.

We have that every exact translation is faithful, and every faithful translation is acceptance-preserving. As we are often not so much interested in the translatability between different semantics, but in the syntactic properties of the translations, we often use translations where $\sigma = \sigma'$, then we just write "Tr is an exact/faithful/acceptance-preserving translation for σ".

Moreover, we need the *naive* semantics to establish some results. A naive extension is a $\subseteq$-maximal conflict-free set; the set of all naive extensions of a SETAF SF is denoted by *naive*(SF). It is known that both $Cred_{naive}$ and $Skept_{naive}$ are in L [12].

Finally, we will distinguish between *active* and *inactive attacks* (cf. [16]). We call an attack (T, a) inactive if there is another attack (S, b) with $S \cup \{b\} \subseteq T$ and all other attacks active. The intuition is that inactive attacks can never be used to defend an argument in a extension and do not contribute to the range of extensions. However, these attacks are still relevant as extensions have to defend their arguments against inactive attacks.

A.1 Proof of Theorem 4

We start by providing a translation for $\sigma \in \{cf, adm, stb, pref, stage, sem\}$ such that for every self-attack-free SETAF SF its translation is primal-symmetric. Moreover we will establish that this translation is efficient and acceptance-preserving. Note that we use Tr_1 only for self-attack-free SETAFs. For an illustration of Tr_1 see Figure 5.

Translation 1. *Let $SF = (A, R)$ be a SETAF. The SETAF-translation Tr_1 is defined as $Tr_1(SF) = (A', R')$ with*

$$\begin{aligned} A' &= A \cup \{a^1_{r,t}, a^2_{r,t} \mid r = (T, h) \in R, t \in T\}, \\ R' &= R \cup \{(a^1_{r,t}, a^2_{r,t}), (a^2_{r,t}, a^1_{r,t}), (\{a^1_{r,t}, a^2_{r,t}, h\}, t), \\ &\qquad (t, a^1_{r,t}), (t, a^2_{r,t}) \mid r = (T, h) \in R, t \in T\} \end{aligned}$$

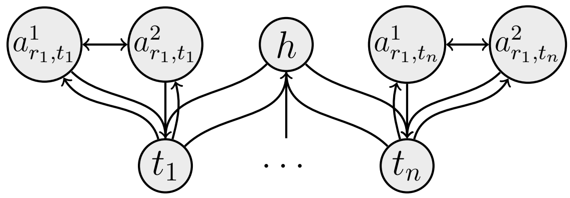

Figure 5: Illustration Translation Tr_1. We have a SETAF with one attack $r_1 = (\{t_1, \dots, t_n\}, h)$.

Intuitively, for each attack in the original SETAF the translation Tr_1 adds an attack towards every attacker. Since the arguments a^1_r and a^2_r attack each other, this added attack is inactive, i.e. cannot be used to defend arguments or extend the range of an extension. Hence, most semantics do not change their extensions. Also in order to preserve admissibility we add an attack towards the added arguments. We have that Tr_1 is efficient and embedding.

Lemma 1. *Let $SF = (A, R)$ be a SETAF and let $SF' = (A', R') = Tr_1(SF)$. Then for every $E' \in cf(SF')$ we have for $E = E' \cap A$ that $E^{\oplus}_R = E'^{\oplus}_{R'} \cap A$.*

Proof. "$\subseteq$": Immediate by the fact that Tr_1 is embedding and the monotonicity of $(.)^{\oplus}$.
"$\supseteq$": Note that the set of active attacks towards arguments in A in SF' is the set of active attacks in SF. The only active attacks towards arguments in A in SF' are from within A. The fact that in the construction of SF' no further attacks between arguments in A is added concludes the proof. □

This brings us to our first result that will allow us to settle the complexity of reasoning for $\sigma \in \{cf, adm, stb, pref, stage, sem\}$ in self-attack-free primal-symmetric SETAFs.

Lemma 2. *Let $\sigma \in \{cf, adm, stb, pref, stage, sem\}$. Then Tr_1 is an acceptance-preserving translation for $\sigma \Rightarrow \sigma$ such that for every self-attack-free SETAF SF its translation $SF' = Tr_1(SF) = (A', R')$ is primal-symmetric.*

Proof. First of all it is easy to verify that SF' is indeed primal-symmetric: For each of the original attacks (T, h) and $t \in T$ there is an attack $(\{a^1_{r,t}, a^2_{r,t}, h\}, t)$ which is accompanied by $(t, a^1_{r,t}), (t, a^2_{r,t})$ in order to make the new attack symmetric. Obviously also the attacks $(a^1_{r,t}, a^2_{r,t})$, $(a^2_{r,t}, a^1_{r,t})$ are symmetric.

We show that the translations is an acceptance-preserving translation separately for each of the semantics. We follow the following scheme for each of the semantics σ: firstly we will show constructively that for any extension $E \in \sigma(SF)$ there exists an extension $E' \in \sigma(SF')$ such that $E' \cap A = E$ ("⇒"). Secondly we will show that for each extension $E' \in \sigma(SF')$ the corresponding extension $E = E' \cap A$ is an extension $E \in \sigma(SF)$ ("⇐").

1. For $\sigma =$ *cf*:
   "⇒": Let $E \in$ *cf*(SF). Then also $E \in$ *cf*(SF'), as there are no attacks between elements of A that are added in the construction.
   "⇐": Let $E' \in$ *cf*(SF) and let $E = E' \cap A$. Then $E \in$ *cf*(SF), as there can be no attack between arguments in A.

2. For $\sigma =$ *adm*:
   "⇒": Let $E \in$ *adm*(SF) and let $E' = E \cup \{a^1_{r,t} \mid r = (T,h), t \in T, E \mapsto_R t\}$. By construction we have $E' \in$ *cf*(SF'). Assume towards contradiction some $a \in E'$ is not defended by E', i.e. there is an attack $(T, a) \in R'$ such that $E' \not\mapsto_{R'} T$. This means either $a \in A' \setminus A$ or $a \in A$. In the first case we have $a = a^1_{r,t}$ for some $r = (T, h) \in R$ with $t \in T$. We have that a defends itself against the attack from $a^2_{r,t}$, the only remaining attack towards a is from t. But since $a \in E'$, by construction we have $E \mapsto_R t$, which also means $E' \mapsto_{R'} t$, so a is defended by E', which is a contradiction. In the second case we have $a \in A$. Since $a \in E$ and $E \in$ *adm*(SF) we know that a is defended against all attacks in R, i.e. all attacks from within A. But since the only active attacks towards a are from within A, we have that a is defended, which is a contradiction.
   "⇐": Let $E' \in$ *adm*(SF') and let $E = E' \cap A$. We know $E \in$ *cf*(SF). Let $a \in E$ and let $(T, a) \in R$ be an attack towards a. Since E' is admissible in SF' we have $E' \mapsto_{R'} T$, i.e. there is an attack $(T', t) \in R'$ such that $t \in T$ and $T' \subseteq E'$. Since the only active attacks towards t are from within A, we also have that $E \mapsto_R t$, which means a is defended by E in SF.

3. For $\sigma =$ *stb*:
   "⇒": Let $E \in$ *stb*(SF) and let $E' = \{a^1_{r,t} \mid r = (T,h), t \in T, t \notin E\}$. We have $E \in$ *cf*(SF') by construction. Moreover, since $E \in$ *stb*(SF), by Lemma 1 we have $A \subseteq E'^{\oplus}_{R'}$, and by construction we have $A' \setminus A \subseteq E'^{\oplus}_{R'}$.
   "⇐": Let $E' \in$ *stb*(SF') and let $E = E' \cap A$. We know $E \in$ *cf*(SF), and, since $E' \in$ *stb*(SF'), by Lemma 1 we have $A \subseteq E^{\oplus}_R$.

4. For $\sigma =$ *pref*:
   "⇒": Let $E \in$ *pref*(SF) and let $E' = \{a^1_{r,t} \mid r = (T,h), t \in T, E \mapsto_R t\}$. We already know $E' \in$ *adm*(SF'). Assume towards contradiction there is a set $S' \in$ *adm*(SF') such that $S' \supset E'$, i.e. there is an argument $a \in A'$ such that $a \in S' \setminus E'$. This means either $a \in A$ or $a \in A' \setminus A$. Let $S = S' \cap A$, we know $S \in$ *adm*(SF). In the first case we would have $S \supset E$, which is a contradiction to the assumption that $S \in$ *pref*(SF). In the second case we have $a \in A' \setminus A$, i.e. $a = a^1_{r,t}$ (or $a = a^2_{r,t}$, in which case the proof continues analogously) for some $r = (T, h) \in R$ and $t \in T$. Since a is attacked by t, in order to defend it we

have $S' \mapsto_{R'} t$. Since the only active attacks towards t are from within A, there must be an attack $(T', t) \in R$ such that $T' \subseteq S'$. We know $E' \not\mapsto_{R'} t$ by construction, so there is an argument $b \in A$ such that $b \in S' \setminus E'$, but since $S \in \mathit{adm}(SF)$ and $S \supset E$ again we have a contradiction to the assumption that $S \in \mathit{pref}(SF)$.
"$\Leftarrow$": Let $E' \in \mathit{pref}(SF')$ and let $E = E' \cap A$. We know $E \in \mathit{adm}(SF)$. Assume towards contradiction there is a set $S \in \mathit{adm}(SF)$ such that $S \supset E$. Let $S' = S \cup (E' \setminus E)$. By construction we have $S' \supset E'$. Moreover we have $S' \in \mathit{adm}(SF')$: assume towards contradiction there is an argument $a \in S'$ that is not defended by S', i.e. there is an attack $(T, a) \in R'$ such that $S' \not\mapsto_{R'} T$. We either have $a \in A$ or $a \in A' \setminus A$. In the first case a defends itself against attacks from $A' \setminus A$, and it is defended against attacks from A, since $a \in S$ and $S \in \mathit{adm}(SF)$. In the second case we have $a = a^1_{r,t}$ (or $a = a^2_{r,t}$, in which case the proof continues analogously) for some $r = (T, h) \in R$ and $t \in T$. We have that a defends itself against the attack from $a^2_{r,t}$. It is also attacked from t, but we have $S' \mapsto_{R'} t$: since $a \in S'$ and $a \in A' \setminus A$ by construction of S' we have $a \in E'$, but since $E' \in \mathit{adm}(SF')$ we have $E' \mapsto_{R'} t$. The argument t can only be actively attacked from within A (since there are no other active attacks towards t in R') and, hence, $S' \mapsto_{R'} t$. This shows $S' \in \mathit{adm}(SF')$, and since $S' \supset E'$ we have a contradiction to the assumption $E' \in \mathit{pref}(SF')$.

5. For $\sigma = \mathit{stage}$:
"$\Rightarrow$": Let $E \in \mathit{stage}(SF)$ and let $E' = \{a^1_{r,t} \mid r = (T, h), t \in T, t \notin E\}$. We have $E' \in \mathit{cf}(SF')$ by construction. Assume towards contradiction there is a set $S' \in \mathit{cf}(SF')$ such that $S'^{\oplus}_{R'} \supset E'^{\oplus}_{R'}$. Let $S = S' \cap A$. We know $S \in \mathit{cf}(SF)$. Moreover we have $S^{\oplus}_R \supseteq E^{\oplus}_R$ by Lemma 1. $S'^{\oplus}_{R'} \supset E'^{\oplus}_{R'}$ means there is an argument $a \in S'^{\oplus}_{R'} \setminus E'^{\oplus}_{R'}$. This means either $a \in A$ or $a \in A' \setminus A$. Since we have $A' \setminus A \subseteq S'^{\oplus}_{R'}$ by construction, the second option is impossible. So there is an argument $a \in A$ such that $a \in S'^{\oplus}_{R'} \setminus E'^{\oplus}_{R'}$, but then $a \in S^{\oplus}_R \setminus E^{\oplus}_R$, so $S^{\oplus} \supset E^{\oplus}$, which is a contradiction to our assumption $E \in \mathit{stage}(SF)$.
"$\Leftarrow$": Let $E' \in \mathit{stage}(SF')$ and let $E = E' \cap A$. We know $E \in \mathit{cf}(SF)$. Assume towards contradiction there is a set $S \in \mathit{cf}(SF)$ such that $S^{\oplus}_R \supset E^{\oplus}_R$. Let $S' = \{a^1_{r,t} \mid r = (T, h), t \in T, t \notin S\}$. We have $S' \in \mathit{cf}(SF')$ by construction. As before we have $A' \setminus A \subseteq S'^{\oplus}_{R'}$. Moreover, by Lemma 1 we have $S'^{\oplus}_{R'} \cap A \supset E'^{\oplus}_{R'} \cap A$, so we have $S'^{\oplus}_{R'} \supset E'^{\oplus}_{R'}$, which is a contradiction to the assumption $E' \in \mathit{stage}(SF')$.

6. For $\sigma = \mathit{sem}$:
"$\Rightarrow$": Let $E \in \mathit{sem}(SF)$ and let $E' = \{a^1_{r,t} \mid r = (T, h), t \in T, E \mapsto_R t\}$. We already know $E' \in \mathit{adm}(SF')$. Assume towards contradiction there is a set $S' \in \mathit{adm}(SF')$ such that $S'^{\oplus}_{R'} \supset E'^{\oplus}_{R'}$. Let $S = S' \cap A$. We know $S \in \mathit{adm}(SF)$. Moreover by Lemma 1 we have $S^{\oplus}_R \supseteq E^{\oplus}_R$. From $S'^{\oplus}_{R'} \supset E'^{\oplus}_{R'}$ we know there is an argument $a \in A'$ such that $a \in S'^{\oplus}_{R'}$ but $a \notin E'^{\oplus}_{R'}$. This means either $a \in A$ or $a \in A' \setminus A$. In the first case by Lemma 1 we get $S^{\oplus}_R \supset E^{\oplus}_R$, which is a contradiction to our assumption $E \in \mathit{sem}(SF)$. In the second case we have $a = a^1_{r,t}$ (or $a = a^2_{r,t}$, in which case the proof continues analogously) for some $r = (T, h) \in R$ and $t \in T$. We have $S' \mapsto_{R'} t$ in order to defend a. But by construction of E' we have $E' \not\mapsto_{R'} t$, hence, $E \not\mapsto_R t$, but since $S \mapsto_R t$ we have $S^{\oplus}_R \supset E^{\oplus}_R$, which is a contradiction to our assumption $E \in \mathit{sem}(SF)$.

Table 4: The complexity for primal-symmetric SETAFs.

	adm	*stb*	*pref*	*com*	*grd*	*stage*	*sem*
$Cred_\sigma$	NP-c	NP-c	NP-c	NP-c	in L	Σ_2^P-c	Σ_2^P-c
$Skept_\sigma$	trivial	coNP-c	Π_2^P-c	in L	in L	Π_2^P-c	Π_2^P-c

"$\Leftarrow$": Let $E' \in \mathit{sem}(SF')$ and let $E = E' \cap A$. We know $E \in \mathit{adm}(SF)$. Assume towards contradiction there is a set $S \in \mathit{adm}(SF)$ such that $S_R^\oplus \supset E_R^\oplus$. Let $S' = \{a_{r,t}^1 \mid r = (T,h), t \in T, S \mapsto_R t\}$. By construction we have $S' \in \mathit{adm}(SF')$. By Lemma 1 we have $S'^\oplus_{R'} \cap A \supseteq E'^\oplus_{R'} \cap A$. Moreover we have $S'^\oplus_{R'} \cap A' \setminus A \supseteq E'^\oplus_{R'} \cap A' \setminus A$: Assume otherwise, i.e. there is an argument $a \in A' \setminus A$ such that $a \in E'^\oplus_{R'} \setminus S'^\oplus_{R'}$. We have $a = a_{r,t}^1$ (or $a = a_{r,t}^2$, in which case the proof continues analogously) for some $r = (T,h) \in R$ and $t \in T$. We either have $a \in E'$ or $t \in E'$. In the first case in order to defend a we would have $E' \mapsto_{R'} t$. The argument t can only be attacked from within A, so we would also have $S \mapsto_R t$ and, hence, $S' \mapsto_{R'} t$, which means $a \in S'^\oplus_{R'}$, which is a contradiction. In the second case we have $t \in E'$, which means $t \in E_R^\oplus$, so by assumption $t \in S_R^\oplus$, and then again by construction $a \in S'^\oplus_{R'}$ (either because $t \in S'$ or because $S \mapsto_R t$).

□

In the following we show that deciding whether an argument is in the grounded extension of a primal-symmetric SETAF is doable efficiently, namely in L.

Lemma 3. *Let $SF = (A,R)$ be a primal-symmetric SETAF. Then an argument $a \in A$ is in the grounded extension G iff a is not in the head of any attack, i.e. there is no attack $(T,a) \in R$.*

Proof. Consider the construction of G with the characteristic function $\mathcal{F}_{SF}$. In the first step, exactly the arguments that are not in the head of an attack are added (which concludes the "$\Leftarrow$"-direction).
Now there are no arguments left that are defended by $\mathcal{F}_{SF}(\emptyset)$: towards contradiction assume there is an argument a that is defended by $\mathcal{F}_{SF}(\emptyset)$, but not in it. There is at least one attack $(T,a) \in R$, otherwise a would be in $\mathcal{F}_{SF}(\emptyset)$. In order to defend a we would have $\mathcal{F}_{SF}(\emptyset) \mapsto_R T$, i.e. there is an attack (T',t) with $T' \subseteq \mathcal{F}_{SF}(\emptyset)$ and $t \in T$. But since SF is primal-symmetric there is another attack $(T'',e) \in R$ such that $t \in T''$ and $e \in T'$, which is a contradiction, since the arguments in $\mathcal{F}_{SF}(\emptyset)$ are not in the head of any attack. This means $\mathcal{F}_{SF}(\emptyset) = \mathcal{F}_{SF}(\mathcal{F}_{SF}(\emptyset))$, which concludes the "$\Rightarrow$"-direction. □

To conclude the proof, the memberships of the respective problems follow from the respective results for arbitrary SETAFs. By Lemma 2 we get the respective hardness results for $Cred_\sigma$ and $Skept_\sigma$ for $\sigma \in \{\mathit{adm}, \mathit{stb}, \mathit{pref}, \mathit{stage}, \mathit{sem}\}$. The hardness of $Cred_{com}$ follows from the identity $Cred_{com} = Cred_{adm}$. Finally, the L-membership of $Cred_{grd}$, $Skept_{grd}$, and $Skept_{com}$ follows from Lemma 3. Summarizing the previous results we have the full complexity landscape for primal-symmetric SETAFs.

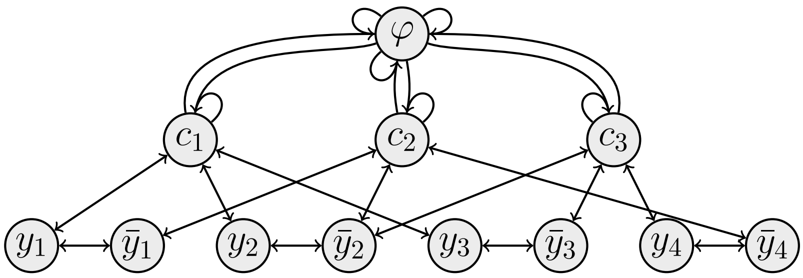

Figure 6: Illustration of SF^1_φ for a formula φ with $Y = \{y_1, y_2, y_3, y_4\}$, and $C = \{\{y_1, y_2, y_3\}, \{\bar{y}_1, \bar{y}_2, \bar{y}_4)\}, \{\bar{y}_2, \bar{y}_3, y_4\}\}$.

A.2 Proof of Theorem 5

First we will show that reasoning on the grounded extension is efficient, in particular we have that the following lemma allows us to decide our reasoning problems w.r.t. *grd* semantics in L.

Lemma 4. *Let $SF = (A, R)$ be a fully-symmetric SETAF. Then an argument $a \in A$ is in the grounded extension G of SF iff it is not involved in any attack.*

Proof. For the $\Rightarrow$-direction consider the construction of G via the characteristic function $\mathcal{F}_{SF}$. In the first step, exactly the arguments that are not involved in any attacks are added. As every other argument is now attacked (i.e. not defended by $\mathcal{F}_{SF}(\emptyset)$), a fix point is reached. The $\Leftarrow$-direction follows from the definition of the grounded extension. □

To show that reasoning w.r.t. *adm* semantics has the full complexity in fully-symmetric SETAFs, consider the following fully-symmetric variation of the standard reduction. For an illustration of the next reduction see Figure 6.

Reduction 1. *Let φ be a CNF-formula with a set of clauses C over propositional atoms Y. We define $SF^1_\varphi = (A', R')$, where*

$$\begin{aligned} A' &= \{\varphi\} \cup C \cup Y \cup \bar{Y} \\ R' &= \{(\{c, \varphi\}, \varphi), (\{c, \varphi\}, c) \mid c \in C\} \cup \{(y, \bar{y}), (\bar{y}, y) \mid y \in Y\} \cup \\ &\quad\ \{(y, c), (c, y) \mid y \in c, c \in C\} \cup \{(\bar{y}, c), (c, \bar{y}) \mid \bar{y} \in c, c \in C\} \end{aligned}$$

The only changes to the standard reduction are some additional attacks in order to make the SETAF fully-symmetric, and that the attacks between φ and arguments $c \in C$ are now self-attacks. These attacks between φ and the arguments in C are always inactive, which means φ cannot defend itself against the attacks from the arguments in C. This means φ can only be in an admissible set, if all $c \in C$ are attacked by arguments $y \in Y$ and $\bar{y} \in \bar{Y}$, which lets us construct a satisfying truth assignment (see next lemma).

Lemma 5. *Let φ be a CNF-formula with a set of clauses C over propositional atoms Y. Then φ is satisfiable iff φ is credulously accepted in SF^1_φ w.r.t. adm semantics.*

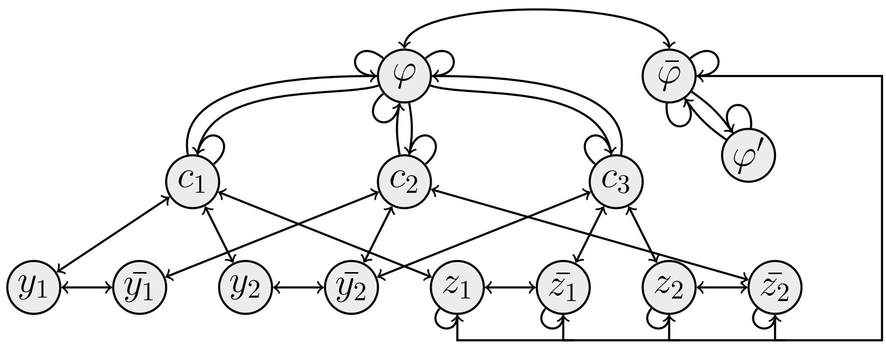

Figure 7: Illustration of SF^2_ϕ for $\phi = \forall Y \exists Z \varphi(Y, Z)$ with $Y = \{y_1, y_2\}$, $Z = \{z_1, z_2\}$, and $C = \{\{y_1, y_2, z_1\}, \{\bar{y}_1, \bar{y}_2, \bar{z}_2)\}, \{\bar{y_2}, \bar{z}_1, z_2\}\}$. Note that the attacks between $\bar{\varphi}$ and $z \in Z$ (and $\bar{z} \in \bar{Z}$ respectively) are of the form $(\{\bar{\varphi}, z\}, z)$, $(\{\bar{\varphi}, z\}, \bar{\varphi})$, $(\{\bar{\varphi}, \bar{z}\}, \bar{z})$, and $(\{\bar{\varphi}, \bar{z}\}, \bar{\varphi})$ respectively, and overlay in this illustration only in the interest of presentability.

Proof. First note that, as there can never be an active attack towards φ in SF^1_φ, no $c \in C$ can be defended against the attack from c and φ, therefore no such c can be in an admissible set.
"⇒": Assume φ is satisfiable, i.e. there is a truth assignment $\mathcal{I}$ such that $\mathcal{I} \vDash \varphi$. We can construct an admissible set in the following way: let $E = \{\varphi\} \cup \{y \mid y \in \mathcal{I}\} \cup \{\bar{y} \mid y \in Y \setminus \mathcal{I}\}$. E is conflict-free by construction. Moreover note that, as E was constructed from a satisfying assignment, each $c \in C$ is attacked, which means φ is defended against all attacks towards it, and also each $y, \bar{y} \in E$ is defended against the attacks from the arguments c. Finally, each of the arguments $y, \bar{y} \in E$ defends itself against the attack from its dual literal.
"⇐": Assume there is an admissible set $E \subseteq A'$ such that $\varphi \in E$. In order to defend φ we have $E \mapsto_{R'} c$ for all $c \in C$. We have that for each $y \in Y$, at most one of y and $\bar{y}$ is in E. Now let $\mathcal{I}$ be an interpretation such that $y \in \mathcal{I} \Leftrightarrow y \in E$ for $y \in Y$. We have $\mathcal{I} \vDash \varphi$, as for each clause c at least one of its literals attacks c in E. □

To further show that also reasoning w.r.t. *pref* semantics has the full complexity in fully-symmetric SETAFs, consider the following fully-symmetric variation of the standard reduction to show Π_2^P-hardness for $Skept_{pref}$. For an illustration of the next reduction see Figure 7.

Reduction 2. *Let $\phi = \forall Y \exists Z C$ be a $QBF^2_\forall$ formula with sets of propositional atoms Y, Z and a conjunctive formula φ over a set of clauses C. We define $SF^2_\phi = (A', R')$, where*

$$\begin{aligned} A' &= \{\varphi, \bar{\varphi}, \varphi'\} \cup C \cup Y \cup \bar{Y} \cup Z \cup \bar{Z} \\ R' &= \{(\{c, \varphi\}, \varphi), (\{c, \varphi\}, c) \mid c \in C\} \cup \{(y, \bar{y}), (\bar{y}, y) \mid y \in Y\} \cup \\ &\quad \{(y, c), (c, y) \mid y \in c, c \in C\} \cup \{(\bar{y}, c), (c, \bar{y}) \mid \bar{y} \in c, c \in C\} \cup \\ &\quad \{(\varphi, \bar{\varphi}), (\bar{\varphi}, \varphi), (\{\bar{\varphi}, \varphi'\}, \bar{\varphi}), (\{\bar{\varphi}, \varphi'\}, \varphi')\} \cup \\ &\quad \{(\{z, \bar{\varphi}\}, z), (\{z, \bar{\varphi}\}, \bar{\varphi}), (\{\bar{z}, \bar{\varphi}\}, \bar{z}), (\{\bar{z}, \bar{\varphi}\}, \bar{\varphi}) \mid z \in Z\} \end{aligned}$$

Similar to Reduction 1, in Reduction 2 we add additional attacks in order to make the SETAF fully-symmetric. As in the standard reduction for the Π_2^P-hardness of $Skept_{pref}$ (cf. [10]) we have that a $QBF_{\forall}^{2}$ formula ϕ is true iff the argument φ is skeptically accepted w.r.t. *pref* semantics in SF_{ϕ}^{2} (see next lemma). The attacks between the argument $\bar{\varphi}$ and the arguments z (or $\bar{z}$ respectively) are also self-attacks, because otherwise the arguments z (or $\bar{z}$ respectively) would defend themselves against the attacks from $\bar{\varphi}$, while only φ should (actively) attack $\bar{\varphi}$. Moreover, we do not want $\bar{\varphi}$ in an admissible set, but a self-loop $(\bar{\varphi}, \bar{\varphi})$ would make SF_{ϕ}^{2} redundant. To ensure $\bar{\varphi}$ is in no admissible set we introduce φ', which would have to be attacked in order to defend $\bar{\varphi}$, but this is impossible. We have that φ' is in a preferred extension iff φ is in the extension.

Lemma 6. *Let $\phi = \forall Y \exists Z C$ be a $QBF_{\forall}^{2}$ formula with sets of propositional atoms Y and Z and a conjunctive formula φ over a set of clauses C. Then ϕ is true iff φ is skeptically accepted in SF_{ϕ}^{2} w.r.t. pref semantics.*

Proof. First note that the argument $\bar{\varphi}$ cannot be in an admissible set, as it is attacked by φ' and the only attack towards φ' is an inactive attacks. As with the previous reduction we have that no argument $c \in C$ can be in an admissible set: the only active attack towards φ is from $\bar{\varphi}$, which is in no admissible set.
"$\Rightarrow$": Assume ϕ is true, i.e. for every partial interpretation $\mathcal{I}_Y \subseteq Y$ there is a partial interpretation $\mathcal{I}_Z \subseteq Z$ such that $\mathcal{I}_Y \cup \mathcal{I}_Z \models \varphi$. Note that each $\mathcal{I}_Y$ corresponds to an admissible set $S = \{y \mid y \in \mathcal{I}_Y\} \cup \{\bar{y} \mid y \in Y \setminus \mathcal{I}_Y\}$. Every admissible set $E \in adm(SF_{\phi}^{2})$ that has $z \in E$ for some $z \in Z$ has to have $\varphi \in E$ in order to defend z against the attack from $\bar{\varphi}$. Now φ is in E iff the arguments from $Y \cup \bar{Y} \cup Z \cup \bar{Z}$ attack all arguments $c \in C$, i.e. if the corresponding interpretation $\mathcal{I} = (E \cap Y) \cup (E \cap Z)$ makes φ true. Since we assumed ϕ to be true, we know that for each partial assignment $\mathcal{I}_Y$ (and, hence, for each admissible set) there is such a partial assignment $\mathcal{I}_Z$, therefore φ is in every preferred extension of SF_{ϕ}^{2}.
"$\Leftarrow$": Assume φ is in every preferred extension of SF_{ϕ}^{2}. As we know each partial assignment $\mathcal{I}_Y \subseteq Y$ corresponds to an admissible set S in SF_{ϕ}^{2} and for each admissible set S there is an extension $E \in pref(SF_{\phi}^{2})$ such that $E \supseteq S$, and since we know φ can only be in an admissible set E is the corresponding interpretation $\mathcal{I} = (E \cap Y) \cup (E \cap Z)$ makes φ true, we get that for each such partial assignment $\mathcal{I}_Y$ there is an assignment $\mathcal{I}_Z \subseteq Z$ such that $\mathcal{I}_Y \cup \mathcal{I}_Z \models \varphi$, i.e. ϕ is true. □

Already we have all information to pinpoint the complexity of reasoning in fully-symmetric SETAFs. The membership for $Cred_{\sigma}$ for $\sigma \in \{cf, adm, stb, pref, com, stage, sem\}$ follows from the general case, likewise the membership for $Skept_{\sigma}$ for $\sigma \in \{stb, pref, stage, sem\}$ follows from the general case. The L-membership for $Cred_{grd}$ follows from Lemma 4, from the identity $Cred_{grd} = Skept_{grd} = Skept_{com}$ we get the respective membership proofs for $Skept_{grd}$ and $Skept_{com}$. The problems $Cred_{\sigma}$ and $Skept_{\sigma}$ for $\sigma \in \{stb, stage, sem\}$ already have their full hardness for symmetric AFs allowing self-attacks (see [10]). The NP-hardness of $Cred_{adm}$ follows from Lemma 5, then the hardness of $Cred_{com}$ and $Cred_{pref}$ immediately follow by the identity $Cred_{adm} = Cred_{com} = Cred_{pref}$. Finally, the Π_2^{P}-hardness of $Skept_{pref}$ follows from Lemma 6.

Table 5: The complexity for redundancy-free self-attack-free fully-symmetric SETAFs.

	adm	*stb*	*pref*	*com*	*grd*	*stage*	*sem*
$Cred_\sigma$	trivial	trivial	trivial	trivial	in L	trivial	trivial
$Skept_\sigma$	trivial	in L	in L	in L	in L	in L	in L

A.3 Proof of Theorem 6

Lemma 7. *Let $SF = (A, R)$ be a self-attack-free, fully-symmetric SETAF. Then we have naive(SF) = stb(SF).*

Proof. We know that every stable extension is naive, it remains to show that for self-attack-free, fully-symmetric SETAFs every naive extension is stable. Towards contradiction assume there is a naive extension $E \in naive(SF)$ such that there is an argument $a \in A \setminus E^{\oplus}_R$. Since E is a naive extension, we have that $E \cup \{a\}$ is not conflict-free, i.e. there is an attack (T, b) with $T \cup \{b\} \subseteq E \cup \{a\}$. As SF is fully symmetric we then also have an attack $((T \cup \{b\}) \setminus \{a\}, a)$. But then we have that $a \in E^{\oplus}_R$, which is a contradiction. □

This suffices to establish the complexity of reasoning in self-attack-free fully-symmetric SETAFs for the semantics under our consideration: since there are no self-attacks, for every argument a the set $\{a\}$ is conflict-free, which means $Cred_{cf}$ is trivially true. Moreover, since for every conflict-free set S there is a naive extension E with $E \subseteq S$ this carries over to $Cred_{naive}$. As by Lemma 7 we have that *naive*(SF) = *stb*(SF), we know that also *stb*(SF) = *stage*(SF) = *sem*(SF) for any self-attack-free fully-symmetric SETAF SF, since there is always at least one naive extension. Likewise we get $Cred_{stb} = Cred_{stage} = Cred_{sem}$. As every stable extension is admissible, preferred, and complete, this also carries over to $Cred_{adm}$, $Cred_{pref}$, and $Cred_{com}$.

Now by Lemma 4 we get that it suffices to check whether an argument is involved in any attack to know if it is in the grounded extension, hence, the problems $Cred_{grd}$, $Skept_{grd}$, and $Skept_{com}$ are in L.

Now note that since for every attack (T, h) the set T is conflict-free (as the SETAF is redundancy-free there cannot be an attack within T), we can construct a naive extension E such that an arbitrary argument a that is involved in at least one attack is not in E. Hence, to decide $Skept_\sigma$ for $\sigma \in \{naive, stb, pref, stage, sem\}$ it also suffices to check whether an argument is involved in any attack, which can be done in L. The results are summarized in Table 5.

A.4 Proof of Theorem 7

The proof of this theorem mainly concerns the correctness and completeness of Algorithm 1. Let $SF = (A, R)$ be a SETAF with a partitioning (Y, Z). As in the algorithm for AFs [6], our adaptation iteratively removes arguments that cannot be defended. This algorithm has to be executed for both the set Y and the set Z to get all credulously accepted arguments of a SETAF SF. Assume we start with Y. In step 6 of the i-th iteration of the of the algorithm we remove every argument y that is attacked via an attack (Z', y) (as SF is primal-bipartite Z' must be a subset of Z) such that there are no defenders against the attack left, i.e. no $z \in Z'$ is attacked by a subset of the arguments left in Y_{i-1}. In step 7 we remove all attacks that origin from already removed arguments; they cannot be part of a defending attack. More formally, the correctness and completeness of Algorithm 1 is shown in the following.

Lemma 8 (cf.[6]). *Let $SF = (A, R)$ be a primal-bipartite SETAF with a partitioning (Y, Z), then an argument $a \in Y$ is credulously accepted w.r.t. pref semantics iff it is in the set returned by Algorithm 1.*
Moreover the set returned by Algorithm 1 is admissible in SF.

Proof. "⇒": We will show inductively that for every iteration of the algorithm the arguments that are removed in step 6 are not defensible and the attacks that are removed in step 7 cannot be part of a defending attack. For the first iteration this is the case, as we construct Y_1 by only removing those arguments $y \in Y$ from Y that are attacked by an attack (Z', y) on which no counter-attack exists. Moreover we remove all attacks (Y', z) towards arguments $z \in Z$ such that for one of the arguments $y' \in Y'$ we already showed it is not defensible, as they cannot defend any argument in an admissible set. Likewise, assuming this property holds for the $i-1$-th iteration, in the i-th iteration we only remove arguments that are not defensible and attacks that cannot play a role in admissible sets.
Assume towards contradiction an argument $y \in Y$ is credulously accepted, but not in the set S that is returned by the algorithm. This means at some iteration i the argument y is removed, but, as established, this means it is not defensible, which is a contradiction to the assumption is it credulously accepted.
"⇐": Let S be the set that is returned by the algorithm. Assume we have $x \in S$ for some argument $x \in Y$. As we have $S \subseteq Y$, we know S is conflict-free in SF. Moreover we know that S defends x: towards contradiction assume otherwise, i.e. there is an attack (Z', x) towards x such that S does not attack Z'. But then x would be removed in step 6, which is a contradiction to the assumption that $x \in S$. □

Algorithm 1 runs in polynomial time: there can be at most $|Y|$ iterations; step 6 is efficient, as all involved sets are bounded by the number of attacks and the number of arguments involved in an attack; step 7 is also efficient, as it suffices to check for every attack towards arguments $z \in Z$.

Note that by symmetry this algorithm also works for arguments $z \in Z$ such that it is sufficient to compute all credulously accepted arguments of a primal-bipartite SETAF SF. Hence, $Cred_{pref}$ is P-easy for this subclass. We will now show that this result carries over to other semantics under our consideration. We know that primal-bipartite SETAFs have no odd-cycles, and therefore are coherent, which implies *pref*(SF) = *stb*(SF). Note that as there always is at least one preferred extension there also always is a stable extension, which further implies *stb*(SF) = *sem*(SF) = *stage*(SF). The following lemma also holds for AFs (see [24]).

Lemma 9. *Let $SF = (A, R)$ be a SETAF with pref(SF) = stb(SF). Then an argument $a \in A$ is skeptically accepted w.r.t. pref semantics iff for every attack $(T, a) \in R$ towards a we have that $T \not\subseteq E$ for every preferred extension $E \in$ pref(SF).*

Proof. "⇒": Assume an argument $a \in A$ is in every preferred extension of SF and let (T, a) be an arbitrary attack towards a. Then T cannot be a subset of any preferred extension E, as we would have $T \cup \{a\} \subseteq E$, which is not conflict-free.
"⇐": Assume in every extension $E \in \mathit{pref}(SF)$ we have for every attack (T, a) towards a that $T \not\subseteq E$. This means for every attack there is some $t \in T$ such that $t \notin E$. But as E is stable by assumption, this means t is attacked, and, hence, a is defended against all attacks. □

By a result of [6] we know that even for bipartite AFs $F = (A, R)$ with a partitioning (Y, Z) it is NP-complete to decide for sets $S \subseteq A$ if the arguments are jointly credulously accepted w.r.t. *pref* semantics. This hardness-result carries over to SETAFs. However, if we restrict the problem to deciding whether a set $S \subseteq Y$ is jointly credulously accepted, this problem becomes P-easy even for SETAFs, as this is the case iff every single argument $a \in S$ is credulously accepted, which we established can be decided in polynomial time with Algorithm 1.

Lemma 10. *Let $SF = (A, R)$ be a primal-bipartite SETAF with a partitioning (Y, Z). Then for any set $Y' \subseteq Y$ there is a preferred extension $E \supseteq Y'$ iff every argument $y' \in Y'$ is credulously accepted w.r.t. pref semantics.*

Proof. By Lemma 8 we have that an argument $a \in Y$ is credulously accepted w.r.t. *pref* semantics iff it is in the set S returned by Algorithm 1, and we have $S \in \mathit{adm}(SF)$. This means that all credulously accepted arguments $a \in Y$ are also jointly credulously accepted in SF, which also means that every subset $Y' \subseteq Y$ that consists only of credulously accepted arguments is jointly credulously accepted w.r.t. *adm*, which in turn means they are jointly credulously accepted w.r.t. *pref* semantics, as every admissible set is part of a subset-maximal admissible set. □

Again, by symmetry, this result also applies for sets $Z' \subseteq Z$. The respective hardness proofs follow from the hardness of bipartite AFs. As already established, Algorithm 1 can be used to efficiently (namely, in polynomial time) compute the set of credulously accepted arguments w.r.t. *pref* semantics, which carries over to *com* and *adm*, and as primal-bipartite SETAFs are coherent and preferred and stable semantics coincide, also to *stb*, *stage*, and *sem*. For the P-membership of $Skept_\sigma$ for $\sigma \in \{\mathit{stb}, \mathit{pref}, \mathit{stage}, \mathit{sem}\}$ we use the same identity $\mathit{stb}(SF) = \mathit{pref}(SF) = \mathit{stage}(SF) = \mathit{sem}(SF)$, and note that to check if an argument $a \in A$ is skeptically accepted w.r.t. *pref* semantics by Lemma 9 we know that it suffices to check if for every attack (T, a) towards a the set T is jointly credulously accepted, which, by Lemma 10, can be done in polynomial time, as it suffices to check if every argument in T is credulously accepted. Finally, the L membership of $Skept_{\mathit{naive}}$ follows from the general case, and the trivial results for $Cred_{\mathit{cf}}$ and $Cred_{\mathit{naive}}$ follow from the fact that every primal-bipartite SETAF has no self-loops (i.e. every argument $a \in A$ is in the conflict-free set $\{a\}$, and, hence, in a naive extension). The results are summarized in Table 6

A.5 Proof of Theorem 8

We will show that we can translate every SETAF SF into a 2-colorable SETAF $Tr(SF)$ with an acceptance-preserving translation Tr. This holds for the semantics $\sigma \in \{\mathit{stb}, \mathit{pref}, \mathit{com}, \mathit{grd}, \mathit{sem}\}$.

Table 6: The complexity of primal-bipartite SETAFs.

	adm	*stb*	*pref*	*com*	*grd*	*stage*	*sem*
$Cred_\sigma$	P-c	P-c	P-c	P-c	P-c	P-c	P-c
$Skept_\sigma$	trivial	P-c	P-c	P-c	P-c	P-c	P-c

To this end we use a translation where we add two fresh arguments as another attacker to the tail of every attack. This is captured by translation Tr_2.

Translation 2. *Let $SF = (A, R)$ be a SETAF. The SETAF translation Tr_2 is defined as $Tr_2(SF) = (A', R')$ with*

$$\begin{aligned} A' &= A \cup \{a_a^*, a_b^*\}, \\ R' &= \{(T \cup \{a_a^*, a_b^*\}, h) \mid (T, h) \in R\} \end{aligned}$$

This translation is efficient. It remains to show that Tr_2 is acceptance-preserving.

Lemma 11. *The SETAF-translation Tr_2 is acceptance preserving for $\sigma \Rightarrow \sigma$ with $\sigma \in \{stb, pref, com, grd, sem\}$ such that $Tr_2(SF)$ is 2-colorable for every SETAF SF.*

Proof. We have that the fresh arguments a_a^*, a_b^* are not attacked and, hence, in the grounded extension, and as every σ-extension contains the grounded extension, they are skeptically accepted w.r.t. σ. Moreover, for any set $S \subseteq A$ it holds that $S \in \sigma(SF)$ iff $S' = (S \cup \{a_a^*, a_b^*\}) \in \sigma(SF')$, which can easily be seen for each of the semantics in question. Furthermore, as a_a^* and a_b^* are part of every attack, every partitioning (Y, Z) with $a_a^* \in Y$ and $a_b^* \in Z$ is a valid 2-coloring. □

Note that this translation does not work as a reduction for semantics based on conflict-free sets, as for every attack (T, h) in the translation the set $T \cup \{h\}$ is conflict-free. To show the hardness of our reasoning tasks for *stage* semantics we introduce another reduction from the Π_2^{P}-hard $QBF_\forall^2$ problem, such that the constructed SETAF is always 2-colorable. For an illustration of SF_3^Φ see Figure 8.

Reduction 3. *Let $\Phi = \forall Y \exists Z C$ be a $QBF_\forall^2$-formula with at least 2 clauses where in each clause at least one positive and at least one negative literal occurs, consisting of a set of clauses C over sets of propositional atoms Y and Z. We define the SETAF $SF_3^\Phi = (A, R)$, where*

$$\begin{aligned} A = {} & \{\varphi, \bar{\varphi}', \bar{\varphi}, \varphi', \varphi'', \varphi'''\} \cup C \cup Y \cup \bar{Y} \cup Z \cup \bar{Z} \ \cup \\ & \{y', y'', y''', \bar{y}', \bar{y}'', \bar{y}''' \mid y \in Y\}, \\ R = {} & \{(x, \bar{x}), (\bar{x}, x) \mid x \in Y \cup Z\} \cup \{(\{x \mid \bar{x} \in c\} \cup \{\bar{x} \mid x \in c\}, c) \mid c \in C\} \ \cup \\ & \{(\{c \mid c \in C\}, \bar{\varphi}'), (\bar{\varphi}', \varphi), (\bar{\varphi}, \varphi), (\varphi, \bar{\varphi})\} \ \cup \\ & \{(\{\varphi, \varphi'\}, \varphi''), (\{\varphi, \varphi'\}, \varphi'''), (\{\varphi'', \varphi'''\}, \varphi''), (\{\varphi'', \varphi'''\}, \varphi''')\} \ \cup \\ & \{(\{y, y'\}, y''), (\{y, y'\}, y'''), (\{y'', y'''\}, y''), (\{y'', y'''\}, y''') \mid y \in Y\} \ \cup \\ & \{(\{\bar{y}, \bar{y}'\}, \bar{y}''), (\{\bar{y}, \bar{y}'\}, \bar{y}'''), (\{\bar{y}'', \bar{y}'''\}, \bar{y}''), (\{\bar{y}'', \bar{y}'''\}, \bar{y}''') \mid \bar{y} \in \bar{Y}\} \end{aligned}$$

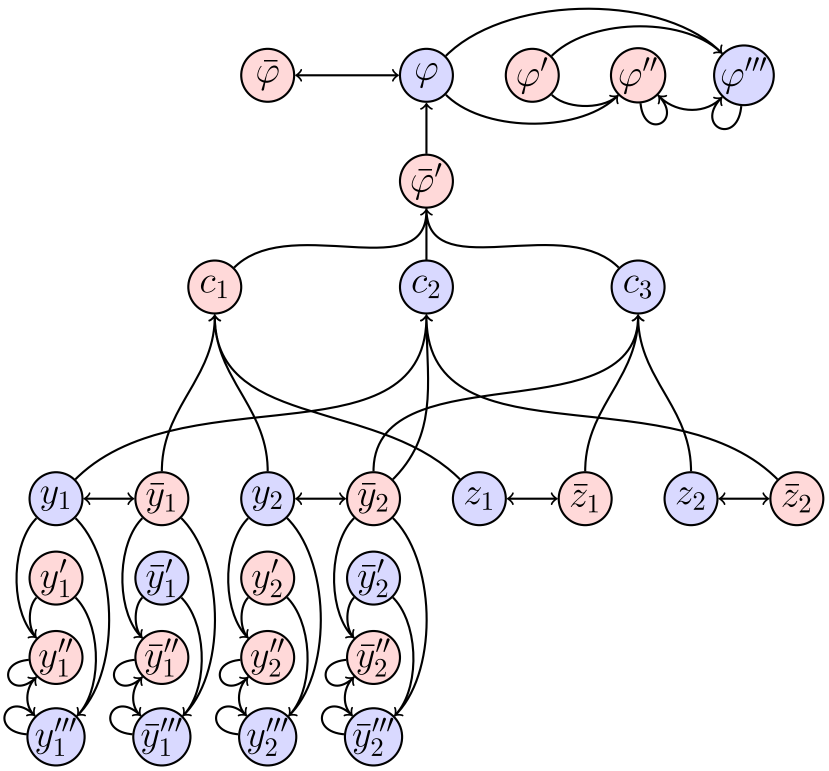

Figure 8: Illustration of SF_3^Φ for $\Phi = \forall Y \exists Z \varphi(Y, Z)$ with $Y = \{y_1, y_2\}$, $Z = \{z_1, z_2\}$, and $\varphi = \{\{y_1, \bar{y}_2, \bar{z}_1\}, \{\bar{y}_1, y_2, z_2)\}, \{y_2, z_1, \bar{z}_2\}\}$. The coloring of the arguments corresponds to a possible partitioning that shows the 2-colorability of SF_3^Φ, i.e. we have that no attack is monochromatic.

We have (as we will show in Lemma 13) that arguments y' and $\bar{y}'$ are in every *stage* extension, and the arguments y'' and y''' (or $\bar{y}''$ and $\bar{y}'''$ respectively) cannot be in a conflict-free set together, so the only way to have both in the range of a *stage* extension is to have y (or $\bar{y}$ respectively) in this extension. This way every combination of arguments from Y and $\bar{Y}$ (that correspond to a partial interpretation over variables Y) is in an incomparable *stage* extension.

It is not immediate why SF_3^Φ is always 2-colorable; for this we need to have for each clause $c \in C$ to have at least one positive and at least one negative literal, as otherwise this partitioning could produce a monochromatic edge (i.e. an edge such that all involved arguments are in just one of Y or Z). Moreover we assume there are at least two clauses; these two constraints do not affect the hardness of the $QBF_\forall^2$ problem. Consider a partitioning (A, B) where $A = (\{c_x, \bar{\varphi}, \bar{\varphi}', \varphi', \varphi''\} \cup \{\bar{y}, y', y'', \bar{y}'' \mid y \in Y\} \cup \{\bar{z} \mid z \in Z\})$ and $B = (\{c \mid c \in C \setminus \{c_x\}\} \cup \{\varphi, \varphi'''\} \cup \{y, \bar{y}', \bar{y}'', y'' \mid y \in Y\} \cup \{z \mid z \in Z\} \cup \{\varphi'''\})$, where c_x is an arbitrary clause. Then one can check that (A, B) is a partitioning such that SF_3^Φ is 2-colorable (the coloring in Figure 8 corresponds to such a partitioning).

The following proof follows the structure of [9].

Lemma 12. *Let Φ be a $QBF_\forall^2$ formula and let $SF_3^\Phi = (A, R)$, then for every extension $E \in$*

stage(SF_3^Φ) *we have* $\{\varphi'', \varphi'''\} \not\subseteq E$, $\{y'', y'''\} \not\subseteq E$, *and* $\{\bar{y}'', \bar{y}'''\} \not\subseteq E$ *for each* $y \in Y$. *Moreover we have* $x \in E$ *iff* $\bar{x} \notin E$ *for each* $x \in Y \cup Z \cup \{\varphi\}$.

Proof. The first statement immediately follows from the fact that E is conflict-free. Moreover we have that at at least one of x and $\bar{x}$ is in E: towards contradiction assume otherwise, i.e. $\{x, \bar{x}\} \cap E = \emptyset$. If $x = \varphi$, then $E' = E \cup \{\bar{\varphi}\}$ is conflict-free with $E'^{\oplus}_R \supset E^{\oplus}_R$. If $x \in Y \cup Z$, then $E' = (E \setminus \{c \mid c \in C, \text{there is some } (T, c) \in R \text{ such that } T \subseteq E \cup \{x\}\} \cup \{x\}$ is conflict-free with $E'^{\oplus}_R \supset E^{\oplus}_R$. By conflict-freeness we also have that at most one of x and $\bar{x}$ is in E. □

Lemma 13. *Let* Φ *be a* $QBF_\forall^2$ *formula and let* $SF_3^\Phi = (A, R)$, *then* $\{x' \mid x \in Y \cup \bar{Y} \cup \{\varphi\}\} \subseteq E$ *for every* $E \in$ *stage*(SF_3^Φ).

Proof. Towards contradiction assume $E \in$ *stage*(SF_3^Φ) and $x' \notin E$ for some $x \in Y \cup \bar{Y} \cup \{\varphi\}$, then we have $E' = (E \cup \{x'\}) \setminus \{x'', x'''\} \in$ *cf*(SF_3^Φ) with $E'^{\oplus}_R \supset E^{\oplus}_R$, which is a contradiction to the assumption $E \in$ *stage*(SF_3^Φ). □

Lemma 14. *Let* Φ *be a* $QBF_\forall^2$ *formula and let* $SF_3^\Phi = (A, R)$, *then* φ *is in every stage extension iff* Φ *is true.*

Proof. "⇒": Assume Φ is false, we show that then there is an extension $E \in$ *stage*(SF_3^Φ) such that $\varphi \notin E$. As Φ is false, there is a partial interpretation I_Y such that for each partial interpretation I_Z we have that at least one clause is not true, i.e. in the corresponding set of arguments at least one argument $c \in C$ is attacked. As by Lemma 12 and since $\bar{\varphi}'$ is not attacked, the only way to have $\{y'', y''' \mid y \in I_Y\} \cup \{\bar{\varphi}'\} \subseteq E^{\oplus}_R$ is if we also have $\bar{\varphi}' \in E$, we know that such a stage extension E with $\bar{\varphi}' \in E$ exists, but this extension can only have $\varphi \notin E$.
"⇐": Assume Φ is true, and let, towards contradiction, $E \in$ *stage*(SF_3^Φ) with $\varphi \notin E$. We know that for each partial interpretation I_Y there is a partial interpretation I_Z such that $I_Y \cup I_Z$ makes φ true. Let $I_Y = E \cap Y$ and let I_Z be such a partial interpretation such that $I_Y \cup I_Z$ makes φ true. Moreover let $E' = I_Y \cup I_Z \cup \{\bar{x} \mid x \in (Y \cup Z) \setminus (I_Y \cup I_Z)\} \cup C \cup (E \cap (Y' \cup Y'' \cup Y''' \cup \bar{Y}' \cup \bar{Y}'' \cup \bar{Y}''')) \cup \{\varphi, \varphi'\}$. One can check that E' is conflict-free in SF_3^Φ, also we have $E'^{\oplus}_R \supset E^{\oplus}_R$: by construction the ranges of E' and E coincide on all arguments but arguments $c \in C$ and on the arguments φ'' and φ''', where we have $C \subseteq E'^{\oplus}_R$ and $\{\varphi'', \varphi'''\} \subseteq E'^{\oplus}_R$, but $\{\varphi'', \varphi'''\} \not\subseteq E^{\oplus}_R$. This is a contradiction to the assumption $E \in$ *stage*(SF_3^Φ). □

These results give us the complexity landscape for 2-colorable SETAFs: they have the full complexity, i.e. 2-colorability does not allow us to reason more efficiently. The membership follows from the general case. We obtain the hardness for $Cred_\sigma$ and for $Skept_\sigma$ with $\sigma \in \{stb, pref, com, grd, sem\}$ by Lemma 11. The hardness of $Cred_{adm}$ follows from the identity $Cred_{adm} = Cred_{pref}$. The hardness of $Cred_{stage}$ follows from Lemma 14, likewise the hardness of $Skept_{stage}$ follows from Lemma 14 and the fact that by Lemma 12 we then have $\bar{\varphi}$ is in every extension $E \in$ *stage*(SF_3^Φ) iff Φ is false.

Table 7: The complexity of 2-colorable SETAFs.

	adm	*stb*	*pref*	*com*	*grd*	*stage*	*sem*
$Cred_\sigma$	NP-c	NP-c	NP-c	NP-c	P-c	Σ_2^P-c	Σ_2^P-c
$Skept_\sigma$	trivial	coNP-c	Π_2^P-c	P-c	P-c	Π_2^P-c	Π_2^P-c

References

[1] Brewka, G., Polberg, S., Woltran, S.: Generalizations of Dung frameworks and their role in formal argumentation. IEEE Intelligent Systems **29**(1), 30–38 (2014). https://doi.org/10.1109/MIS.2013.122

[2] Coste-Marquis, S., Devred, C., Marquis, P.: Symmetric argumentation frameworks. In: Godo, L. (ed.) Proceedings of the 8th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU 2005). Lecture Notes in Computer Science, vol. 3571, pp. 317–328. Springer (2005). https://doi.org/10.1007/11518655_28

[3] Diller, M., Keshavarzi Zafarghandi, A., Linsbichler, T., Woltran, S.: Investigating subclasses of abstract dialectical frameworks. Argument & Computation **11**, 191–219 (2020). https://doi.org/10.3233/AAC-190481

[4] Dimopoulos, Y., Torres, A.: Graph theoretical structures in logic programs and default theories. Theor. Comput. Sci. **170**(1-2), 209–244 (1996). https://doi.org/10.1016/S0304-3975(96)80707-9

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

[6] Dunne, P.E.: Computational properties of argument systems satisfying graph-theoretic constraints. Artif. Intell. **171**(10-15), 701–729 (2007). https://doi.org/10.1016/j.artint.2007.03.006

[7] Dunne, P.E., Bench-Capon, T.J.M.: Complexity and combinatorial properties of argument systems. Technical report, Dept. of Computer Science, University of Liverpool (2001)

[8] Dunne, P.E., Bench-Capon, T.J.M.: Coherence in finite argument systems. Artif. Intell. **141**(1/2), 187–203 (2002). https://doi.org/10.1016/S0004-3702(02)00261-8

[9] Dvořák, W.: Computational Aspects of Abstract Argumentation. Ph.D. thesis, Vienna University of Technology, Institute of Information Systems (2012), `http://permalink.obvsg.at/AC07812708`

[10] Dvořák, W., Dunne, P.E.: Computational problems in formal argumentation and their complexity. FLAP **4**(8) (2017), http://www.collegepublications.co.uk/downloads/ifcolog00017.pdf

[11] Dvořák, W., Fandinno, J., Woltran, S.: On the expressive power of collective attacks. Argument Comput. **10**(2), 191–230 (2019). https://doi.org/10.3233/AAC-190457

[12] Dvořák, W., Greßler, A., Woltran, S.: Evaluating SETAFs via answer-set programming. In: Thimm, M., Cerutti, F., Vallati, M. (eds.) Proceedings of the Second International Workshop on Systems and Algorithms for Formal Argumentation (SAFA 2018) co-located with the 7th International Conference on Computational Models of Argument (COMMA 2018), Warsaw, Poland, September 11, 2018. CEUR Workshop Proceedings, vol. 2171, pp. 10–21. CEUR-WS.org (2018), http://ceur-ws.org/Vol-2171/paper_2.pdf

[13] Dvořák, W., Ordyniak, S., Szeider, S.: Augmenting tractable fragments of abstract argumentation. Artificial Intelligence **186**(0), 157–173 (2012). https://doi.org/10.1016/j.artint.2012.03.002

[14] Dvořák, W., Pichler, R., Woltran, S.: Towards fixed-parameter tractable algorithms for abstract argumentation. Artif. Intell. **186**, 1 – 37 (2012). https://doi.org/10.1016/j.artint.2012.03.005

[15] Dvořák, W., Rapberger, A., Wallner, J.P.: Labelling-based algorithms for SETAFs. In: Gaggl, S.A., Thimm, M., Vallati, M. (eds.) Proceedings of the Third International Workshop on Systems and Algorithms for Formal Argumentation co-located with the 8th International Conference on Computational Models of Argument (COMMA 2020), September 8, 2020. CEUR Workshop Proceedings, vol. 2672, pp. 34–46. CEUR-WS.org (2020), http://ceur-ws.org/Vol-2672/paper_4.pdf

[16] Dvořák, W., Rapberger, A., Woltran, S.: On the different types of collective attacks in abstract argumentation: Equivalence results for SETAFs. Journal of Logic and Computation **30**(5), 1063–1107 (2020). https://doi.org/10.1093/logcom/exaa033

[17] Dvořák, W., Woltran, S.: Complexity of semi-stable and stage semantics in argumentation frameworks. Inf. Process. Lett. **110**(11), 425–430 (2010). https://doi.org/10.1016/j.ipl.2010.04.005

[18] Dvořák, W., Woltran, S.: On the intertranslatability of argumentation semantics. J. Artif. Intell. Res. (JAIR) **41**, 445–475 (2011)

[19] Flouris, G., Bikakis, A.: A comprehensive study of argumentation frameworks with sets of attacking arguments. Int. J. Approx. Reason. **109**, 55–86 (2019). https://doi.org/10.1016/j.ijar.2019.03.006

[20] König, M.: Graph-Classes of Argumentation Frameworks with Collective Attacks. Master's thesis, TU Wien (2020), http://permalink.obvsg.at/AC15750327

[21] Nielsen, S.H., Parsons, S.: Computing preferred extensions for argumentation systems with sets of attacking arguments. In: Dunne, P.E., Bench-Capon, T.J.M. (eds.) Computational Models of Argument: Proceedings of COMMA 2006, September 11-12, 2006, Liverpool, UK. Frontiers in Artificial Intelligence and Applications, vol. 144, pp. 97–108. IOS Press (2006), http://www.booksonline.iospress.nl/Content/View.aspx?piid=1930

[22] Nielsen, S.H., Parsons, S.: A generalization of Dung's abstract framework for argumentation: Arguing with sets of attacking arguments. In: Maudet, N., Parsons, S., Rahwan, I. (eds.) Argumentation in Multi-Agent Systems, Third International Workshop, ArgMAS 2006, Hakodate, Japan, May 8, 2006, Revised Selected and Invited Papers. Lecture Notes in Computer Science, vol. 4766, pp. 54–73. Springer (2006). https://doi.org/10.1007/978-3-540-75526-5_4

[23] Polberg, S.: Developing the Abstract Dialectical Framework. Ph.D. thesis, Vienna University of Technology, Institute of Information Systems (2017), https://permalink.obvsg.at/AC13773888

[24] Vreeswik, G.A.W., Prakken, H.: Credulous and sceptical argument games for preferred semantics. In: Ojeda-Aciego, M., de Guzmán, I.P., Brewka, G., Moniz Pereira, L. (eds.) Proceedings of Logics in Artificial Intelligence. pp. 239–253. Springer Berlin Heidelberg, Berlin, Heidelberg (2000)

[25] Yun, B., Vesic, S., Croitoru, M.: Toward a more efficient generation of structured argumentation graphs. In: Modgil, S., Budzynska, K., Lawrence, J. (eds.) Computational Models of Argument - Proceedings of COMMA 2018, Warsaw, Poland, 12-14 September 2018. Frontiers in Artificial Intelligence and Applications, vol. 305, pp. 205–212. IOS Press (2018). https://doi.org/10.3233/978-1-61499-906-5-205