Journal of Artificial Intelligence Research 79 (2024) 69-136 Submitted 04/2023; published 01/2024

Principles and their Computational Consequences for Argumentation Frameworks with Collective Attacks

Wolfgang Dvořák WOLFGANG.DVORAK@TUWIEN.AC.AT
TU Wien, Institute of Logic and Computation

Matthias König MATTHIAS.KOENIG@TUWIEN.AC.AT
TU Wien, Institute of Logic and Computation

Markus Ulbricht MULBRICHT@INFORMATIK.UNI-LEIPZIG.DE
Leipzig University, Department of Computer Science, ScaDS.AI

Stefan Woltran STEFAN.WOLTRAN@TUWIEN.AC.AT
TU Wien, Institute of Logic and Computation

Abstract

Argumentation frameworks (AFs) are a key formalism in AI research. Their semantics have been investigated in terms of *principles*, which define characteristic properties in order to deliver guidance for analyzing established and developing new semantics. Because of the simple structure of AFs, many desired properties hold almost trivially, at the same time hiding interesting concepts behind syntactic notions. We extend the principle-based approach to argumentation frameworks with collective attacks (SETAFs) and provide a comprehensive overview of common principles for their semantics. Our analysis shows that investigating principles based on decomposing the given SETAF (e.g. directionality or SCC-recursiveness) poses additional challenges in comparison to usual AFs. We introduce the notion of the reduct as well as the modularization principle for SETAFs which will prove beneficial for this kind of investigation. We then demonstrate how our findings can be utilized for incremental computation of extensions and show how we can use graph properties of the frameworks to speed up these algorithms.

1. Introduction

In the field of knowledge representation & reasoning there is a variety of formalisms to capture argumentation processes and discussions. Many of these notions are based on abstract argumentation frameworks (AFs) as proposed by Dung (1995), where arguments are interpreted as abstract entities. Consequently, the focus is solely on the relationship between the arguments, i.e. which arguments are in conflict with each other. To classify and distinguish the various semantics which are used to define consistent sets of arguments (so called extensions) the *principle-based analysis* is an established method in formal argumentation research. Principle-based investigations have recently been performed e.g. for AFs (van der Torre & Vesic, 2017), ranking semantics (Bonzon, Delobelle, Konieczny, & Maudet, 2017), preference-based argumentation frameworks (Kaci, van der Torre, & Villata, 2018), quantitative bipolar argumentation frameworks (Baroni, Rago, & Toni, 2019), and abstract agent argumentation frameworks (Yu, Chen, Qiao, Shen, & van der Torre, 2021).

In the present paper we consider argumentation frameworks with collective attacks (SETAFs), introduced by Nielsen and Parsons (2006). SETAFs generalize Dung-style AFs in the sense that some arguments can only be effectively defeated by a collection of attackers, yielding a natural representation as a directed hypergraph. Many key semantic properties of AFs have been shown to carry over to SETAFs, see e.g. (Nielsen & Parsons, 2006; Flouris & Bikakis, 2019). Moreover, work has been done on expressiveness (Dvořák, Fandinno, & Woltran, 2019), and translations from SETAFs to AFs (Polberg, 2017; Flouris & Bikakis, 2019). The hypergraph structure of SETAFs has recently been subject of investigation (Dvořák, König, & Woltran, 2021b, 2022a, 2022b). Recent applications of SETAFs include instantiations of inconsistent knowledge bases and similar formalisms (Yun, Vesic, & Croitoru, 2020; König, Rapberger, & Ulbricht, 2022). However, a thorough principle-based analysis of SETAF semantics is still unavailable. In this paper, we will close this gap by investigating the common SETAF semantics w.r.t. a comprehensive selection of principles.

Although we will see that in many cases the behavior generalizes from AFs to the setting with collective attacks, our study also reveals situations where caution is required and thus emphasizes properties we deem natural for AFs. In fact, many AF principles like SCC-recursiveness (Baroni, Giacomin, & Guida, 2005) or the recently introduced modularization property (Baumann, Brewka, & Ulbricht, 2020a) are concerned with partial evaluation of the given graph and step-wise computation of extensions. We will pay special attention to these kind of principles since (a) they require to establish novel technical foundations when generalizing the underlying structure from simple graphs to hypergraphs and (b) have immediate implications for the design of solvers. Along the way, we will also introduce a SETAF version for the reduct of an AF (Baumann, Brewka, & Ulbricht, 2020b) which has proven to be a handy tool when investigating argumentation semantics.

Finally, we will define and characterize graph classes for SETAFs and utilize these findings in conjunction with the formerly established principles to provide a framework for efficient computation. Again, as a starting point we use the existing literature on graph properties for AFs, and generalize the relevant notions to SETAFs. Along the way we point out the exact requirements for tractability in these classes and provide reasonable alternatives that do not allow for efficient reasoning in the case of collective attacks. Finally, we will apply these results in the context of SCC-recursiveness to establish the computational speedup, providing novel algorithms for the evaluation of frameworks along the way.

The main contribution of this paper is to show that our natural extensions of the AF principles are well-behaving for SETAFs and can be utilized for efficient computation. We show that basic properties are preserved, as well as their implications in terms of the structure of extensions. More specifically, this paper is structured as follows.

- First we provide the necessary preliminaries in Section 2. We then generalize and analyze basic principles of abstract argumentation for SETAFs in Section 3. Moreover we introduce novel principles that are trivial for standard AFs, but provide additional insights in the case of SETAFs.

- We propose the E-reduct SF^E for a SETAF SF and a set E of arguments and investigate its core properties, including the modularization property (Section 4). Moreover, we use the reduct to provide alternative characterizations of SETAFs semantics.

- We introduce uninfluenced sets of arguments in SETAFs as the counterpart of unattacked sets in AFs. We then propose and investigate a SETAF version of the directionality and non-interference principles (Section 5) and SCC-recursiveness (Section 6).

- We discuss the computational implications of modularization, directionality and SCC-recursiveness in Section 7. In particular we illustrate the potential for incremental algorithms. We then refine these results in order to be applicable in even more cases. We introduce and analyze graph classes for SETAFs and exemplify their use for efficient computation using the SCC-recursive scheme, generalizing known (parameterized) tractability results from the literature.

- Finally, we conclude in Section 8.

Some technical proofs are moved to the appendix in order to enhance readability of the paper. Sections 3-6 of this article evolved from the paper (Dvořák, König, Ulbricht, & Woltran, 2022) presented at the KR 2022 conference, and from the paper (Dvořák, König, Ulbricht, & Woltran, 2021) presented at the NMR 2021 workshop. Section 7 continues the investigations of (Dvořák, König, & Woltran, 2021a) presented at the JELIA 2021 conference.

2. Background

We briefly recall the definitions of SETAFs and its semantics (see, e.g., (Bikakis, Cohen, Dvořák, Flouris, & Parsons, 2021)). Throughout the paper, we assume a countably infinite domain $\mathfrak{A}$ of possible arguments.

Definition 2.1. *A SETAF is a pair $SF = (A, R)$ where $A \subseteq \mathfrak{A}$ is a finite set of* arguments*, 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.*

If the tail T of an attack (T, h) is a single argument, we usually write (t, h) to denote the set-attack $(\{t\}, h)$. The class of SETAFs where all attacks are of this form amounts to (standard Dung) AFs. Given a SETAF (A, R) and $S, S' \subseteq A$, we write $S \mapsto_R a$ if there is a set $T \subseteq S$ with $(T, a) \in R$. Furthermore, we write $S' \mapsto_R S$ if $S' \mapsto_R a$ for some $a \in S$. For $S \subseteq A$, we use S^+_R to denote the set $\{a \in A \mid S \mapsto_R a\}$ and define the *range* of S (w.r.t. R), denoted $S^\oplus_R$, as the set $S \cup S^+_R$. Moreover, we use $A(SF)$ and $R(SF)$ to identify its arguments A and its attack relation R, respectively.

Example 2.2. *Consider the SETAF $SF = (A, R)$ with arguments $A = \{a, b, c, d, e, f, g, h\}$ and attack relation*

$$R = \{(a, b), (\{b, d\}, c), (b, d), (d, b), (d, e), (e, d),\\ (\{d, f\}, h), (f, g), (g, f), (g, h), (h, g)\};$$

the collective attacks $(\{b, d\}, c), (\{d, f\}, h)$ are highlighted.

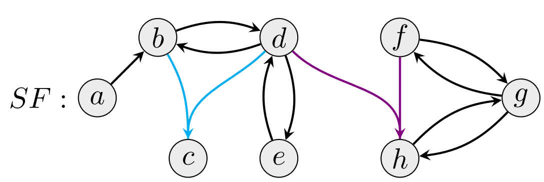

For example, the arguments $\{b, d\}$ are only effectively attacking c (through the blue attack) if both *b and d are accepted. Consequently, it suffices to defeat* either *b or d to defend c.*

The well-known notions of conflict and defense from classical Dung-style AFs naturally generalize to SETAFs.

Definition 2.3. *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 2.4. *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).*

Moreover, we make use of the *characteristic function* Γ_{SF} of a SETAF $SF = (A, R)$, defined as $\Gamma_{SF}(S) = \{a \in A \mid S \text{ defends } a \text{ in } SF\}$ for $S \subseteq A$.

Remark 2.5. *We briefly highlight the difference between the nature of collective attacks in SETAFs and the concept of* accrual of arguments *which is, intuitively speaking, the combination of arguments that have (or attack) the same statement (see e.g., (Bikakis et al., 2021; Rossit, Mailly, Dimopoulos, & Moraitis, 2021)). Notice that collective attacks in SETAFs do not consider any kind of argument strength or preferences between arguments. Moreover, each argument in the tail of a collective attack is essential for the attack to be effective. In contrast the line of research on accrual of arguments allows for different strengths of arguments and considers how the strengths of arguments attacking the same statement are combined and when an attack is effective. That is, adding an additional argument to an accrual, i.e., a set of arguments attacking the same statement, increases the strength of the attack.*

The semantics we study in this work are grounded, admissible, complete, preferred, stable, naive, stage, semi-stable, ideal, and eager semantics, which we will abbreviate by *grd*, *adm*, *com*, *pref*, *stb*, *naive*, *stage*, *sem*, *ideal*, and *eager* respectively (see e.g. (Bikakis et al., 2021)). We denote the set of semantics under our consideration by Σ.

Definition 2.6. *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 $\nexists T \in adm(SF)$ s.t. $T \supset S$,*
- $S \in stb(SF)$, *if $S \mapsto_R a$ for all $a \in A \setminus S$,*
- $S \in naive(SF)$, *if $\nexists T \in cf(SF)$ with $T \supset S$,*
- $S \in stage(SF)$, *if $\nexists T \in cf(SF)$ with $T_R^{\oplus} \supset S_R^{\oplus}$,*
- $S \in sem(SF)$, *if $S \in adm(SF)$ and $\nexists T \in adm(SF)$ s.t. $T_R^{\oplus} \supset S_R^{\oplus}$,*

- $S \in \mathit{ideal}(SF)$, *if* $S \in \mathit{com}(SF)$, $S \subseteq \bigcap_{E \in \mathit{pref}(SF)} E$ *and* $\nexists T \in \mathit{com}(SF)$ *s.t.* $T \subseteq \bigcap_{E \in \mathit{pref}(SF)} E$ *and* $T \supset S$, *and*

- $S \in \mathit{eager}(SF)$, *if* $S \in \mathit{com}(SF)$, $S \subseteq \bigcap_{E \in \mathit{sem}(SF)} E$ *and* $\nexists T \in \mathit{com}(SF)$ *s.t.* $T \subseteq \bigcap_{E \in \mathit{sem}(SF)} E$ *and* $T \supset S$.

The relationship between the semantics has been clarified in (Dvořák, Greßler, & Woltran, 2018; Flouris & Bikakis, 2019; Nielsen & Parsons, 2006) and matches with the relations between the semantics for Dung AFs, i.e. for any SETAF SF:

$$\mathit{stb}(SF) \subseteq \mathit{sem}(SF) \subseteq \mathit{pref}(SF) \subseteq \mathit{com}(SF) \subseteq \mathit{adm}(SF)$$
$$\mathit{stb}(SF) \subseteq \mathit{stage}(SF) \subseteq \mathit{naive}(SF) \subseteq \mathit{cf}(SF)$$

Finally, we introduce the notion of the *projection*, which we will revisit and redefine in Sections 5 and 6.

Definition 2.7 (Projection)**.** *Let* $SF = (A, R)$ *be a SETAF and* $S \subseteq A$. *We define the* projection $SF{\downarrow_S}$ *of* SF *on* S *as* $(S, \{(T', h) \mid (T, h) \in R, h \in S, T' = T \cap S, T' \neq \emptyset\})$.

We illustrate the semantics and the concept of projection in the following example.

Example 2.8. *Consider again the SETAF SF from Example 2.2 (left) and its extensions w.r.t. some semantics* $\sigma \in \Sigma$ *(right).*

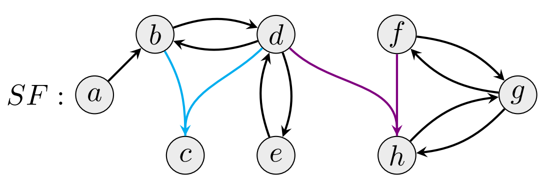

$$\begin{aligned}\mathit{com}(SF) &= \{\{a,c\},\{a,c,e\},\{a,c,d\},\{a,c,g\},\\ &\quad\{a,c,e,g\},\{a,c,d,g\},\{a,c,f\},\\ &\quad\{a,c,d,f\},\{a,c,e,f,h\}\}\\ \mathit{pref}(SF) &= \{\{a,c,e,g\},\{a,c,d,g\},\{a,c,d,f\},\\ &\quad\{a,c,e,f,h\}\}\\ \mathit{grd}(SF) &= \{\{a,c\}\}\end{aligned}$$

Intuitively, the projection "hides" parts of the SETAF while we only concentrate on some remaining arguments. Note however, that the extensions do not in general carry over from the "full" SETAF to its part. We project SF to the arguments $\{b, c, d, h\}$ *(left) and see that the extensions are incomparable to the original framework.*

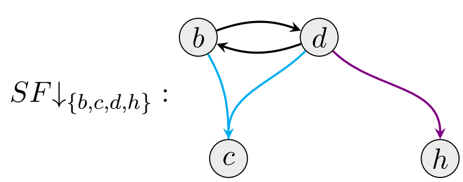

$$\begin{aligned}\mathit{com}(SF) &= \{\emptyset,\{b,c,h\},\{c,d\}\}\\ \mathit{pref}(SF) &= \{\{b,c,h\},\{c,d\}\}\\ \mathit{grd}(SF) &= \{\emptyset\}\end{aligned}$$

After the projection, the argument b becomes acceptable, and c is no longer in every complete extension. Among others, these issues are formally captured and ultimately fixed in different ways in the next sections.

	cf	*grd*	*adm*	*com*	*stb*	*pref*	*naive*	*sem*	*stage*	*ideal*	*eager*
Conflict-freeness	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Defense	✗	✓	✓	✓	✓	✓	✗	✓	✗	✓	✓
Admissibility	✗	✓	✓	✓	✓	✓	✗	✓	✗	✓	✓
Reinstatement	✗	✓	✗	✓	✓	✓	✗	✓	✗	✓	✓
CF-reinstatement	✗	✓	✗	✓	✓	✓	✓	✓	✓	✓	✓
Weak reinstatement	✗	✓	✗	✓	✓	✓	✗	✓	✗	✓	✓
Naivety	✗	✗	✗	✗	✓	✗	✓	✗	✓	✗	✗
I-maximality	✗	✓	✗	✗	✓	✓	✓	✓	✓	✓	✓
Tightness	*✗*	✓	✗	✗	*✗*	✗	*✗*	✗	*✗*	✓	✓
Allowing abstention	✓	✓	✓	✓	✗	✗	✗	✗	✗	✓	✓
Crash resistance	✓	✓	✓	✓	✗	✓	✓	✓	✓	✓	✓
Modularization	✗	✓	✓	✓	✓	✓	✗	✓	✗	✓	✓
Directionality	✓	✓	✓	✓	✗	✓	✗	✗	✗	✓	✗
Weak-directionality	✓	✓	✓	✓	✓	✓	✗	✗	✗	✓	✗
Semi-directionality	✓	✓	✓	✓	✗	✓	✓	✗	✗	✓	✗
Non-interference	✓	✓	✓	✓	✗	✓	✓	✓	✓	✓	✓
SCC-recursiveness	✗	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗
Allowing partial conflicts I	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Allowing partial conflicts II	✓	✓	✓	✓	✓	✓	✓	✗	✗	✓	✗
Allowing partial conflicts III	✓	✗	✗	✗	✓	✗	✓	✗	✗	✗	✗
Tail strengthening	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Attack weakening	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table 1: An overview of our results regarding SETAF principles. Differences from the respective results for AFs are highlighted (*✗*). Note that the SETAF-specific principles Allowing Partial Conflicts I-III, Tail Strengthening, and Attack Weakening are not applicable to AFs.

3. Basic Principles

We start our principle-based analysis of SETAF semantics by generalizing basic principles from AFs. Satisfaction (or non-satisfaction) of principles allows us to distinguish semantics with respect to fundamental properties that are crucial in certain applications.

The principles we consider have natural counterparts for Dung-style AFs, simply by applying them to SETAFs where $|T| = 1$ for each tail. Hence, if the AF counterpart of a principle is violated by a semantics, this carries over to the SETAF principle. We therefore formalize the following observation:

Observation 3.1. *Let P be a SETAF-principle that properly generalizes an AF-principle P^{AF} in the sense that for SETAFs SF with $|T| = 1$ for each $(T, h) \in R(SF)$, every semantics σ satisfies P iff it satisfies P^{AF}. In this case, if a semantics σ does not satisfy P^{AF}, then σ does not satisfy P.*

As all of our principles properly generalize the respective AF-principles, whenever a principle is not satisfied for AFs, this translates to the corresponding SETAF principle as well.

3.1 Basic Properties

Now we follow (van der Torre & Vesic, 2017) and introduce analogous principles for SETAFs. Our first set of principles is concerned with basic properties of semantics.

Principle 3.2 (Conflict-freeness (Dung, 1995; Nielsen & Parsons, 2006))**.** *A semantics σ satisfies* conflict-freeness *if and only if for all SETAFs SF, every $E \in \sigma(SF)$ is conflict-free.*

As conflict-freeness is a basic principle that underlies most semantics by definition, it is not surprising that all semantics under our consideration satisfy this principle.

Proposition 3.3. *Each $\sigma \in \Sigma$ satisfies* conflict-freeness.

The concept of *defense* is central to most classical semantics of abstract argumentation. This central notion is the core of Dung's framework (Dung, 1995) and has been adapted by Nielsen and Parsons (Nielsen & Parsons, 2006) for SETAFs to take collective attacks into account.

Principle 3.4 (Defense (Dung, 1995; Nielsen & Parsons, 2006))**.** *A semantics σ satisfies* defense *if and only if for all SETAFs SF, we have that $E \in \sigma(SF)$ implies $E \subseteq \Gamma_{SF}(E)$.*

Most semantics that satisfy defense are refinements of *adm*. Thus satisfaction of defense is encoded explicitly within their definition. For stable semantics we recall the well-known relation $stb(SF) \subseteq pref(SF)$ for any SETAF SF. The semantics based on conflict-freeness do not satisfy admissibility (as it is the case in AFs), as can be easily seen in the following Example 3.5.

Example 3.5. *Consider the following SETAF SF. We have $\{a, c\} \in cf(SF)$, as well a $\{a, c\} \in naive(SF)$, and $\{a, c\} \in stage(SF)$, but a is not defended by $\{a, c\}$, which means that $\{a, c\}$ is not an extension in any of the admissibility-based semantics.*

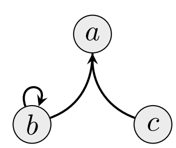

Proposition 3.6. *The principle* defense *is satisfied by grd, adm, com, stb, pref, sem, ideal, and eager, and violated by cf, naive, and stage.*

The *admissibility* principle combines the former two, defense and conflict-freeness.

Principle 3.7 (Admissibility (Dung, 1995; Nielsen & Parsons, 2006))**.** *A semantics σ satisfies* admissibility *if and only if for all SETAFs SF, every $E \in \sigma(SF)$ is admissible.*

Since our semantics $\sigma \in \Sigma$ all satisfy conflict-freeness by Proposition 3.3, we have that these semantics satisfy the *admissibility* principle if and only if they satisfy defense. Hence, the results from Proposition 3.6 carry over.

Proposition 3.8. *The principle* admissibility *is satisfied by grd, adm, com, stb, pref, sem, ideal, and eager, and violated by cf, naive, and stage.*

In the following, we generalize different versions of *reinstatement*. This principle is concerned the question whether or not all defended arguments are indeed included in a given σ-extension E. The principle is thus inspired by the definition of completeness which requires $a \in E$ whenever E defends a. Speaking in terms of the characteristic function, admissible sets satisfy $E \subseteq \Gamma_{SF}(E)$ whereas complete extensions refine this to $E = \Gamma_{SF}(E)$. The reinstatement principle formalizes the "$\supseteq$"-direction.

Principle 3.9 (Reinstatement (Baroni & Giacomin, 2007))**.** *A semantics σ satisfies* reinstatement *if and only if for all SETAFs SF, we have that $E \in \sigma(SF)$ implies $E \supseteq \Gamma_{SF}(E)$.*

Example 3.10. *Recall the SETAF SF from Example 2.2. It is easy to check that $\{a\}$ is admissible in SF. Since $\{a\}$ defends c, we have $\Gamma_{SF}(\{a\}) = \{a, c\} \not\subseteq \{a\}$, which means that admissible semantics violates reinstatement.*

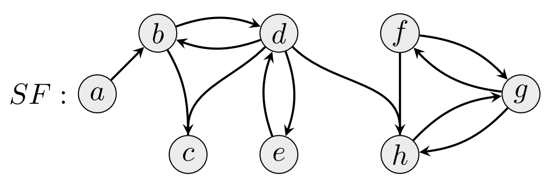

Complete semantics satisfies reinstatement by definition, the other results follow from the known relations $stb(SF) \subseteq sem(SF) \subseteq pref(SF) \subseteq com(SF)$, as well as $ideal(SF) \subseteq com(SF)$, and $eager(SF) \subseteq com(SF)$.

Proposition 3.11. *The principle* reinstatement *is satisfied by grd, com, stb, pref, sem, ideal, and eager, and violated by cf, adm, naive, and stage.*

For admissibility-based semantics, the fundamental lemma (originally due to (Dung, 1995), for SETAFs in (Nielsen & Parsons, 2006)) ensures conflict-freeness for additional defended arguments. Formally, if $E \in adm(SF)$ and $a \in \Gamma_{SF}(E)$, then $E \cup \{a\} \in cf(SF)$. For semantics based on conflict-freeness such as naive or stage, it might happen that some extension E is in conflict with some argument a, although $a \in \Gamma_{SF}(E)$. However, if $E \cup \{a\} \notin cf(SF)$ is the case, then we do not expect it to be a σ-extension anymore (if σ is cf-based). Therefore, the following refinement of reinstatement has been proposed, which explicitly requires $E \cup \{a\}$ to be conflict-free.

Principle 3.12 (CF-Reinstatement (Baroni & Giacomin, 2007))**.** *A semantics σ satisfies* CF-reinstatement *if and only if for all SETAFs SF, we have that $E \in \sigma(SF)$, $a \in \Gamma_{SF}(E)$, and $E \cup \{a\} \in cf(SF)$ imply $a \in E$.*

Due to the fundamental lemma, for admissibility-based semantics this notion simply coincides with reinstatement. However, also *naive* and *stage* satisfy CF-reinstatement, which can be inferred from their respective maximality requirements: assume for $a \notin E$ it holds $E \cup \{a\} \in cf(SF)$, then E cannot be a naive extension as $E \cup \{a\} \supset E$. Finally, recall that $stage(SF) \subseteq naive(SF)$ for every SETAF SF.

Proposition 3.13. *The principle* CF-reinstatement *is satisfied by grd, com, stb, pref, naive, sem, stage, ideal, and eager, and violated by cf, and adm.*

Another possible way to refine reinstatement is by restricting our attention to so-called *strongly defended* arguments. Strong defense was initially defined as the underlying defense notion for strong admissibility (Baroni & Giacomin, 2007; Caminada, 2014). So instead of imposing $E \cup \{a\} \in cf(SF)$ as a premise, we take fewer candidates a into consideration. Strong defense for SETAFs is defined as follows[1].

Definition 3.14. *Given a SETAF $SF = (A, R)$, an argument $a \in A$ is* strongly defended *(in SF) by a set $S \subseteq A$ if for each $(B, a) \in R$ there is an argument $b \in B$ and a set $S' \subseteq S$ such that $(S', b) \in R$, and each $s \in S'$ is strongly defended by $S \setminus \{a\}$.*

Naturally, the induced weakening of reinstatement is given as follows.

Principle 3.15 (Weak reinstatement (Baroni & Giacomin, 2007))**.** *A semantics σ satisfies* weak reinstatement *if and only if for all SETAFs SF, if $E \in \sigma(SF)$ and E strongly defends $a \in A$, then $a \in E$.*

As in the AF case, if a set strongly defends an argument, then it also (classically) defends said argument. Hence, if a semantics satisfies reinstatement, also weak reinstatement is satisfied. The positive results in Table 1 are due to this property. The negative cases follow from Observation 3.1 and the respective counter-examples from AFs.

Proposition 3.16. *The principle* weak reinstatement *is satisfied by grd, com, stb, pref, sem, ideal, and eager, and violated by cf, adm, naive and stage.*

The final two principles we consider in this subsection are concerned with the structure of the σ-extensions. First, naivety checks whether each $E \in \sigma(SF)$ is maximal conflict-free.

Principle 3.17 (Naivety (van der Torre & Vesic, 2017))**.** *A semantics σ satisfies* naivety *if and only if for all SETAFs SF, $E \in \sigma(SF)$ implies that E is $\subseteq$-maximal in $cf(SF)$.*

Again, the negative results are due to Observation 3.1; the positive ones follow from the relation $stb(SF) \subseteq stage(SF) \subseteq naive(SF)$.

Proposition 3.18. *The principle* naivety *is satisfied by stb, naive , and stage and violated by cf, grd, adm, com, pref, sem, ideal, and eager.*

Second, I-maximality is satisfied iff $\sigma(F)$ forms an anti-chain, i.e. no two extensions are in proper subset relation to each other. Here, I-maximality is due to Baroni and Giacomin (2007).

Principle 3.19 (I-maximality (Baroni & Giacomin, 2007))**.** *A semantics σ satisfies* I-maximality *if and only if for all SETAFs SF, if $E, E' \in \sigma(SF)$ and $E \subseteq E'$, then $E = E'$.*

1. Strong admissibility has been generalized to Abstract Dialectical Frameworks (ADFs) (Keshavarzi Zafarghandi, Verbrugge, & Verheij, 2022) and SETAFs can be interpreted as special kind of ADF (Polberg, 2016; Linsbichler, Pührer, & Strass, 2016). In fact, our definition of strong defense is compatible with the respective notions on ADFs.

Oftentimes, I-maximality is directly implemented in the definition of the σ-extensions (most famously *grd* and *pref*). Further results for SETAFs have been shown in (Dvořák et al., 2019).

Proposition 3.20. *The principle* I-maximality *is satisfied by grd, stb, pref, naive, sem, stage, ideal, and eager, and violated by cf, adm, and com.*

3.2 Advanced Principles

The next principle we discuss is called *allowing abstention* (Baroni, Caminada, & Giacomin, 2011). As the name suggests, it allows the underlying semantics to be indecisive in certain scenarios. Formally, suppose we have some target argument a and two extensions $E \in \sigma(SF)$ as well as $E' \in \sigma(SF)$ where $a \in E$, but $a \in (E')^+$; that is, E accepts a, but E' rejects it. In this case, since the status of a is not determined, one might argue that σ should also admit an extension where a is neither accepted nor rejected. This idea is formalized by the allowing abstention principle.

Principle 3.21 (Allowing abstention (Baroni et al., 2011)). *A semantics σ satisfies* allowing abstention *if and only if for all SETAFs $SF = (A, R)$, for all $a \in A$, if there exist $E, E' \in \sigma(SF)$ with $a \in E$ and $a \in (E')^+$, then there also exists some $E'' \in \sigma(SF)$ such that $a \notin (E'')^\oplus$.*

As *grd*, *ideal*, and *eager* always admit a single extension, the principle is trivially satisfied by these semantics. Moreover, allowing abstention is satisfied by complete semantics, since—as in AFs—if there exist $E, E' \in com(SF)$ with $a \in E$ and $a \in E'^+$, this means $a \notin G^\oplus$ where $G \in grd(SF)$. For $\sigma \in \{cf, adm\}$, this follows from $\emptyset \in \sigma(SF)$ for all SETAFs SF.

Proposition 3.22. *The principle* allowing abstention *is satisfied by cf, grd, adm, com, ideal, and eager, and violated by stb, pref, naive, sem, and stage.*

The next principle we discuss is called *crash resistance* (Caminada, Carnielli, & Dunne, 2012). It formalizes that it should not be possible to render certain parts of an argumentation framework completely meaningless by adding a particular set of (disjoint) arguments. This idea is formalized in the definition of a contaminating SETAF.

Definition 3.23. *We call a SETAF $SF' = (A', R')$* contaminating *for a semantics σ if for every SETAF $SF = (A, R)$ with $A \cap A' = \emptyset$, it holds that $\sigma(SF \cup SF') = \sigma(SF')$, where $SF \cup SF'$ is the SETAF $(A \cup A', R \cup R')$.*

That is, the semantics of the given SETAF $SF = (A, R)$ are entirely overwritten due to the presence of SF'. Observe that SF' has this influence on *every* conceivable SETAF SF. The crash resistance principle forbids the existence of such a contaminating SETAF.

Principle 3.24 (Crash resistance). *A semantics σ satisfies* crash resistance *if there is no contaminating SETAF for σ.*

As in the case for AFs, *stb* is the only semantics considered in this paper which is not crash-resistant. The reason is that one can choose SF' to be an isolated odd cycle, yielding $stb(SF \cup SF') = \emptyset$ for any SETAF SF. The other semantics are more robust in this regard and yield $\sigma(SF \cup SF') = \{E \cup E' \mid E \in \sigma(SF), E' \in \sigma(SF')\}$ whenever $A \cap A' = \emptyset$.

Proposition 3.25. *The principle* crash resistance *is satisfied by cf, grd, adm, com, pref, naive, sem, stage, ideal, and eager, and violated by stb.*

The last principle we consider in this subsection is inspired by research on expressiveness in abstract argumentation (Dunne, Dvořák, Linsbichler, & Woltran, 2015; Dvořák, Rapberger, & Woltran, 2020). In this context, the notion of *tightness* has been introduced. It formalizes that if E is a σ-extension and $a \notin E$, then some $b \in E$ must be the culprit for a not being acceptable. Towards formalizing this, we need the notion of *pairs*, i.e. jointly acceptable arguments.

Definition 3.26. *Let SF be a SETAF and σ some semantics. We define the set of* pairs *as $\mathcal{P}airs_\sigma(SF) = \{(a,b) \mid \exists E \in \sigma(SF) \text{ s.t. } \{a,b\} \subseteq E\}$.*

A semantics σ satisfies the tightness principle if for an argument a that does not belong to an extension $E \in \sigma(SF)$ there is some $b \in E$ such that $\{a,b\}$ is not part of any σ-extension, i.e. there is a single argument b in E which can be considered responsible for excluding a.

Principle 3.27 (Tightness)**.** *A semantics σ is* tight *if for all SETAFs $SF = (A,R)$, for all $E \in \sigma(SF)$ and all credulously accepted $a \in A$, the following implication holds: if $E \cup \{a\} \notin \sigma(SF)$, then there is some $b \in E$ such that $(a,b) \notin \mathcal{P}airs_\sigma(SF)$.*

Clearly, any unique status semantics σ, i.e. $|\sigma(SF)| = 1$ for each SF, is tight. However, while Dunne et al. showed that on AFs tightness holds also for conflict-freeness, naive, stable, and stage semantics, this is not the case for SETAFs, as the following example illustrates (see also (Dvořák et al., 2019)).

Example 3.28. *Consider the following SETAF SF.*

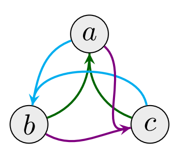

We have naive(SF) = stb(SF) = stage$(SF) = \{\{a,b\},\{b,c\},\{a,c\}\}$. Consider for example $c \notin \{a,b\}$. Tightness would require $(a,c) \notin \mathcal{P}airs_\sigma(SF)$ or $(b,c) \notin \mathcal{P}airs_\sigma(SF)$, but both $\{a,c\}$ and $\{b,c\}$ are σ-extensions. Likewise, the same counter-example illustrates that conflict-free sets are not tight.

Thus, we end up with only the unique status semantics *grd*, *ideal*, and *eager* being tight.

Proposition 3.29. *The principle* tightness *is satisfied by grd, ideal, and eager, and violated by cf, adm, com,pref, naive, sem, stage, and stb.*

3.3 SETAF-Specific Principles

The principles we discussed up until this points where inspired by known AF principles and have been suitably adjusted to SETAFs. In this section we want to introduce *genuine* SETAF principles, i.e., we discuss properties which are not applicable or trivialize for standard Dung-AFs.

Towards our first SETAF principle, observe that a conflict within some set E of arguments requires the whole tail of a corresponding attack to be contained in E; that is, there has to be some $(T, a) \in R$ with $a \in E$ and $T \subseteq E$. The underlying intuition is that attacks are only "active" if the whole tail is accepted. Semantics which adhere to this intuition should be able to distinguish between attacks that are fully active, i.e., $T \subseteq E$ and attacks which are only partially active, i.e., $T \cap E \neq \emptyset$, but $T \nsubseteq E$. We therefore consider the following notion of a partial conflict.

Definition 3.30. *Let $SF = (A, R)$ be a SETAF and $E \subseteq A$. We say E contains a* partial conflict *whenever there is some $(T, a) \in R$ with $a \in E$ and $T \cap E \neq \emptyset$ as well as $T \nsubseteq E$.*

The way SETAF semantics are designed, semantics should usually allow partial conflicts (APC).

Principle 3.31 (Allowing partial conflicts I)**.** *A semantics σ satisfies the principle* allowing partial conflicts I *if there is some SETAF SF and some extension $E \in \sigma(SF)$ s.t. E contains some partial conflict.*

Observe that for Dung-AFs partial conflicts never exist, since the conditions $T \cap E \neq \emptyset$ as well as $T \nsubseteq E$ can never be met simultaneously for a singleton T. Hence this principle is trivially violated for AFs. For SETAFs, it is also easy to see that all semantics under our consideration satisfy APC I.

Proposition 3.32. *Each $\sigma \in \Sigma$ satisfies* allowing partial conflicts I.

We can strengthen this requirement as follows. Intuitively, we say for a given extension E we can add a new attack (T, h) and still have E as an extension in the remaining framework if at least one argument in T is already attacked by E.

Principle 3.33 (Allowing partial conflicts II)**.** *A semantics σ satisfies the principle* allowing partial conflicts II *if for every SETAF $SF = (A, R)$ and every $E \in \sigma(SF)$ it holds for all $h \in E, T_1 \subseteq E, \emptyset \subsetneq T_2 \subseteq E^+$ also $E \in \sigma(SF')$ where $SF' = (A, R \cup \{(T_1 \cup T_2, h)\})$.*

Since in admissibility-based semantics this added attack has no effect (as the tail is attacked), these semantics satisfy the principle. The exception to this rule is semi-stable semantics, as the introduction of the new attack might lead to a different preferred extension with a larger range. Finally, since no conflict is introduced, also *cf* and *naive* satisfy APC II. The counterexamples for *sem*, *stage*, and *eager* are illustrated in Example 3.36.

Proposition 3.34. *The principle* allowing partial conflicts II *is satisfied by cf, grd, adm, com, stb, pref, naive, and ideal, and violated by sem, stage, and eager.*

In APC II we require that for the introduced attack $(T_1 \cup T_2, h)$ there is at least one argument in T_2, i.e., there is at least one argument in $T_1 \cup T_2$ that is *attacked by E*. However, if we only require an argument that is *not in E* (instead of *attacked by E*), we end up with a stronger requirement, captured in the following principle.

Principle 3.35 (Allowing partial conflicts III)**.** *A semantics σ satisfies the principle* allowing partial conflicts III *if for every SETAF $SF = (A, R)$ and every $E \in \sigma(SF)$ it holds for all $h \in E, T_1 \subseteq E, \emptyset \subsetneq T_2 \subseteq A \setminus E$ also $E \in \sigma(SF')$ where $SF' = (A, R \cup \{(T_1 \cup T_2, h)\})$.*

First note that for stable APC II and APC III coincide, as for any $E \in stb(SF)$ it holds $E^+ = A \setminus E$ by definition. Similarly, for the conflict-freeness based semantics *cf* and *naive* it plays no role whether an argument is attacked or not, hence, APC III is still satisfied for these semantics. Most admissibility-based semantics under our consideration violate APC III, as the introduction of an attack (T, h) might lead to a situation where h is not defended, as Example 3.36 illustrates. Clearly, we have that if σ satisfies ACP III then σ satisfies APC II, and if σ satisfies APC II then σ satisfies APC I. The reverse does not hold, as the results in Table 1 illustrate.

Example 3.36. *Consider the SETAF (a) SF_a (first without the attack $(\{a, b\}, c)$). It is easy to check that $\{a, c, f\}$ is a semi-stable, stage, and eager extension. However, if we add the attack $(\{a, b\}, c)$ the set $\{a, b, d\}$ becomes a stable extension, and is in fact the only stable extension of the resulting SETAF. Hence, $\{a, b, d\}$ is also the only semi-stable,stage, and eager extension, i.e., $\{a, c, f\}$ is no longer an extension. This violates APC II (and, hence, APC III) for sem, stage, and eager. In SETAF (b) SF_b (first without the attack $(\{a, b\}, c)$) the set $\{a, c\}$ is grounded, admissible, complete, preferred, and ideal. If we again add the attack $(\{a, b\}, c)$ the only extension w.r.t. these semantics is $\{a\}$, violating APC III.*

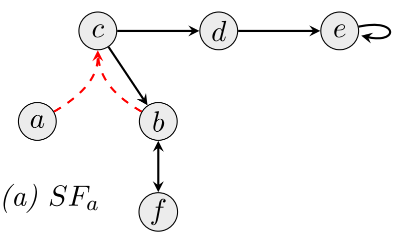

(a) SF_a

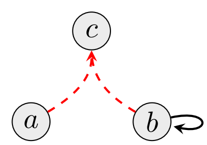

(b) SF_b

Combining these considerations, we get the following results for APC III.

Proposition 3.37. *The principle* allowing partial conflicts III *is satisfied by cf,stb, and naive, and violated by grd, adm, com, pref,sem, stage, ideal, and eager.*

The underlying idea of a collective attack (T, a) is that *all* arguments in T are required in order to defeat a. Hence, an attack (T', a) is in a certain sense stronger than (T, a) if $T' \subseteq T$. In the same spirit, if $T \subseteq E$ for some extension $E \in \sigma(F)$, then we make E stronger if (T, a) is replaced by some stronger attack.

Principle 3.38 (Tail Strengthening)**.** *A semantics σ satisfies* tail strengthening *if for all SETAFs $SF = (A, R)$ and for all $E \in \sigma(SF)$ the following implication holds: if $(T, a) \in R$ with $T \subseteq E$, then we also have $E \in \sigma(SF')$ where $SF' = (A, R')$ with $R' = (R \setminus \{(T, a)\}) \cup \{(T', a)\}$ for some $T' \subseteq T$.*

Vice versa, suppose we have an argument $a \in E$ and some in-coming attack $(T, a) \in R$. If we make this attack weaker, we expect E still to be represent a jointly acceptable point of view. Formally:

Principle 3.39 (Attack Weakening)**.** *A semantics σ satisfies* attack weakening *if for all SETAFs $SF = (A, R)$ and for all $E \in \sigma(SF)$ the following implication holds: if $(T, a) \in R$ with $a \in E$, then we also have $E \in \sigma(SF')$ where $SF' = (A, R')$ with $R' = (R \setminus \{(T, a)\}) \cup \{(T', a)\}$ for some $T \subseteq T'$.*

For the semantics considered in this paper, it follows by definition that both properties are satisfied.

Proposition 3.40. *Each $\sigma \in \Sigma$ satisfies* tail strengthening *and* attack weakening.

4. Reduct and Modularization

In the remainder of this paper, our analysis will put strong emphasis on computational aspects and the partial evaluation of SETAFs. In this section, we will provide the first steps into this direction. First we will introduce the so-called SETAF reduct which corresponds to the resulting SETAF after the status of a certain subset of the arguments is decided. Based on this, we will generalize the modularization property (Baumann et al., 2020a), which formalizes how to compute extensions step-wise by means of the reduct. As an aside, the modularization property yields concise alternative characterizations for the classical semantics.

4.1 The SETAF Reduct

For many of our subsequent results, the *reduct* of a SETAF w.r.t. a given set E of arguments will play a central role. Intuitively, the reduct w.r.t. E represents the SETAF that result from "accepting" E and rejecting what is defeated by E, while not deciding on the remaining arguments. To illustrate the idea, consider the following example:

Example 4.1. *Recall the SETAF SF from Example 2.2. Consider the singleton $\{a\}$. If we view a as accepted, then b is rejected. This means that the attack from b to d can be disregarded. However, we also observe that c cannot be attacked anymore since attacking it requires both b and d, but b is rejected. Now consider $\{f\}$. Interpreting f as accepted renders g rejected. In order to attack h, only d is still required. Thus, if we let $E = \{a, f\}$, then we expect the SETAF reduct SF^E –with the intuitive meaning that a and f are set to true– to look as follows.*

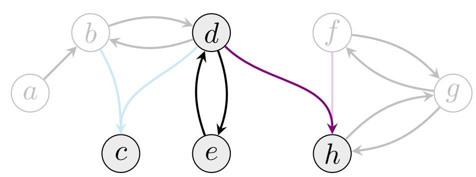

That is, in the reduct SF^E, we only need to consider arguments that are still undecided, i.e. all arguments neither in E nor attacked by E. In contrast to the AF-reduct (Baumann et al., 2020a), it might happen that some attacks are preserved that involve deleted arguments, i.e. the attack is only partially evaluated. In particular, if the arguments in the tail of an attack are "accepted" (i.e. in E), the attack can still play a role in attacking or defending. If the tail of an attack (T, h) is already attacked by E, we can disregard (T, h). By this, we get the following formal definition of the SETAF reduct.

Definition 4.2. *Given a SETAF $SF = (A, R)$ and $E \subseteq A$, the E-*reduct *of SF is the SETAF $SF^E = (A', R')$, with*

$$A' = A \setminus E_R^{\oplus}$$
$$R' = \{(T \setminus E, h) \mid (T, h) \in R,\ T \cap E_R^{+} = \emptyset,\ T \not\subseteq E,\ h \in A'\}$$

Thereby, the condition $T \cap E_R^+ = \emptyset$ captures cases like the attack $(\{b,d\},c)$ from our example: b is attacked by E, and thus, the whole attack gets removed. The reason why we take $(T \setminus E, h)$ as our novel attacks is the partial evaluation as in the attack $(\{d,f\},h)$ after setting f to true: when additionally accepting d, we "activate" the attack against h.

Example 4.3. *Given the SETAF SF from Example 4.1 as well as $E = \{a,f\}$ as before, the reduct SF^E is the SETAF depicted above, i.e. $SF^E = \{A', R'\}$ with $A' = \{c,d,e,h\}$ and $R' = \{(d,e),(e,d),(d,h)\}$.*

We start our formal investigation of the reduct with a technical lemma to settle some basic properties.

Lemma 4.4. *Given a SETAF $SF = (A,R)$ and two disjoint sets $E, E' \subseteq A$. Let $SF^E = (A', R')$.*

1. *If there is no $S \subseteq A$ s.t. $S \mapsto_R E'$, then the same is true in SF^E.*
2. *Assume E does not attack $E' \in cf(SF)$. Then, E defends E' in SF iff there is no $S' \subseteq A'$ s.t. $S' \mapsto_{R'} E'$.*
3. *Let $E \in cf(SF)$. If $E \cup E'$ does not attack E in SF and $E' \subseteq A'$, with $E' \in cf\left(SF^E\right)$ then $E \cup E' \in cf(SF)$.*
4. *Let $E \cup E' \in cf(SF)$. If $E' \mapsto_{R'} a$, then $E \cup E' \mapsto_R a$.*
5. *If $E \cup E' \in cf(SF)$, then $SF^{E \cup E'} = \left(SF^E\right)^{E'}$.*

4.2 The Modularization Property

Having established the basic properties of the SETAF reduct, we are now ready to introduce the modularization property (Baumann et al., 2020a).

Principle 4.5 (Modularization)**.** *A semantics σ satisfies* modularization *if for all SETAFs SF, for every $E \in \sigma(SF)$ and $E' \in \sigma(SF^E)$, we have $E \cup E' \in \sigma(SF)$.*

Modularization allows us to build extensions iteratively. After finding such a set $E \subseteq A$ we can efficiently compute its reduct SF^E and pause before computing an extension E' for the reduct in order to obtain a larger extension $E \cup E'$ for SF. Hence, this first step can be seen as an intermediate result that enables us to reduce the computational effort of finding extensions in SF, as the arguments whose status is already determined by accepting E do not have to be considered again. Instead, we can reason on the reduct SF^E (see Section 7). In the following, we establish the modularization property for admissible and complete semantics.

Theorem 4.6. *Let SF be a SETAF, $\sigma \in \{adm, com\}$ and $E \in \sigma(SF)$.*

1. *If $E' \in \sigma(SF^E)$, then $E \cup E' \in \sigma(SF)$.*
2. *If $E \cap E' = \emptyset$ and $E \cup E' \in \sigma(SF)$, then $E' \in \sigma(SF^E)$.*

Proof. (for $\sigma = adm$) Let $SF^E = (A', R')$.

1) Since E is admissible and $E' \subseteq A'$, E' does not attack E. By Lemma 4.4, item 3, $E \cup E' \in cf(SF)$. Now assume $S \mapsto_R E \cup E'$. If $S \mapsto_R E$, then $E \mapsto_R S$ by admissibility of E. If $S \mapsto_R E'$, there is $T \subseteq S$ s.t. $(T, e') \in R$ for some $e' \in E'$. In case $E \mapsto_R T$, we are done. Otherwise, $(T \setminus E, e') \in R'$ and by admissibility of E' in SF^E, $E' \mapsto_{R'} T \setminus E$. By Lemma 4.4, item 4, $E \cup E' \mapsto_R T \setminus E$.

2) Now assume $E \cup E' \in adm(SF)$. We see $E' \in cf\left(SF^E\right)$ as follows: If $(T', e') \in R'$ for $T' \subseteq E'$ and $e' \in E'$, then there is some $(T, e') \in R$ with $T' = T \setminus E$. Hence $E \cup E' \mapsto_R E'$, contradiction. Now assume E' is not admissible in SF^E, i.e. there is $(T', e') \in R'$ with $e' \in E'$ and E' does not counterattack T' in SF^E. Then there is some $(T, e') \in R$ with $T' = T \setminus E$ and $T \cap E_R^+ = \emptyset$. By admissibility of $E \cup E'$, $E \cup E' \mapsto_R T$, say $(T^*, t) \in R$, $T^* \subseteq E \cup E'$ and $t \in T$. Since $E \cup E'$ is conflict-free, $T^* \cap E_R^+ = \emptyset$ and thus we either have a) $T^* \subseteq E$, contradicting $T \cap E_R^+ = \emptyset$, or b) $(T^* \setminus E, t) \in R'$ and $t \in T'$, i.e. E' counterattacks T' in SF^E contradicting the above assumption.

For *com* semantics we utilize the results for *adm*:

1) We have $E \cup E' \in adm(SF)$. Moreover, E' is complete, i.e. $(SF^E)^{E'}$ does not contain unattacked arguments in the reduct SF^E (see Proposition 4.7). Lemma 4.4, item 5, implies that $SF^{E \cup E'}$ does not contain unattacked arguments, either. Hence $E \cup E' \in com(SF)$.

2) Given $E \cup E' \in com(SF)$ we have $E' \in adm\left(SF^E\right)$, as established. Regarding completeness, we again use the fact that $SF^{E \cup E'} = (SF^E)^{E'}$ does not contain unattacked arguments. □

Note that the modularization property also holds for *stb*, *pref*, and *sem* semantics. However, the only admissible set in the reduct w.r.t. a stable/preferred/semi-stable extension is the empty set, rendering the property trivial. The exact relation is captured by the following alternative characterizations of the semantics under our consideration.

Proposition 4.7. *Let $SF = (A, R)$ be a SETAF, $E \in cf(SF)$ and $SF^E = (A', R')$.*

1. *$E \in stb(SF)$ iff $SF^E = (\emptyset, \emptyset)$,*
2. *$E \in adm(SF)$ iff $S \mapsto_R E$ implies $S \setminus E \not\subseteq A'$,*
3. *$E \in pref(SF)$ iff $E \in adm(SF)$ and $adm(SF^E) = \{\emptyset\}$,*
4. *$E \in com(SF)$ iff $E \in adm(SF)$ and $grd(SF^E) = \{\emptyset\}$,*
5. *$E \in sem(SF)$ iff $E \in pref(SF)$ and there is no $E' \in pref(SF)$ s.t. $A(SF^{E'}) \subsetneq A(SF^E)$.*

From the characterization of complete semantics provided in Proposition 4.7 we infer that for any SETAF SF the complete extensions $E \in com(SF)$ satisfy $grd(SF^E) = \{\emptyset\}$ implying modularization for *grd*. Moreover, as the grounded extension G is the least complete extension, we can utilize modularization of *adm* and obtain G by the following procedure: (1) add the set of unattacked arguments U into G, (2) repeat step (1) on SF^U until there are no unattacked arguments.

We have two cases left to discuss, namely *eager* and *ideal* semantics. Both satisfy the modularization property, because they only admit the empty set as admissible extension in their corresponding reduct SF^E (as in the case of e.g. *sem* semantics). Since this is however not as easy to see, we will give the necessary proofs in detail here. We follow the proof technique of the AF case (Friese & Ulbricht, 2021). First we show that the property formalized in Theorem 4.6 also holds for semi-stable semantics. This will be useful later since *eager* semantics build upon semi-stable extensions.

Proposition 4.8. *Let SF be a SETAF and let $E \in sem(SF)$. Suppose $E = E' \cup E''$ with $E' \cap E'' = \emptyset$ for some $E' \in adm(SF)$. Then $E'' \in sem(SF^{E'})$.*

Proof. We already know $E'' \in adm(SF^{E'})$ since $sem(SF) \subseteq adm(SF)$. Now assume E'' is not semi-stable in $F^{E'}$. Then there is some admissible $S \in adm\left(SF^{E'}\right)$ with $(E'')^{\oplus} \subsetneq S^{\oplus}$. Since E'' and S occur in $SF^{E'}$, this immediately yields $E^{\oplus} = (E' \cup E'')^{\oplus} \subsetneq (E' \cup S)^{\oplus}$. Since by modularization we have $E' \cup S \in adm(SF)$, we infer $E \notin sem(SF)$, a contradiction. □

Next we show that the reduct w.r.t. some eager extension admits only $\emptyset$ as admissible set.

Proposition 4.9. *If $E \in eager(SF)$, then $eager\left(SF^E\right) = \{\emptyset\}$.*

Proof. Let $SF = (A, R)$ be a SETAF and let $E \in eager(SF)$. Consider the reduct SF^E and assume $E' \in eager(SF^E)$ is not empty. Let S be a semi-stable extension of SF. By definition of *eager*, we have that $E \subseteq S$. Our goal is to show $E' \subseteq S$ as well, yielding a contradiction since $E \cup E' \in com(SF)$ by modularization of *com*; since S is arbitrary, the eager extension of SF must then contain $E \cup E'$. To this end note that $S = E \cup S'$ for $E \in adm(SF)$ and some set S' of arguments. By the above Proposition 4.8, $S' \in sem\left(SF^E\right)$ and hence $E' \subseteq S' \subseteq S$ and we are done. □

Since $\emptyset$ is thus the only candidate extension in the reduct SF^E, we immediately get satisfaction of the modularization property.

Corollary 4.10. *The eager semantics satisfies modularization.*

In order to lift the above proof technique to *ideal* as well it suffices to note the following adjustment to Proposition 4.8.

Proposition 4.11. *Let SF be a SETAF and let $E \in pref(SF)$. Suppose $E = E' \cup E''$ with $E' \cap E'' = \emptyset$ for some $E' \in adm(SF)$. Then $E'' \in pref(SF^{E'})$.*

Proof. According to Proposition 4.7, we have that $E \in pref(SF)$ if and only if $E \in adm(SF)$ and SF^E does not possess any admissible argument. We already know admissibility of E'' in $SF^{E'}$. Moreover, $SF^E = (SF^{E'})^{E''}$ does not contain admissible arguments; thus we are done. □

This yields the same behavior for *ideal* as well. First, we again infer that the reduct does not tolerate any non-empty extension.

Proposition 4.12. *If $E \in ideal(F)$, then $ideal\left(F^E\right) = \{\emptyset\}$.*

Proof. We reason as in the proof of Proposition 4.9 with $S \in pref(SF)$ instead of $S \in sem(SF)$. □

Corollary 4.13. *The ideal semantics satisfies modularization.*

5. Directionality and Non-Interference

In this section we discuss the principles directionality and non-interference. Intuitively, these principles give information about the behavior of separate parts of a framework. Beside being informative regarding the behavior of semantics, this principles also have computational implications. In order to formalize this separation-property, we start of with the notion of unattacked sets of arguments[2]. For directionality (Baroni & Giacomin, 2007) we have to carefully consider this notion in order to obtain a natural generalization of the AF case preserving the intended meaning. A naive definition of unattacked sets will lead to nonsensical results: assume a set S is unattacked in a SETAF $SF = (A, R)$ whenever it is not attacked from "outside", i.e. if the condition $A \setminus S \not\mapsto_R S$ holds.

Example 5.1. *Consider now the following SETAF (a) and its projections (b), (c) w.r.t. the "unattacked" set $S = \{a, c\}$.*

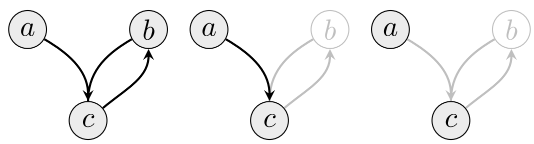

(a) *(b)* *(c)*

Note that $\{a, c\}$ is stable in (a). If we now consider the projection $SF{\downarrow_S}$—see (b)—we find that $\{a, c\}$ is not stable, falsifying directionality. However, one might argue that this is due to the credulous nature of our projection-notion. We could easily consider a different proper generalization of the projection, namely $SF{\downarrow^*_S} = (S, \{(T, h) \mid (T, h) \in R, T \cup \{h\} \subseteq S\})$. In this more skeptical version we delete attacks if any of the arguments in the tail are not in the projected set—see (c). However, we still cannot obtain the desired results: in (a) we find $\{a\}$ to be the unique grounded extension, while in (c) $\{a, c\}$ is grounded, again falsifying directionality. As for the directionality principle we do not want to add additional arguments or attacks and we exhausted all possible reasonable projection notions for this small example, we conclude that the underlying definition of unattacked *sets was improper. We therefore suggest a different definition—and at the same time suggest to think of these sets rather as "uninfluenced" than "unattacked". In AFs, clearly both notions coincide. However, we still argue that the concept of "influence" captures the true nature of directionality in a more intuitive and precise manner. Moreover, note that in the case of uninfluenced sets both notions of projection coincide, as well as the notion of* restriction *(see Definition 6.6) for arbitrary sets $D \subseteq A \setminus S$.*

Towards the formal definition of influence, we utilize the notion of the *primal graph* of a SETAF (Dvořák et al., 2021a). We will use this extension of our graph-related terminology to the directed hypergraph structure of SETAFs several times in the remaining part of this paper as a starting point for structural properties. Intuitively, collective attacks are "split up" in order to obtain a directed graph with a similar structure as the original SETAF.

Definition 5.2 (Primal Graph)**.** *Let $SF = (A, R)$ be a SETAF. Its* primal graph *is defined as $primal(SF) = (A, R')$ with $R' = \{(t, h) \mid (T, h) \in R, t \in T\}$.*

2. While in the previous section we used "unattacked arguments", i.e. arguments that are not the head of any attack, unattacked *sets of arguments* allow for attacks within the set.

Example 5.3. *Recall the SETAF SF. Its primal graph primal(SF) looks as follows.*

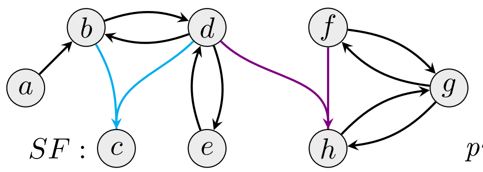

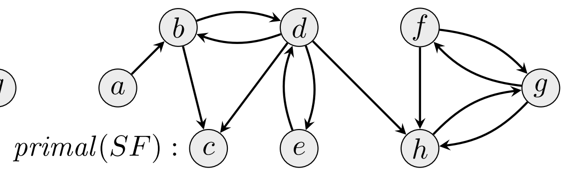

Definition 5.4 (Influence)**.** *Let $SF = (A, R)$ be a SETAF. An argument $a \in A$* influences *$b \in A$ if there is a directed path from a to b in $primal(SF)$. A set $U \subseteq A$ is* uninfluenced *in SF if no $a \in A \setminus U$ influences any $b \in U$. We denote the set of uninfluenced sets by $US(SF)$.*

Utilizing this notion, we can properly generalize directionality (Baroni & Giacomin, 2007).

Principle 5.5 (Directionality)**.** *A semantics σ satisfies* directionality *if for all SETAFs SF and every $U \in US(SF)$ it holds $\sigma(SF{\downarrow_U}) = \{E \cap U \mid E \in \sigma(SF)\}$.*

Moreover, weaker versions of directionality have been proposed which require only a subset relation (van der Torre & Vesic, 2017):

Principle 5.6 (Semi Directionality)**.** *A semantics σ satisfies* semi directionality *if for all SETAFs SF and every $U \in US(SF)$ it holds $\sigma(SF{\downarrow_U}) \subseteq \{E \cap U \mid E \in \sigma(SF)\}$.*

Principle 5.7 (Weak Directionality)**.** *A semantics σ satisfies* weak directionality *if for all SETAFs SF and every $U \in US(SF)$ it holds $\sigma(SF{\downarrow_U}) \supseteq \{E \cap U \mid E \in \sigma(SF)\}$.*

We will revisit directionality at the end of the next section, as we can utilize SCC-recursiveness to show that *grd*, *com*, and *pref* satisfy directionality. In contrast, this is not possible for *eager* and *ideal* semantics, so we investigate these two cases directly.

Let us start with *eager* semantics. As the following examples show, *eager* satisfies neither semi directionality nor weak directionality.

Example 5.8. *Let $F = (A, R)$ be the following AF (van der Torre & Vesic, 2017, Figure 3):*

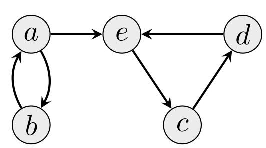

Let $U = \{a, b\}$. We have $sem(F) = \{\{a, c\}\}$ and thus $eager(F) = \{\{a, c\}\}$ as well. Thus $\{E \cap U \mid E \in eager(F)\} = \{\{a\}\}$. On the other hand, $sem(SF{\downarrow_U}) = \{\{a\}, \{b\}\}$ and thus, $eager(SF{\downarrow_U}) = \{\emptyset\}$, i.e. weak directionality is violated.

Example 5.9. *Now let Let $F = (A, R)$ be the following AF (van der Torre & Vesic, 2017, Figure 8):*

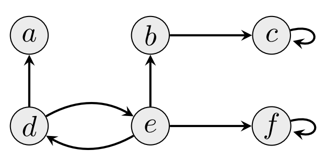

Let $U = \{d, e, f\}$*. We have* $sem(F) = \{\{a, e\}, \{d, b\}\}$ *and thus* $eager(F) = \{\emptyset\}$*. On the other hand,* $sem(SF{\downarrow_U}) = \{\{e\}\}$ *and thus,* $eager(SF{\downarrow_U}) = \{\{e\}\}$*. Hence semi directionality is violated.*

Now let us turn to ideal semantics. We show that directionality is satisfied. Our proof follows the technique from the AF case (Baroni & Giacomin, 2007). The required structural properties also hold for SETAFs. Therefore, we only require minor adjustments to reason analogously in our setting.

Lemma 5.10. *Let* $SF = (A, R)$ *be a SETAF. The the unique ideal extension* S *satisfies*

$$S = \bigcup_{E \in adm(SF) : \forall P \in pref(SF) : E \subseteq P} E$$

This auxiliary lemma is a convenient characterization of *ideal* in order to infer directionality as follows.

Proposition 5.11. *The semantics ideal satisfies directionality.*

Proof. Let $SF = (A, R)$ be a SETAF and suppose $U \in US(SF)$. We have to show $ideal(SF{\downarrow_U}) = \{E \cap U \mid E \in ideal(SF)\}$. Due to Lemma 5.10 it suffices to show

$$\bigcup_{E \in adm(SF) : \forall P \in pref(SF) : E \subseteq P} E \cap U = \bigcup_{E \in adm(SF{\downarrow_U}) : \forall P \in pref(SF{\downarrow_U}) : E \subseteq P} E$$

($\subseteq$) Let E be an arbitrary set in $adm(SF)$. We show the claim for this particular set and thus, the same holds for the union over all extensions in $adm(SF)$ as well. Due to directionality of *adm* semantics, $E \in adm(SF)$ implies $E \cap U \in adm(SF{\downarrow_U})$. Therefore, we have to show that $E \cap U$ is a subset of each preferred extension in $SF{\downarrow_U}$ and thus, $E \cap U$ is part of the union of the right-hand side.

Now, for each $P \in pref(SF)$ we have $E \cap U \subseteq P \cap U$, i.e.

$$\forall P \in pref(SF) : E \cap U \subseteq P \cap U. \tag{1}$$

By directionality of *pref* semantics, $\{P \cap U \mid P \in pref(SF)\} = pref(SF{\downarrow_U})$. This turns (1) into

$$\forall P \in pref(SF{\downarrow_U}) : E \cap U \subseteq P.$$

which we had to show.

($\supseteq$) Now let $E \in adm(SF{\downarrow_U})$. By definition of admissibility, it is clear that $E \in adm(SF)$ follows. For each $P \in pref(SF{\downarrow_U})$ it follows that $E \subseteq P$, i.e.

$$\forall P \in pref(SF{\downarrow_U}) : E \subseteq P. \tag{2}$$

Again by directionality, we turn (2) into

$$\forall P \in pref(SF) : E \subseteq P \cap U \subseteq P.$$

which proves the claim. □

Similarly, we generalize *non-interference* (Caminada et al., 2012), which has an even stronger requirement. $U \subseteq A$ is *isolated* in $SF = (A, R)$, if U is uninfluenced and $A \setminus U$ is uninfluenced, i.e. there are no edges in $primal(SF)$ between U and $A \setminus U$.

Principle 5.12 (Non-interference)**.** *A semantics σ satisfies* non-interference *iff for all SETAFs SF and all isolated $S \subseteq A(SF)$, it holds $\sigma(SF{\downarrow_U}) = \{E \cap U \mid E \in \sigma(SF)\}$.*

Clearly, directionality implies non-interference. It is easy to see from the respective definitions that also naive, semi-stable, ideal, eager, and stage semantics satisfy non-interference.

6. SCC-Recursiveness

In graph theory, the notion of a strongly connected component (SCC) is a widely known concept. An SCC consists of a set S of nodes s.t. for any $a, b \in S$ there is a directed path from a to b within the given graph. SCC-recursiveness (Baroni et al., 2005) formalizes the intuition that the acceptance status of an argument depends only on its ancestors—i.e., the arguments that feature a directed path to the argument in question. If some semantics satisfies SCC-recursiveness, one can construct all SCCs of a given graph and then compute resp. verify extensions step-wise, by working along the SCCs. This provides theoretical insights as it formalizes the independence of arguments of their child SCCs, but also provides us with computational advantages as we will see in Section 7.

In Section 5 we considered the concept of *influence*. In a nutshell, an argument a "influences" an argument b in a SETAF SF if there is a directed path from a to b in $primal(SF)$. It is therefore reasonable to investigate SCCs with this idea in mind. In particular, our definition of SCCs captures the equivalence classes w.r.t. the influence relation.

Definition 6.1 (SCCs)**.** *Let SF be a SETAF. By SCCs(SF) we denote the set of strongly connected components of SF, which we define as the sets of arguments contained in the strongly connected components of primal(SF).*

Example 6.2. *Recall our SETAF from before.*

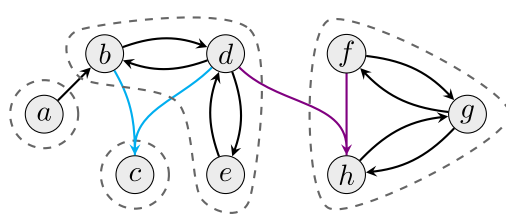

In this SETAF, we have the four SCCs $\{a\}, \{b, d, e\}, \{c\}$, and $\{f, g, h\}$, as depicted in dashed lines.

Analogously to (Baroni et al., 2005), we partition the arguments in defeated, provisionally defeated and undefeated ones. Intuitively, accepting a defeated argument would lead to a conflict, the provisionally defeated cannot be defended and will therefore be rejected (while not being irrelevant for defense of other arguments), and the undefeated form the candidates for extensions. We obtain the following formal definition of the sets we just described.

Definition 6.3. *Let $SF = (A, R)$ be a SETAF. Moreover, let $E \subseteq A$ be a set of arguments and $S \in SCCs(SF)$ be an SCC. We define the set of defeated arguments $D_{SF}(S, E)$, provisionally defeated arguments $P_{SF}(S, E)$, and undefeated arguments $U_{SF}(S, E)$ w.r.t. S, E as*

$$\begin{aligned}
D_{SF}(S,E) &= \{a \in S \mid E \setminus S \mapsto_R a\},\\
P_{SF}(S,E) &= \{a \in S \mid A \setminus (S \cup E^+) \mapsto_R a\} \setminus D_{SF}(S,E),\\
U_{SF}(S,E) &= S \setminus (D_{SF}(S,E) \cup P_{SF}(S,E)).
\end{aligned}$$

Moreover, we set $UP_{SF}(S,E) = U_{SF}(S,E) \cup P_{SF}(S,E)$.

It is important to note that all these sets are calculated w.r.t. a given set candidate E, i.e. the purpose is to verify whether E is some σ-extension.

Example 6.4. *Recall the SETAF from above. Let $S = \{b,d,e\}$ be the SCC under consideration.*

Take the admissible extension $E = \{a, e\}$. We have that $D_{SF}(S, E) = \{b\}$ since the argument a from the parent SCC $\{a\}$ defeats b; observe that $d \notin D_{SF}(S, E)$ since d is only defeated by e which is part of the given SCC S. Moreover, $P_{SF}(S, E) = \emptyset$ and hence $U_{SF}(S, E) = \{d, e\}$.

Consider now $E' = \{d\}$. Then $D_{SF}(S, E') = \emptyset$ because no argument within S is defeated from an argument in E' occurring in a parent SCC. However, $P_{SF}(S, E') = \{b\}$ reflecting that b cannot be defended (for this we would have to defeat a, but from within the given SCC S this is impossible). Therefore, $U_{SF}(S, E') = \{b, d, e\} = S$.

In order to formalize SCC-recursiveness, we need the notion of the *restriction*. It will be convenient in order to evaluate our given extension SCC-wise, since in each step we can remove the defeated arguments $D_{SF}(S, E)$ and thus restrict our attention to $U_{SF}(S, E)$. For classical Dung-AFs, the restriction coincides with the *projection* from Definition 2.7. However, in the following we will argue that the projection does not capture the intricacies of this process. Ultimately, we will see that for a reasonable restriction we need semantic tools that are similar to the reduct SF^E. For that, we revisit Example 5.1.

Example 6.5. *Consider the following SETAFs SF and SF'.*

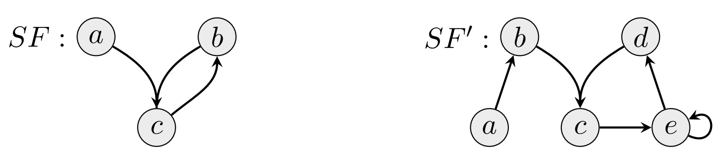

Assume we accept the argument a in SF. Now for the remaining SCC $\{b, c\}$ the projection $SF{\downarrow_{\{b,c\}}}$ contains the attacks (b, c) and (c, b), as one might expect.

Regarding SF', assume we accept a and therefore reject b. The projection $SF'{\downarrow_{\{c,d,e\}}}$ yields an odd cycle where none of the remaining arguments c, d, e can be accepted. However, as b is defeated, the attack $(\{b, d\}, c)$ is counter-attacked and thus, c is defended. Hence we would expect c to be acceptable in the restriction w.r.t. a.

One might argue that this notion of projection is therefore too credulous, i.e., attacks survive that should be discarded. Recall Example 5.1 where we defined the alternative projection

$$SF{\downarrow}^*_S = (S, \{(T,h) \mid (T,h) \in R,\ T \cup \{h\} \subseteq S\}).$$

Now, one can check that we get the expected results in SF'*. However,* $SF{\downarrow}^*_{\{b,c\}}$ *only features the attack* (c,b)*, which incorrectly suggests that we cannot accept* b*.*

We solve this problem by adapting the notion of a *restriction* such that both cases are handled appropriately. We keep track of a set of rejected arguments and discard attacks once an argument in its tail is discarded—these attacks are irrelevant to the further evaluation of the SETAF. This leads to the following notion,

Definition 6.6 (Restriction)**.** *Let* $SF = (A, R)$ *be a SETAF and let* $S, D \subseteq A$*. We define the* restriction *of* SF *w.r.t.* S *and* D *as the SETAF* $SF{\Downarrow}^D_S = (S', R')$ *where*

$$S' = (A \cap S) \setminus D$$
$$R' = \{(T \cap S', h) \mid (T,h) \in R,\ h \in S',\ T \cap D = \emptyset,\ T \cap S' \neq \emptyset\}.$$

Let us work through the conditions:

- D will be $D_{SF}(S, E)$ later on, i.e. the set of defeated arguments; thus $S =' A \cap S \setminus D$ is the set of non-defeated arguments in the current SCC S.
- The set R' of attacks reduces the tail T of a given attack to the set S' of consideration, i.e. (T,h) is reduced to $(T \cap S', h)$, but only if:
  - the attacked argument h belongs to S',
  - none of the arguments in the tail are defeated, $T \cap D = \emptyset$, and
  - at least one argument in the tail belongs to the current set S', $T \cap S' \neq \emptyset$.

Example 6.7. *The restriction handles both cases of Example 6.5 according to our intuition.*

- *The SETAF* $SF{\Downarrow}^{\{a\}}_{\{b,c\}}$ *contains* b *and* c*, and as we accepted* a*, i.e. the part tail of* $(\{a,b\},c)$ *outside* $\{b,c\}$*, the attack* (b,c) *is kept.*
- *The restriction* $SF'{\Downarrow}^{\{b\}}_{\{c,d,e\}}$ *contains the attacks* (c,e)*,* (e,d)*, and* (e,e)*; as* $b \in D = \{b\}$ *the tail of* $(\{b,d\},c)$ *is already defeated and we therefore do not include* (d,c)*.*

We want to emphasize that this example illustrates how the notion of projection is akin to the SETAF-reduct: indeed, constructing $SF{\Downarrow}^D_S$ consists in projecting to a certain set of arguments and then i) removing attacks where defeated arguments are involved as well as ii) partially evaluating the remaining tails. Formally, the connection is as follows.

Lemma 6.8. *Let* $SF = (A,R)$ *be a SETAF and let* $E, S \subseteq A$*. Then* $SF{\Downarrow}^{(E\setminus S)^+}_S = SF^{(E\setminus S)}{\downarrow}_S$*.*

Let us now formally introduce SCC-recursiveness (Baroni et al., 2005) as a SETAF principle. Extensions satisfying this property can be recursively characterized as follows: if the SETAF SF consists of a single SCC, the *base function* $\mathcal{BF}$ of the semantics yields the extensions. For SETAFs that consist of more SCCs, we apply the *generic selection function* $\mathcal{GF}$, where SF is evaluated separately on each SCC by means of our *restriction*, taking into account arguments that are defeated by previous SCCs.

Principle 6.9 (SCC-recursiveness)**.** *A semantics σ satisfies* SCC-recursiveness *if for all SETAFs $SF = (A, R)$, it holds that $\sigma(SF) = \mathcal{GF}(SF)$, where $\mathcal{GF}(SF) \subseteq 2^A$ is defined as follows: $E \subseteq A \in \mathcal{GF}(SF)$ if and only if*

- *if $|SCCs(SF)| = 1$, then $E \in \mathcal{BF}(SF)$;*
- *otherwise, $\forall S \in SCCs(SF)$ it holds that $E \cap S \in \mathcal{GF}\left(SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E)}\right)$,*

where $\mathcal{BF}$ is a function that maps a SETAF $SF = (A, R)$ with $|SCCs(SF)| = 1$ to a subset of 2^A.

In the following subsections we will investigate and refine SCC-recursiveness for the different semantics under our consideration. For the proofs we loosely follow the structure of (Baroni et al., 2005), incorporating our SETAF-specific notions.

6.1 Stable Semantics

We start with stable semantics, as this is the easiest case.

Example 6.10. *Recall Example 6.2.*

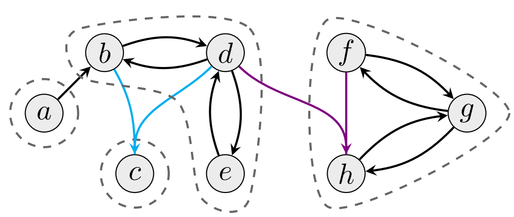

We use the base function

$$\mathcal{BF}(SF) = stb(SF).$$

Consider the stable extension $E = \{a, c, d, f\}$ of SF. Let $S = \{b, d, e\}$. The projected SETAF $SF{\Downarrow}^{(E\setminus S)^+}_S$ is given as

$$SF{\Downarrow}^{(E\setminus S)^+}_S = SF{\Downarrow}^{\{b\}}_S = (\{d, e\}, \{(d, e), (e, d)\}).$$

This projected SETAF consists of one SCC only, and we apply the base case of $\mathcal{GF}$, i.e.

$$\mathcal{GF}\left(SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E)}\right) = \mathcal{BF}\left(SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E)}\right) = stb\left(SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E)}\right).$$

Since $E \cap S = \{d\}$ is indeed a stable extension of $SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E)}$, the required condition w.r.t. the SCC S is met.

In this section we will show that this is no coincidence, i.e. *stb* satisfies SCC-recursiveness (with the base function *stb*). For the investigation of SCC-recursiveness in stable semantics we use the fact that there are no undecided arguments. Thus, in each step we do not have to keep track of as much information from previous SCCs. Formally, we obtain the following auxiliary lemma.

Lemma 6.11. *Let SF be a SETAF and $E \in stb(SF)$, then for all $S \in SCCs(SF)$ it holds $P_{SF}(S,E)=\emptyset$.*

We continue with the main technical underpinning for the SCC-recursive characterization of stable semantics. Intuitively, Proposition 6.12 states that an extension E is "globally" stable in SF if and only if for each of its SCCs S it is "locally" stable in $SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E)}$.

Proposition 6.12. *Let $SF = (A, R)$ be a SETAF and let $E \subseteq A$, then $E \in stb(SF)$ if and only if $\forall S \in SCCs(SF)$ it holds $(E \cap S) \in stb\left(SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E)}\right)$.*

This leads us to the characterization of stable extensions. As the base function is $stb(SF)$, the base case is immediate. The composite case follows from Proposition 6.12.

Theorem 6.13. *Stable semantics is SCC-recursive.*

6.2 Admissible Sets

As already mentioned, when investigating stable semantics we can use the observation that each argument is either in E or defeated by E, i.e. stable extensions correspond to two-valued models of the AF. For admissibility-based semantics, we might also have "undecided" arguments, i.e. arguments which are not in the range $E^\oplus$ of the given extension. These arguments make handling the SCC-recursive procedure more involved.

To this end we add a second component to $\mathcal{GF}$ which intuitively collects all arguments C that can still be defended within the current SCC S. On the other hand, arguments which occur in the restriction $SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E)}$ but not in C cannot be accepted anymore; however, we have to defend our extension against them. We account for this in Definition 6.22 by maintaining a set of candidate arguments C.

While this is all similar in spirit to the AF case, there is however another crucial observation we make. That is, the particular case of SETAFs gives rise to a novel scenario, where certain attacks are present in an SCC, but not applicable.

Example 6.14. *Recall our SETAF from before.*

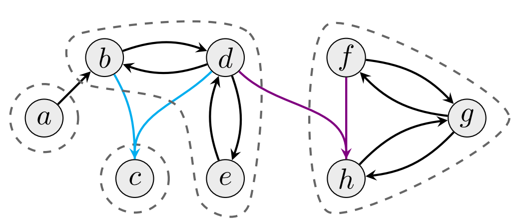

This time, consider $S = \{f, g, h\}$ *with given extension* $E = \emptyset$. *Then*

$$\begin{aligned} &D_{SF}(S,E) = \emptyset && SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E)} = SF{\Downarrow}^{\emptyset}_{S} = (A', R') \\ &P_{SF}(S,E) = \emptyset && A' = \{f, g, h\} \\ &U_{SF}(S,E) = \{f, g, h\} && R' = \{(f,h), (g,h), (h,g), (f,g), (g,f)\} \end{aligned}$$

We now observe that although there is an attack from f *to* h *in* $SF{\Downarrow}^{\emptyset}_{S}$, *the argument* h *can actually not be defeated by* f, *because this would require* d *to be present in our extension. Note however that we cannot delete the attack* (f, h), *as this would mean we could accept* h*—without defending* h *against the attack from* $\{d, f\}$.

Consequently, we will keep track of these attacks that have to be considered for defense, but cannot themselves be used to defeat an argument. We will call these attacks *mitigated*.

Definition 6.15 (Mitigated Attacks). *Let* $SF = (A, R)$ *be a SETAF. Moreover, let* $E \subseteq A$ *and* $S \in SCCs(SF)$. *The set* $M_{SF}(S, E)$ *of* mitigated attacks *is given as*

$$M_{SF}(S,E) = \{(T,h) \in R\left(SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E)}\right) \mid \forall (T',h) \in R : T' \supseteq T \Rightarrow (T' \setminus T) \not\subseteq E\}.$$

One can check that with this definition, for Example 6.14 indeed the attack towards h is identified in the SCC $S = \{f, g, h\}$, and the resulting attack (f, h) in $SF{\Downarrow}^{\emptyset}_{S}$ is mitigated (in particular, we have $M_{SF}(S, \emptyset) = \{(f, h)\}$). The intuition behind the condition

$$\forall (T',h) \in R : T' \supseteq T \Rightarrow (T' \setminus T) \not\subseteq E$$

is that (T, h) might stem from some modified attack (T', h) in the SETAF with $T \subseteq T'$: the attack (T', h) is suitably modified when computing the restriction $SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E)}$ and yields (T, h). Then, $T' \setminus T \not\subseteq E$ ensures that the attack is not active in E, independent of the choice of arguments within the SCC S. Intuitively, this accounts for a scenario where an attack (T, h) that appears in a sub-framework generated from projecting to an SCC has two or more possible origins: at least one attack $(T', h) \in R$ with $T' \supseteq T$ where some argument $t \in T' \setminus T$ is not in $E^{\oplus}$ (i.e., causing the resulting (T, h) to be mitigated), and at least one attack $(T'', h) \in R$ with $T'' \supseteq T$ where $T'' \setminus T \subseteq E$ (i.e., causing the resulting (T, h) to be non-mitigated). In this case the non-mitigated interpretation "overrides" the mitigated interpretation, as this attack can clearly be used to defend other arguments.

To account for the novel scenarios arising from the context of mitigated attacks we adapt the notion of acceptance. We have to assure that the "counter-attacks" used for defense are not mitigated. Recall that in addition our generic selection function also stores some set C of acceptable arguments, with the consequences mentioned above. Putting all of this together yields the following notions:

- If M is the set of mitigated attacks, then some extension E defends $a \in A$ if for each arbitrary attacker $(T, a) \in R$ there is some non-mitigated counter-attack $(X, t) \in R \setminus M$ with $X \subseteq E$ and $t \in T$, i.e. E counters the attack without relying on any mitigated attack;
- Each extension E must be a subset of the set C of acceptable arguments.

Formally, we obtain the following semantics considering C, M.

Definition 6.16 (Semantics Considering C, M)**.** *Let $SF = (A, R)$ be a SETAF, and let $E, C \subseteq A$ and $M \subseteq R$. We say that*

- *E is* conflict-free in C considering M*, denoted by $E \in cf(SF, C, M)$, if $E \subseteq C$ and there is no $(T, h) \in R \setminus M$ s.t. $T \cup \{h\} \subseteq E$;*
- *an argument $a \in A \setminus C$ is* acceptable considering M w.r.t. E *if for all $(T, a) \in R$ there is $(X, t) \in R \setminus M$ s.t. $X \subseteq E$ and $t \in T$;*
- *E is* admissible in C considering M*, denoted by $E \in adm(SF, C, M)$, if $E \subseteq C$, $E \in cf(SF, C, M)$, and each $a \in E$ is acceptable considering M w.r.t. E;*
- *E is* complete in C considering M*, denoted by $E \in com(SF, C, M)$, if it holds $E \in adm(SF, C, M)$ and E contains all $a \in C$ acceptable considering M w.r.t. E;*
- *E is* preferred in C considering M*, denoted by $E \in pref(SF, C, M)$, if it holds $E \in adm(SF, C, M)$ and there is no $E' \in adm(SF, C, M)$ with $E \subsetneq E'$.*

The characteristic function $F^M_{SF,C}$ of SF in C considering M *is the mapping $F^M_{SF,C} : 2^C \to 2^C$ where $F^M_{SF,C}(E) = \{a \in C \mid a \text{ is acceptable considering } M \text{ w.r.t. } E\}$.*

- *E is* grounded in C considering M*, denoted by $E \in grd(SF, C, M)$, if E is the least fixed point of $F^M_{SF,C}$.*

Setting $C = A$ and $M = \emptyset$ recovers the original semantics, in these cases we will omit writing the respective parameter. Let us discuss some properties of the semantics in C considering M. First, if we deal with admissibility-based semantics, we can actually restrict our attention to the usual notion of conflict-freeness.

Proposition 6.17. *Let $SF = (A, R)$ be a SETAF, and let $E, C \subseteq A$ and $M \subseteq R$. Let $E \subseteq C$ be a set of arguments s.t. each $a \in E$ is acceptable considering M w.r.t. E. Then $E \in cf(SF, C, M)$ if and only if $E \in cf(SF)$.*

Since we restrict our attention to admissibility-based semantics, we will for ease of notation in the following assume that $E \in cf(SF)$ instead of $E \in cf(SF, C, M)$. Next we establish that important basic properties of the characteristic function also hold in this generalized setting.

Theorem 6.18. *Let SF be a SETAF, and let $C \subseteq A$ and $M \subseteq R$. Then,*

1. *$F^M_{SF,C}$ is monotonic,*
2. *the fundamental lemma holds, i.e. if $E \in adm(SF, C, M)$ and $a \in A \cap C$ is acceptable w.r.t. E considering M, then $E \cup \{a\} \in adm(SF, C, M)$,*
3. *$E \in grd(SF, C, M)$ is the least set in $com(SF, C, M)$ w.r.t. $\subseteq$, and*
4. *$E \in pref(SF, C, M)$ are the maximal sets in $com(SF, C, M)$ w.r.t. $\subseteq$.*

To adequately characterize the defeated, provisionally defeated, and undefeated arguments in this setting we now also have to consider mitigated attacks. We illustrate this using the following example.

Example 6.19. *Consider the following SETAF SF (the dashed lines indicate the SCCs).*

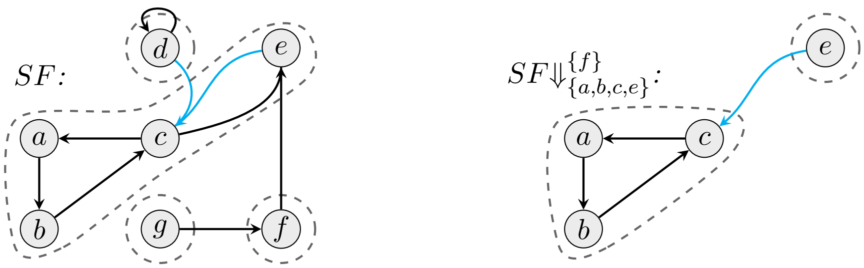

The set $\{a, e, g\}$ is not admissible, and should therefore not be characterized by our (yet to be formally defined) notion of SCC-recursiveness. Intuitively, the singleton SCCs $\{g\}, \{f\}, \{d\}$ are unsurprisingly evaluated w.r.t. $\{a, e, g\}$ in the sense that g is accepted, f is defeated, and d is undecided. To characterize the remaining SCC $\{a, b, c, e\}$ we have to take the defeated arguments into account (namely, f), resulting in $SF\Downarrow^{\{f\}}_{\{a,b,c,e\}}$. As f is defeated, we delete the attack towards e and as a result "split" the SCC into two SCCs $\{a, b, c\}, \{e\}$. It is important to see that the remaining attack (e, c) is mitigated, but in contrast to the situation illustrated in Example 6.14 the mitigated attack did not origin in the current recursion step—because we split the original SCC, we will invoke the general function of the SCC-recursive scheme on the sub-framework $SF\Downarrow^{\{f\}}_{\{a,b,c,e\}}$. Consequently, the attack (e, c) is not indicated as mitigated by the set $M_{SF}(E, S)$; to still have this relevant information we generalize this set to also take the mitigated attacks from earlier recursion steps into account. Otherwise, the attack (e, c) is not marked as mitigated. Moreover, for the same reason we have to take the mitigated attacks into account when we calculate the set of defeated arguments: e is not sufficient to defeat c—if we would not account for this "inherited" mitigated attack (e, c), we would conclude that c is defeated and therefore a is acceptable, mistakenly characterizing $\{a, e, g\}$ as admissible.

Formally, we capture this in the following slightly adapted version of Definition 6.3. Note that the only difference to the former definition is that arguments that are only attacked by E via mitigated attacks do not count as defeated, but provisionally defeated.

Definition 6.20. *Let $SF = (A, R)$ be a SETAF. Moreover, let $E \subseteq A$ be a set of arguments, $M \subseteq R$ a set of attacks, and $S \in SCCs(SF)$ be an SCC. We define the set of defeated arguments $D_{SF}(S, E, M)$, provisionally defeated arguments $P_{SF}(S, E, M)$, and undefeated arguments $U_{SF}(S, E, M)$ w.r.t. S, E, M as*

$$D_{SF}(S, E, M) = \{a \in S \mid \exists(T, a) \in R \setminus M \text{ s.t. } T \subseteq E \setminus S\},$$
$$P_{SF}(S, E, M) = \{a \in S \mid A \setminus (S \cup E^+) \mapsto_R a\} \setminus D_{SF}(S, E, M),$$
$$U_{SF}(S, E, M) = S \setminus (D_{SF}(S, E, M) \cup P_{SF}(S, E, M)).$$

Moreover, we set $UP_{SF}(S, E, M) = U_{SF}(S, E, M) \cup P_{SF}(S, E, M)$.

We have to make similar adjustments to the notion of mitigated attacks. Due to Definition 6.15, for the computation of mitigated attacks only the ancestor SCCs are relevant. In particular, the set $(T' \setminus T)$ is contained in ancestor SCCs of S for each attack $(T', h) \in R$. However, when we apply the concept of mitigated attacks to characterize SCC-recursiveness in admissibility-based semantics, we will face situations where we already know for the original SETAF that some attacks are mitigated. To account for this set of given mitigated attacks M, we slightly modify the condition for mitigated attacks, s.t. only the non-mitigated attacks $(T', h) \in R$ can "override" the status of a mitigated attack as non-mitigated.

Definition 6.21 (Mitigated Attacks, refined)**.** *Let $SF = (A, R)$ be a SETAF. Moreover, let $E \subseteq A$ and $S \in SCCs(SF)$. The set $M_{SF}(S, E, M)$ of* mitigated attacks *is given as*

$$M_{SF}(S, E, M) = \{(T, h) \in R\left(SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}\right) \mid \forall (T', h) \in R\setminus M : T' \supseteq T \Rightarrow (T' \setminus T) \not\subseteq E\}.$$

Indeed, for Example 6.19 we get that the attack (e, c) is mitigated in $SF\Downarrow^{\{f\}}_{\{a,b,c,e\}}$. Next we redefine Definition 6.9 in order to capture the admissibility-based semantics. For this, we need to take into account that in each recursive call of the generic selection function $\mathcal{GF}$ we will also have to pass the current set M of mitigated attacks.

Principle 6.22 (SCC-recursiveness, refined)**.** *A semantics σ satisfies* SCC-recursiveness[3] *if and only if for all SETAFs $SF = (A, R)$ it holds $\sigma(SF) = \mathcal{GF}(SF, A, \emptyset)$, where the* generic selection function *$\mathcal{GF}(SF, C, M) \subseteq 2^A$ is defined as: $E \subseteq A \in \mathcal{GF}(SF, C, M)$ if and only if*

- *if $|SCCs(SF)| = 1$, then $E \in \mathcal{BF}(SF, C, M)$,*
- *otherwise, $\forall S \in SCCs(SF)$ it holds that*

$$E \cap S \in \mathcal{GF}\left(SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S, E, M) \cap C, M_{SF}(S, E, M)\right),$$

where $\mathcal{BF}$ maps $SF = (A, R)$ with $|SCCs(SF)| = 1$ and sets $C \subseteq A$, $M \subseteq R$ to a subset of 2^A.

Towards an SCC-recursive characterization of admissible sets we discuss the following auxiliary results. Lemma 6.23 shows that global acceptability implies local acceptability, Lemma 6.24 shows the converse direction.

Lemma 6.23. *Let $SF = (A, R)$ be a SETAF, let $M \subseteq R$, $C \subseteq A$, and let $E \in adm(SF, C, M)$ be an admissible set of arguments and let $a \in A \cap C$ be acceptable w.r.t. E considering M in SF, where $a \in S$ for some SCC S. Then*

1. *we have $a \in U_{SF}(S, E, M)$ and a is acceptable w.r.t. $(E \cap S)$ in $SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}$ considering $M_{SF}(S, E, M)$;*
2. *it holds that $(E \cap S)$ is conflict-free in $SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}$.*

3. Note that in the conference version of this work (Dvořák et al., 2022) we erroneously do not account for the mitigated attacks that are already passed down from previous recursion steps. To fix this issue we introduced the respective third parameters in $D_{SF}(S, E, M)$, $U_{SF}(S, E, M)$, $P_{SF}(S, E, M)$, $UP_{SF}(S, E, M)$, $M_{SF}(S, E, M)$ keeping track of exactly these attacks. As a consequence, we adapted Definition 6.21 and Definition 6.20 to account for the same issue. The situation is illustrated in Example 6.19.

Lemma 6.24. *Let $SF = (A, R)$ be a SETAF, let $M \subseteq R$, let $E \subseteq A$ such that*

$$(E \cap S) \in adm\left(SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S, E, M), M_{SF}(S, E, M)\right)$$

for all $S \in SCCs(SF)$. Moreover, let $S' \in SCCs(SF)$ and let $a \in U_{SF}(S', E, M)$ be acceptable w.r.t. $(E \cap S')$ in $SF\Downarrow^{(E\setminus S')^+}_{UP_{SF}(S',E,M)}$ considering $M_{SF}(S', E, M)$. Then a is acceptable w.r.t. E in SF considering M.

Combining these two results we obtain the SCC-recursive characterization of admissible sets.

Proposition 6.25. *Let $SF = (A, R)$ be a SETAF and let $E \subseteq A$ be a set of arguments. Then for each $C \subseteq A$ and $M \subseteq R$ it holds $E \in adm(SF, C, M)$ if and only if $\forall S \in SCCs(SF)$ it holds $(E \cap S) \in adm(SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S, E, M) \cap C, M_{SF}(S, E, M))$.*

The base function for admissible sets is $adm(SF, C, M)$. We will utilize this result to obtain the characterizations of the other (admissibility-based) semantics.

Theorem 6.26. *Admissible semantics is SCC-recursive.*

Proof. The base function $\mathcal{BF}(SF)$ is $adm(SF, C, M)$. The case for $|SCCs(SF)| = 1$ is immediate, the composite case follows from Proposition 6.25. □

Example 6.27. *Recall our example with the (not admissible) set $E = \{a, e, g\}$.*

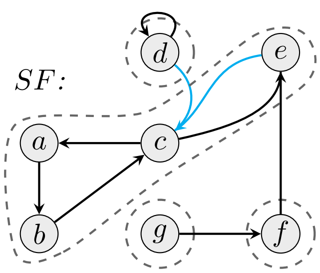

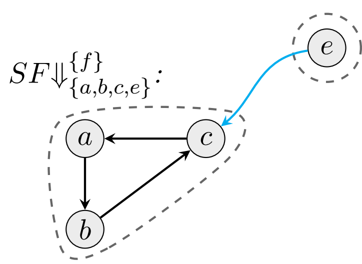

Indeed, in $SF\Downarrow^{\{f\}}_{\{a,b,c,e\}}$ we have that (e, c) is a mitigated attack. Therefore, in the next recursive step, in the SCC $\{a, b, c\}$ the argument c is detected as provisionally defeated. Hence a is not admissible in the corresponding sub-framework and thus, E is rightfully detected as non-admissible.

Let us now consider SF' the same SETAF as SF, but without the self-attacker d.

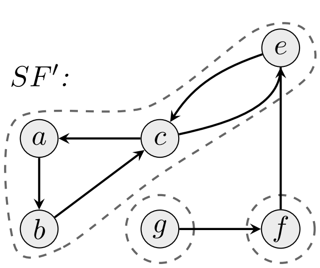

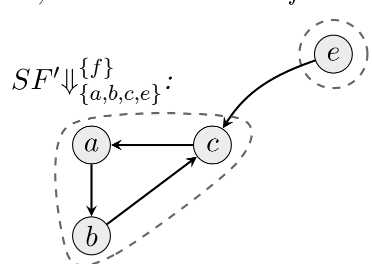

Let E be as above. Then, the attack (e, c) is not mitigated in $SF'\Downarrow^{\{f\}}_{\{a,b,c,e\}}$ and hence, c is detected as defeated. Hence in order to evaluate $\{a, b, c\}$ we require another recursive step and find acceptance of $\{a\}$ in an sub-SCC consisting only of the single undefeated argument a. Therefore, it is rightfully detected that $E \in adm(SF')$.

6.3 Complete Semantics

We already have the tools to characterize complete extensions: Proposition 6.25 proves the desired properties for admissible sets, in addition we can apply Lemma 6.23 and Lemma 6.24 to show that complete extensions contain all arguments they defend (i.e., for an SCC S', an extension E, and a set of mitigated attacks M, exactly those arguments from $U_{SF}(S', E, M)$ that are *acceptable* w.r.t. $(E \cap S')$ in $SF{\Downarrow}^{(E \setminus S')^+}_{UP_{SF}(S',E,M)}$ considering $M_{SF}(S', E, M)$).

Proposition 6.28. *Let $SF = (A, R)$ be a SETAF, let $M \subseteq R$, and let $E \subseteq A$ be a set of arguments. Then $\forall C \subseteq A$ it holds $E \in com(SF, C, M)$ if and only if $\forall S \in SCCs(SF)$ it holds $(E \cap S) \in com(SF{\Downarrow}^{(E \setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S, E, M) \cap C, M_{SF}(S, E, M))$.*

From this we get the desired result regarding complete extensions. The base function is $com(SF, C, M)$.

Theorem 6.29. *Complete semantics is SCC-recursive.*

6.4 Preferred Semantics

The next lemma illustrates that if we already found a globally admissible set E and find a (larger) locally admissible set $E' \supsetneq E \cap S$ in an SCC S, then we can find a globally admissible set incorporating this set E'. This idea underlies the incremental computation of extensions (see Section 7).

Lemma 6.30. *Let $SF = (A, R)$, let $M \subseteq R$, and let $E \in adm(SF, A, M)$, let $S \in SCCs(SF)$ be an SCC. Moreover, let $E' \subseteq A$ be a set of arguments such that $(E \cap S) \subseteq E' \subseteq U_{SF}(S, E, M)$, and $E' \in adm(SF{\Downarrow}^{(E \setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S, E, M), M_{SF}(S, E, M))$. Then $E \cup E'$ is admissible in SF considering M.*

Given this lemma, we are ready to show SCC-recursiveness for preferred semantics.

Proposition 6.31. *Let $SF = (A, R)$ be a SETAF, let $M \subseteq R$ and let $E \subseteq A$ be a set of arguments. Then $\forall C \subseteq A$ it holds $E \in pref(SF, C, M)$ if and only if $\forall S \in SCCs(SF)$ it holds $(E \cap S) \in pref(SF{\Downarrow}^{(E \setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S, E, M) \cap C, M_{SF}(S, E, M))$.*

From this we get the desired result regarding preferred extensions. The base function is $pref(SF, C, M)$.

Theorem 6.32. *Preferred semantics is SCC-recursive.*

6.5 Grounded Semantics

For the characterization of grounded semantics we exploit the fact that also in our setting the grounded is the unique minimal complete extension (see Theorem 6.18). Hence, we can apply Proposition 6.28 and utilize the fact that that for the unique grounded extension minimality has to hold for each SCC to prove minimality of the whole extension.

Proposition 6.33. *Let $SF = (A, R)$ be a SETAF, $M \subseteq R$, and let $E \subseteq A$ be a set of arguments. Then $\forall C \subseteq A$ it holds $E \in grd(SF, C, M)$ if and only if $\forall S \in SCCs(SF)$ it holds $(E \cap S) \in grd(SF{\Downarrow}^{(E \setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S, E, M) \cap C, M_{SF}(S, E, M))$.*

Theorem 6.34. *Grounded semantics is SCC-recursive.*

6.6 Connection to Directionality

As it is the case in AFs, we can obtain results regarding directionality using SCC-recursiveness if the base function always admits at least one extension (Baroni & Giacomin, 2007). First note that for an uninfluenced set U any SCC S with $S \cap U \neq \emptyset$ has to be contained in U, as well as all ancestor SCCs of S. Then, by the SCC-recursive characterization we get the following general result, subsuming the semantics under our consideration.

Proposition 6.35. *Let σ be a semantics such that for all SETAFs SF and all $C \subseteq A(SF)$, $M \subseteq R(SF)$ it holds $\mathcal{BF}(SF, C, M) \neq \emptyset$. If σ satisfies SCC-recursiveness then it satisfies directionality.*

Proof. We use the fact that for an uninfluenced set U any SCC S with $S \cap U \neq \emptyset$ has to be contained in U, as well as all ancestor SCCs of S. Let $\mathcal{S}$ be the set of SCCs S with $S \subseteq U$. Considering the SCC-recursive characterization, this yields $\sigma(SF\Downarrow^{\emptyset}_{U}) = \{E \subseteq U \mid \forall S \in \mathcal{S} : (E \cap S) \in \mathcal{GF}(SF\Downarrow^{\emptyset}_{UP_{SF}(S,E,M)}, U_{SF}(S, E, M), M_{SF}(S, E, M))\}$. We have to show that $\sigma(SF\Downarrow^{\emptyset}_{U}) = \{E \cap U \mid E \in \sigma(SF)\}$.

We get the "$\subseteq$" direction from the fact that $U_{SF}(S, E, M) = U_{SF}(S, E \cap U, M)$ and $P_{SF}(S, E, M) = P_{SF}(S, E \cap U, M)$ for all $S \in \mathcal{S}$. The "$\supseteq$" direction is immediate: as we assume that $\mathcal{BF}(SF, C, M)$ always yields at least one extension, we can extend any set $(E \cap U)$ according to the SCC-recursive scheme (see Section 7 for details). □

7. Incremental Computation & Computational Graph Fragments

In this section we discuss the computational implications of a semantics satisfying directionality, modularization, or SCC-recursiveness, and how we can improve the asymptotic runtime of the resulting algorithms by utilizing structures in the graph of the SETAFs. To this end, we first briefly recall the basic notions of complexity analysis in the context of abstract argumentation (Section 7.1). We then establish the basic idea of our algorithms exploiting structural properties (Section 7.2) and define and analyze graph classes to further refine the relevant structures (Sections 7.3-7.6), Finally, we show how the presence of these properties leads to computational ease (Section 7.7) and generalize this result to be applicable in the general case in the context of SCCs (Section 7.8).

7.1 Complexity of Reasoning in Abstract Argumentation

We assume the reader to have basic knowledge in computational complexity theory[4], in particular we make use of the complexity classes L (logarithmic space), P (polynomial time), NP (non-deterministic polynomial time), coNP, DP ($L_1 \cap L_2$ for $L_1 \in \mathsf{NP}$, $L_2 \in \mathsf{coNP}$), Θ_2^{P} ($\mathsf{P}^{\mathsf{NP}[log(n)]}$), Σ_2^{P} ($\mathsf{NP}^{\mathsf{NP}}$), and Π_2^{P} ($\mathsf{coNP}^{\mathsf{NP}}$). We have the following relationships between these classes (an arrow from class $\mathcal{C}$ to $\mathcal{C}'$ means $\mathcal{C} \subseteq \mathcal{C}'$, we omit some arrows that are immediate due to transitivity):

4. For a gentle introduction to complexity theory in the context of formal argumentation, see (Dvořák & Dunne, 2017).

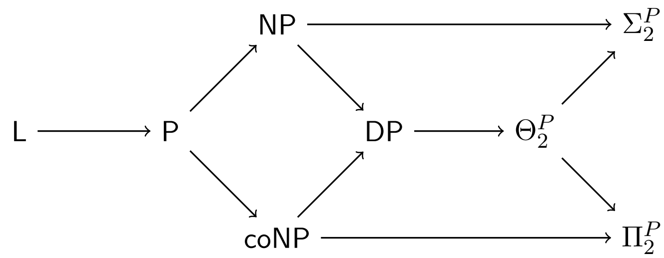

For a given SETAF $SF = (A, R)$ and an argument $a \in A$, we consider the decision problems (under semantics σ) in formal argumentation:

- Credulous acceptance $Cred_\sigma$: Given $SF = (A, R)$ and $a \in A$, is it true that $a \in E$ for some $E \in \sigma(SF)$?
- Skeptical acceptance $Skept_\sigma$: Given $SF = (A, R)$ and $a \in A$, is it true that $a \in E$ for each $E \in \sigma(SF)$?
- Verification Ver_σ: Given $SF = (A, R)$ and $E \subseteq A$, is it true that $E \in \sigma(SF)$?

The complexity landscape of SETAFs coincides with that of Dung AFs and is depicted in Table 2 ("General"). As SETAFs generalize Dung AFs the hardness results for Dung AFs (Coste-Marquis, Devred, & Marquis, 2005; Dimopoulos & Torres, 1996; Dunne & Bench-Capon, 2002; Dvořák, 2012; Dvořák & Woltran, 2010, 2011; Dunne, 2009; Caminada et al., 2012) (for a survey see (Dvořák & Dunne, 2017)) carry over to SETAFs. Also the same upper bounds hold for SETAFs (Dvořák et al., 2018).

Many of the above mentioned hardness-results are based on (variations of) the so-called *standard reduction*, see e.g. (Dvořák & Dunne, 2017, Reduction 3.6). The idea is to express the complexity of the boolean satisfiability problem in an AF. As we will base some of our complexity results for SETAFs on the standard reduction, we briefly recall the construction and some known results.

Reduction 7.1 (Standard Reduction)**.** *Let φ be a formula in CNF (conjunctive normal form) with clauses C over atoms Y. We construct the AF F_φ as follows:*

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

An example of the standard reduction can be found in Figure 1.

Some of the main results from the literature regarding the semantics under our consideration can be summarized as follows.

Theorem 7.2. *Let φ be a propositional formula in CNF and F_φ the corresponding AF from the standard reduction. The following statements are equivalent:*

1. *the formula φ is satisfiable,*
2. *the argument φ is credulously accepted in F_φ w.r.t. $\sigma \in \{adm, com, stb, pref, sem, stage\}$,*

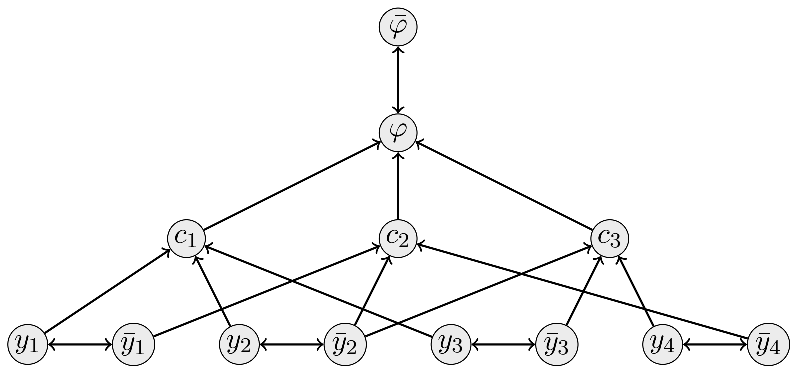

Figure 1: The standard reduction applied to φ with atoms $Y = \{y_1, y_2, y_3, y_4\}$, and clauses $C = \{\{y_1, y_2, y_3\}, \{\bar{y}_1, \bar{y}_2, \bar{y}_4)\}, \{\bar{y}_2, \bar{y}_3, y_4\}\}$.

		cf	*grd*	*adm*	*com*	*stb*	*pref*	*naive*	*sem*	*stage*	*ideal*	*eager*
General	$Cred_\sigma$	in P	P-c	NP-c	NP-c	NP-c	NP-c	in P	Σ_2^P-c	Σ_2^P-c	in Θ_2^P	Π_2^P-c
	Ver_σ	in P	P-c	in L	in L	in L	coNP-c	in P	coNP-c	coNP-c	in Θ_2^P	DP-c
	$Skept_\sigma$	in P	P-c	triv.	P-c	coNP-c	Π_2^P-c	in P	Π_2^P-c	Π_2^P-c	in Θ_2^P	Π_2^P-c
Acyclicity	$Cred_\sigma$	in P	P-c	P-c	P-c	P-c	P-c	in P	P-c	P-c	in P	in P
	Ver_σ	in P	in P	in L	in L	in L	in L	in P	in L	in L	in P	in P
	$Skept_\sigma$	in P	P-c	triv.	P-c	P-c	P-c	in P	P-c	P-c	in P	in P
Even-cycle-freeness	$Cred_\sigma$	in P	P-c	P-c	P-c	P-c	P-c	in P	P-c	Σ_2^P-c	in P	in P
	Ver_σ	in P	in P	in L	in P	in P	in P	in P	in P	in coNP	in P	in P
	$Skept_\sigma$	in P	P-c	triv.	P-c	P-c	P-c	in P	P-c	Π_2^P-c	in P	in P
Odd-cycle-freeness	$Cred_\sigma$	triv.	P-c	NP-c	NP-c	NP-c	NP-c	in P	NP-c	NP-c	coNP-c	coNP-c
	Ver_σ	in P	in P	in L	in P	in P	in P	in P	in P	in P	coNP-c	coNP-c
	$Skept_\sigma$	in P	P-c	triv.	P-c	coNP-c	coNP-c	in P	coNP-c	coNP-c	coNP-c	coNP-c
Self-attack-free full-symmetry	$Cred_\sigma$	triv.	in L	triv.	triv.	triv.	triv.	triv.	triv.	triv.	in P	in P
	Ver_σ	in P	in P	in L	in L	in P	in P	in P	in P	in P	in P	in P
	$Skept_\sigma$	in P	in L	triv.	in L	in L	in L	in L	in L	in L	in P	in P
Primal-bipartiteness	$Cred_\sigma$	triv.	P-c	P-c	P-c	P-c	P-c	in P	P-c	P-c	in Θ_2^P	in Π_2^P
	Ver_σ	in P	in P	in L	in L	in L	in L	in P	in L	in L	in Θ_2^P	in DP
	$Skept_\sigma$	in P	P-c	triv.	P-c	P-c	P-c	in P	P-c	P-c	in Θ_2^P	in Π_2^P

Table 2: Graph fragments in SETAFs. $\mathcal{C}$-c denotes completeness for class $\mathcal{C}$; "triv." denotes a trivial problem (either all instances are positive or all instances are negative).

3. *the argument $\bar{\varphi}$ is not skeptically accepted in F_φ w.r.t. $\sigma \in \{stb, pref, sem, stage\}$,*

4. *the set $\{\bar{\varphi}\}$ is not the unique ideal/eager extension, and*

5. *the argument $\bar{\varphi}$ is not in the unique ideal/eager extension.*

Note that some of the results reported in Table 2 stem from (Dvořák et al., 2021a). However, we added the analysis of the verification problem as well as several semantics. For the remainder of the paper we will omit proof details and rather give general ideas on how the respective complexity results can be obtained. Proof details for the theorems that are not immediate from these explanations can be found in Appendix D.

7.2 Basic Computational Speedup

First, for a semantics σ satisfying directionality an argument a is in some extension (in all extensions) if and only if it is in some extension (in all extensions) of the framework that is restricted to the arguments that influence a. That is, when deciding credulous or skeptical acceptance of an argument, in a preprocessing step, we can shrink the framework to the relevant part. The property of modularization is closely related to CEGAR style algorithms for preferred semantics that can be implemented via iterative SAT-solving (Dvořák, Järvisalo, Wallner, & Woltran, 2014). In order to compute a preferred extension we can iteratively compute a non-empty admissible set of the current framework, build the reduct w.r.t. this admissible set, and repeat this procedure on the reduct until the empty set is the only admissible set. The preferred extension is then given by the union of the admissible sets.

Finally, for SCC-recursive semantics we can iteratively compute extensions along the SCCs of a given framework (see (Baumann, 2011; Liao, Jin, & Koons, 2011; Baroni, Giacomin, & Liao, 2014; Cerutti, Giacomin, Vallati, & Zanella, 2014) for such approaches for AFs). It is well known that the SCCs of any directed graph form a partial order w.r.t. reachability: in Example 7.3 (b) the SCC S_1 is an initial SCC and precedes S_2 and S_3, and S_2 precedes S_3. In (one of) the initial SCCs we simply compute the extensions and then for each of these extensions we proceed on the preceding SCCs. We then iteratively continue this process on SCCs in their order. To evaluate an SCC that is attacked by other ones we have to take the attacks from earlier SCCs into account and, as we have already fixed our extension there, we can simply follow the SCC-recursive schema. We next illustrate this for stable semantics.

Example 7.3. *Consider (a) the SETAF SF and (b) the order of its SCCs.*

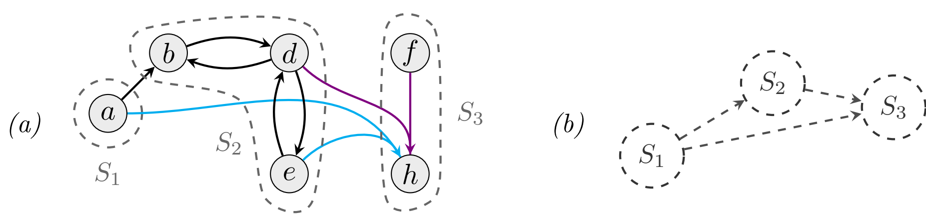

We can iteratively compute the stable extensions of SF as follows: in the first SCC $S_1 = \{a\}$ we simple compute all the stable extensions, i.e., stb($SF{\Downarrow}^{\emptyset}_{S_1}$) $= \{\{a\}\}$. We then proceed with $\{a\}$ as extension E for the part of the SETAF considered so far. Next we consider S_2 and adapt it to take E into account. As $(E \setminus S_2)^+ = \{b\}$ we only have to delete the argument b from S_2 before evaluating the SCC and thus we obtain $SF{\Downarrow}^{\{b\}}_{S_2} = (\{d, e\}, \{(d, e), (e, d)\})$. Combining these with E we obtain two stable extensions $E_1 = \{a, d\}$, $E_2 = \{a, e\}$ for $SF{\Downarrow}^{\emptyset}_{S_1 \cup S_2}$. We proceed with S_3 and first consider E_1. As $(E_1 \setminus S_3)^+ = \{b, e\}$ we do not remove arguments from S_3. However, as $d \in E_1$ we cannot delete the attack $(\{d, f\}, h)$ but have to replace it by the attack (f, h). We then have stb($SF{\Downarrow}^{\{b,e\}}_{S_3}$) $= \{\{f\}\}$ and thus obtain the first stable extensions of SF $\{a, d, f\}$. Now consider E_2. We have that E_2 attacks h, i.e., $(E_2 \setminus S_3)^+ = \{b, d, h\}$, and thus we have to remove h before evaluating S_3 and thus obtain $SF{\Downarrow}^{\{b,d,h\}}_{S_3} = (\{f\}, \emptyset)$. We end up with $\{a, e, f\}$ as the second stable extension of SF.

The computational advantage of the incremental computation is that certain computations are performed over single SCCs instead of the whole framework. This is in particular significant for preferred semantics where the $\subseteq$-maximality check can be done within the SCCs. Notice that verifying a preferred extension is in general coNP-complete (Dimopoulos & Torres, 1996; Dvořák et al., 2018). However, given our results regarding the SCC-recursive scheme, the following parameterized tractability result is easy to obtain: it is well known that computing the SCCs of a directed graph can be done efficiently. It then suffices to verify a given extension along the SCCs of the framework, whereby we only need to consider one SCC at the time.

Theorem 7.4. *Let SF be a SETAF where $|S| \leq k$ for all $S \in SCCs(SF)$. Then we can verify a given preferred extensions in $O\big(2^k \cdot poly(|SF|)\big)$ for some polynomial poly.*

Next, we will define and analyze graph classes for SETAFs and illustrate how we can exploit them to reason more efficiently. To this end we generalize situations that are known to yield computationally easy fragments in the special case of AFs, and provide positive as well as negative results to illustrate the border cases of tractability. Moreover we will show how we can utilize these findings in the context of SCC-recursiveness to achieve a computational speedup even if the framework is heterogeneous, i.e., does not as a whole belong to one of these tractable fragments. Instead, in the spirit of Theorem 7.4 we follow the SCC-recursive scheme and pose the respective restrictions only on the strongly connected components, resulting in a more flexible setting.

7.3 Towards Graph Classes for SETAFs

The directed hypergraph-structure of SETAFs is rather specific and to the best of our knowledge the hypergraph literature does not provide generalizations of common graph classes to this kind of directed hypergraphs. Thus we first identify suitable 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. Next, 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.*

7.4 Acyclicity, Even- and Odd-Cycle-Freeness

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. Cycles in SETAFs in the context of restricted length and their effect on the computational complexity have been investigated in (Dvořák et al., 2021b).

Definition 7.5. *A* cycle C *of length* $|C| = n$ *is a sequence of pairwise distinct arguments* $C = (a_1, a_2, \dots, a_n, a_1)$ *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 called* acyclic *if it contains no cycle (otherwise it is called* cyclic*),* 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 (Nielsen & Parsons, 2006), i.e. there is no infinite sequence of sets $B_1, B_2, \dots$, such that for all i, the argument B_i is the tail of an attack towards an argument in B_{i-1}. As shown in (Nielsen & Parsons, 2006), 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. Finally, this means that the grounded extension also coincides with the ideal and eager extension in this case. Together with the hardness from AFs, we immediately obtain our first result from Table 2.

Theorem 7.6. *For acyclic SETAFs the problems* $Cred_\sigma$ *and* $Skept_\sigma$ *for* $\sigma \in \{grd, com, pref, stb, stage, sem, ideal, eager\}$ *are* P*-complete. Moreover* $Cred_{adm}$ *is* P*-complete. Finally,* Ver_σ *is in* P *for all semantics* σ *under our consideration.*

For AFs we have that the class of all no-even graphs 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 (Dvořák, 2012, Proposition 15). This property also holds for SETAFs, which in turn means even-cycle-free SETAFs admit the grounded extension as their unique complete set, which is then also the only preferred and semi-stable extension (and, hence, ideal and eager). Our proof of this property follows along the lines of the respective known proof for AFs (see Appendix D).

Proposition 7.7. *Let* $SF = (A, R)$ *be a SETAF. If* $|com(SF)| \geq 2$ *then* SF *contains an even-cycle.*

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. (Dvořák & Dunne, 2017)) to SETAFs. This yields the next line of results in Table 2.

Theorem 7.8. *For even-cycle-free SETAFs the problems* $Cred_\sigma$ *and* $Skept_\sigma$ *for* $\sigma \in \{com, pref, stb, sem, ideal, eager\}$ *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. Finally, the problems* Ver_σ *for* $\sigma \in \{com, pref, stb, sem, ideal, eager\}$ *are in* P*.*

For odd-cycle free SETAFs the situation is just like with odd-cycle-free AFs (Dunne & Bench-Capon, 2002). If there is a sequence of arguments $(a_1, a_2, \dots)$, we say a_1 *indirectly attacks* the arguments a_{2i-1} and *indirectly defends* the arguments a_{2i} for $i \geq 1$ (cf. (Nielsen

& Parsons, 2006)). As odd-cycle-free SETAFs are *limited controversial* (Nielsen & Parsons, 2006), 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. As a consequence, ideal and eager semantics coincide—for both semantics the reasoning and verification problems drop to coNP: hardness can be shown with the standard reduction (already in the special case of AFs, see Theorem 7.2); membership is clear, as for the complementary problem it suffices to guess a preferred extension that attacks a given argument—because all preferred extensions are stable, we can verify such an extension in polynomial time.

Theorem 7.9. *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. Moreover, the problems $Cred_\sigma$, $Skept_\sigma$, Ver_σ for $\sigma \in \{ideal, eager\}$ are* coNP*-complete.*

7.5 Symmetry

In the following we provide two natural generalizations of the notion of symmetry[5] for SETAFs. The first definition by means of 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 Example 7.12.

Definition 7.10. *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 $h \in H$.*

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 (Diller, Keshavarzi Zafarghandi, Linsbichler, & Woltran, 2020). 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 full symmetry we distinguish between self-attacks and attacks which are not self-attacks.

Definition 7.11. *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 illustrate the two symmetry notions with the following example.

Example 7.12. *Consider the following SETAFs. The framework in (a) is primal-symmetric (but not fully-symmetric), the framework in (b) is fully-symmetric. It is easy to see from the respective definitions that every fully-symmetric SETAF is primal symmetric.*

5. Further symmetry-notions for SETAFs have been investigated in (König, 2020).

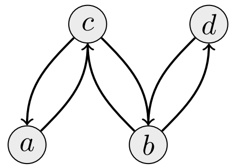

(a)

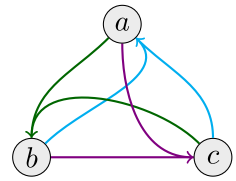

(b)

For usual AFs, symmetry is a rather strong notion: Every argument defends itself against all incoming attacks and hence, admissible sets coincide with conflict-free sets, and it becomes computationally easy to reason with admissible, complete, and preferred extensions. However, this is not the case for our notions of symmetry for SETAFs. Consider the fully-symmetric (and thus also primal-symmetric) SETAF from Example 7.12(b): 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 under our consideration, when making no further restrictions on the graph structure.

For both notions of symmetry we have that an argument is in the grounded extension if and only if it is not in the head of any attack, which be can easily checked in logarithmic space. This is by the characterization of the grounded extension as least fixed point of the *characteristic function* (Nielsen & Parsons, 2006), i.e. the grounded extension can be computed by starting from the empty set and iteratively adding all defended arguments (cf. our characterization via the reduct in Section 4.2). 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 of reasoning in 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 $(t, a^1_{r,t}), (t, a^2_{r,t})$, as illustrated in the following example.

Example 7.13. *Consider the attack from t_1, t_2 to h (illustrated in violet color). We introduce the four additional arguments and the attacks as discussed. Observe that the cyan attacks are never efficient, as it is impossible to accept the tail in a conflict-free semantics. Hence, it is easy to see that this operation does not change the stable, preferred, semi-stable, and stage extensions.*

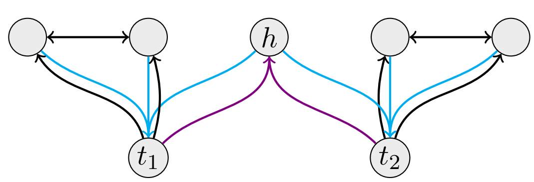

It can be verified that the resulting SETAF SF' is primal-symmetric, does not introduce self-attacks and preserves the acceptance status of the original arguments (shown in Appendix D).

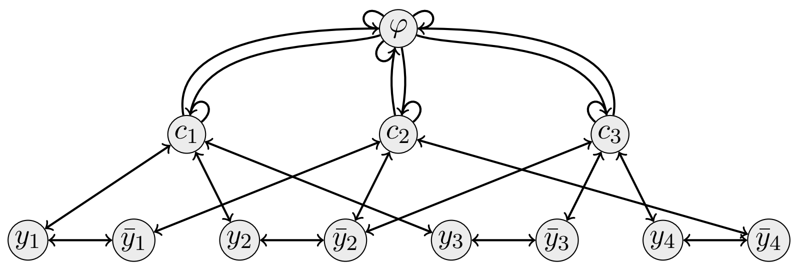

Figure 2: Reduction 7.15 applied to φ with atoms $Y = \{y_1, y_2, y_3, y_4\}$, and clauses $C = \{\{y_1, y_2, y_3\}, \{\bar{y}_1, \bar{y}_2, \bar{y}_4)\}, \{\bar{y}_2, \bar{y}_3, y_4\}\}$.

Theorem 7.14. *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 (Dvořák & Dunne, 2017). For reasoning with admissible, complete and preferred semantics, hardness can be shown with adjustments to the standard reductions (cf. Reduction 7.1). 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 $(\{a, b\}, a), (\{a, b\}, b)$.

Reduction 7.15. *Let φ be a CNF-formula with clauses C over atoms Y, we define the fully-symmetric SETAF $SF_\varphi = (A, R)$ as follows:*

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

An example of this construction can be found in Figure 2.

Proposition 7.16. *For fully-symmetric SETAFs the problems $Cred_\sigma$ for $\sigma \in \{adm, stb, pref, com\}$ are* NP*-complete.*

It is no coincidence that our reduction features self-attacks, as we will find out that fully-symmetric SETAFs without self-attacks indeed allow us to reason in polynomial time. We utilize the fact that an attack $(\{a, b\}, b)$ cannot be used to defeat argument b.

Note that we can extend Reduction 7.15 to show that verifying preferred extensions remains hard in this context by simply adding (a) a self-attacking argument $\bar{\varphi}$, (b) symmetric attacks $(\bar{\varphi}, \varphi), (\varphi, \bar{\varphi})$, and (c) symmetric attacks $(\{x, \bar{\varphi}\}, x), (\{x, \bar{\varphi}\}, \bar{\varphi})$ for each $x \in Y \cup \bar{Y}$. The argument $\bar{\varphi}$ has to be defeated by any non-empty admissible set, and the only way to do so is to accept φ. The empty set $\emptyset$ is preferred if and only if the original formula φ is unsatisfiable. As all semi-stable extensions are preferred, coNP-hardness carries over to semi-stable semantics.

If instead of the attacks $(\{x, \bar{\varphi}\}, x), (\{x, \bar{\varphi}\}, \bar{\varphi})$ we add "classical" symmetric attacks $(\bar{\varphi}, x), (\bar{\varphi}, x)$ for each $x \in Y \cup \bar{Y}$, it is easy to see that the set $\{\bar{\varphi}\}$ is admissible. Moreover, by the same reasoning as before, $\{\bar{\varphi}\}$ is the only preferred extension if and only if φ is unsatisfiable. If φ is satisfiable, then there are other preferred (semi-stable) extensions containing the argument φ. Hence, $\{\bar{\varphi}\}$ is ideal/eager if and only if φ is unsatisfiable, proving coNP-hardness in these cases.

We can summarize the complexity of reasoning and verification in fully-symmetric SETAFs as follows (see also Table 2). Note that intractability for symmetric AFs allowing self-attacks has already been observed in (Dvořák, 2012).

Theorem 7.17. *For fully-symmetric SETAFs (allowing self-attacks) the problems $Cred_{grd}$, $Skept_{grd}$, Ver_{grd}, and $Skept_{com}$ are in* L*, the complexity of $Cred_\sigma$ and $Skept_\sigma$ for $\sigma \in \{adm, com, stb, pref, naive, sem, stage\}$ coincides with the complexity for the general problems. The problem Ver_σ for $\sigma \in \{pref, sem\}$ is* coNP*-complete. Moreover, Ver_σ for $\sigma \in \{ideal, eager\}$ is* coNP*-hard.*

Investigations on symmetric AFs often distinguish between frameworks with and without self-attacks (Dvořák & Dunne, 2017). As in self-attack-free symmetric AFs, for *self-attack-free* fully-symmetric SETAFs we have that all naive extensions 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: if we assume the SETAF is *redundancy-free*, i.e., for all attacks (T, h) it holds there is no stronger attack (T', h) with $T' \subset T$, then for self-attack free fully-symmetric SETAFs it holds that every tail T of any attack (T, h) is conflict-free. Hence in this case an argument is skeptically accepted if and only if it is not attacked. If we loosen this restriction and allow for redundant attacks, the following algorithm decides whether an argument a is skeptically accepted: let $(T_1, a), \dots, (T_n, a)$ be all attacks towards a. Then a is skeptically accepted if and only if no set T_i is conflict-free: if some set T_k is conflict-free we can extend it to a naive extension that attacks a. If not, there is no way to attack a, and hence, a has to be in every stable extension.

Finally, the grounded extension coincides with the ideal and eager extension.

Theorem 7.18. *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 $\sigma \in \{grd, com, pref, stb, stage, sem, ideal, eager\}$. Moreover, Ver_σ and $Cred_\sigma$ for $\sigma \in \{grd, ideal, eager\}$ is in* L*.*

7.6 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

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 $Y_i = Y_{i-1}$;

primal-bipartiteness makes credulous and skeptical reasoning P-easy. For an illustration of the respective definitions see Example 7.23.

Definition 7.19. *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), s.t.*

- *$Y \cup Z = A$, $Y \cap Z = \emptyset$, 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 (Dunne, 2007); 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 (Nielsen & Parsons, 2006), 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.

Lemma 7.20. *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.*

Note that primal-bipartite SETAFs are odd-cycle-free and therefore coherent. These results suffice to pin down the complexity of credulous and skeptical reasoning for the semantics under our consideration.

Theorem 7.21. *For primal-bipartite SETAFs the problems $Cred_\sigma$ and $Skept_\sigma$ for $\sigma \in \{com, pref, stb, stage, sem\}$ are P-complete. Moreover the problems $Cred_{adm}$ and Ver_{pref} are 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 (Dunne, 2007); 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.

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 7.22. *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, i.e. (a, a) with a single argument in the tail, but 2-colorable SETAFs may contain self-attacks (T, h) with $|T| \geq 2$, $h \in T$.

Example 7.23. *Consider the following SETAFs (a) SF_1 and (b) SF_2. SF_1 is primal-bipartite with partition illustrations of (a) primal-bipartiteness (the primal graph is bipartite with partition $(\{y_1, y_2\}, \{z_1, z_2\})$, and, hence, also 2-colorable. SF_2 is also 2-colorable (assign different colors to y and z, the only edge $(\{y, z\}, z)$ is not monochromatic). However, SF_2 is not primal-bipartite, as z has an edge to itself in the primal graph.*

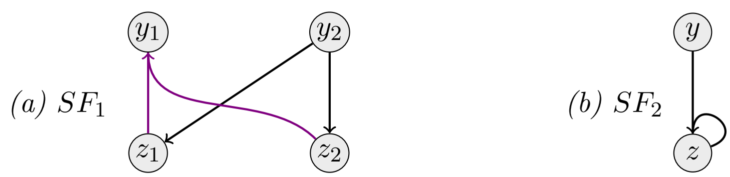

For admissibility-based semantics that preserve the grounded extension (such as *grd*, *com*, *pref*, *stb*, *sem*) it is easy to see that the problems remain hard in 2-colorable SETAFs: intuitively, one can add two fresh arguments to any SETAF and add them to the tail T of every attack (T, h) - they will be in each extension of the semantics in question, and other than that the extensions will coincide with the original SETAF (as this is an instance of the *attack-weakening principle*, see Principle 3.39). 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 (see Appendix D for details).

Theorem 7.24. *For 2-colorable SETAFs the complexity of $Cred_\sigma$ and $Skept_\sigma$ for $\sigma \in \{grd$, adm, com, stb, $pref$, $naive$, sem, $stage\}$ coincides with the complexity of the general problem.*

7.7 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 (Dvořák, Ordyniak, & Szeider, 2012), the respective AF-classes can be efficiently decided by graph algorithms. As for acyclicity, even-cycle-freeness, and primal-bipartiteness it suffices 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 7.25. *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 2.

7.8 Utilizing Tractable Fragments for Efficient Computation Along SCCs

In this section we will show which of our tractable fragments we can exploit in the context of SCC-recursiveness to speed up computation. In particular, we exemplify the speedup with the Ver_{pref} problem. We will show for acyclicity, even-cycle-freeness, and primal-bipartiteness, that we obtain a speedup when every SCC of a SETAF belongs to one of these tractable fragments.

On the other hand, we show that for full-symmetry this is not the case. Deleting (parts of) attacks from a fully-symmetric SETAF might lead to a situation where the remaining framework is no longer fully symmetric. We will show that therefore this fragment does not allow a speedup in the SCC-recursive scheme. The key idea is that prior SCCs can "disable" arbitrary attacks in a given SCC. In the reduction from the general verification problem we use to prove this (illustrated in Example 7.27) we have at least 3 SCCs: the one containing x, the one containing y, and the ones containing our original framework (if the original framework is not connected, we obtain more than one). Each SCC is fully symmetric, but the symmetric counter-attacks in the SCCs corresponding to the original framework are irrelevant, as the argument y in the tail is always defeated.

Proposition 7.26. *The problem Ver_{pref} remains* coNP*-complete even for self-attack-free SETAFs SF where all SCCs $S \in SCCs(SF)$ are fully-symmetric, i.e., $SF{\downarrow_S}$ is fully symmetric.*

Example 7.27. *Consider the SETAF in (a) with preferred extensions $\{\{a, c\}, \{b\}\}$. By adding the arguments x, y we make every SCC primal-symmetric, while preserving the preferred extensions under projection (i.e., the added arguments and attacks have no practical effect, in particular, if we construct the reduct w.r.t. $\{x\}$ we recover the original framework): (b) has preferred extensions $\{\{x, a, c\}, \{x, b\}\}$.*

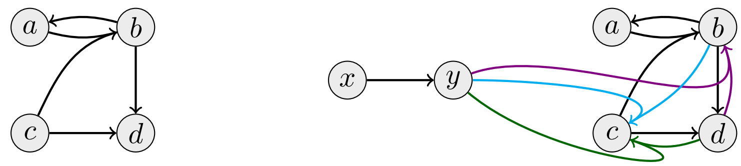

(a) (b)

In contrast to this negative result, it is indeed possible to verify preferred extensions efficiently in primal-bipartite SCCs. To establish this result, we generalize our notion of the reduct and modularization to the semantics considering a candidate set $C \subseteq A$ and a set of mitigated attacks $M \subseteq R$ (see SCC-recursiveness, Definition 6.22).

Definition 7.28. *Given a SETAF $SF = (A, R)$, $M \subseteq R$, and $E \subseteq A$, the* E-reduct of SF considering M *is the SETAF* $SF^E_M = (A', R')$, *with*

$$\begin{aligned} A' &= A \setminus E^{\oplus}_{R\setminus M} \\ R' &= \{(T \setminus E, h) \mid (T, h) \in R,\, T \cap E^{+}_{R\setminus M} = \emptyset,\, T \not\subseteq E,\, h \in A'\} \end{aligned}$$

Note the parallels to the definition of the restricted frameworks in the SCC-recursive scheme. We now show that the modularization property also holds in this context. The idea is similar to the special case of $C = A, M = \emptyset$ that we discussed in Theorem 4.6. What we have left to consider is that an admissible set E could attack an argument x in its reduct via mitigated attacks. As in the SCC-recursive scheme, we cannot accept x in the reduct. Hence, we add such arguments to the set C.

Proposition 7.29. *Let $SF = (A, R)$ be a SETAF and $C \subseteq A$, $M \subseteq R$, and $E \in adm(SF, C, M)$. Let $SF' = SF^E_M = (A', R')$.*

1. *If $E' \in adm(SF', C', M')$ with*

$$\begin{aligned} C' &= C \cap \{a \in A \mid \nexists (T, a) \in R \cap M : (T \subseteq E)\} \\ M' &= \{(T, h) \in R' \mid \forall (T', h) \in R \setminus M : T' \supseteq T \Rightarrow (T' \setminus T) \not\subseteq E\} \end{aligned}$$

   then $E \cup E' \in adm(SF, C, M)$.

2. *If $E \cap E' \neq \emptyset$ and $E \cup E' \in adm(SF, C, M)$, then $E' \in adm(SF', C', M')$.*

Proof. 1. Since E is admissible in SF considering M, E' does not attack E via non-mitigated attacks in SF. By construction of SF', E does not attack E' via non-mitigated attacks either. Since $E' \subseteq C'$ we know E also does not attack E' via mitigated attacks. Towards contradiction assume E' attacks E via a mitigated attack. By admissibility, then E attacks E' in SF which is not the case (as just established). Hence, $E \cup E' \in cf(SF)$.

Now assume $S \mapsto_R E \cup E'$. If $S \mapsto_R E$, then $S \mapsto_{R\setminus M}$ by admissibility. If $S \mapsto_R E'$, then there is $T \subseteq S$ s.t. $(T, e') \in R$ with $e' \in E'$. If now $E \mapsto_{R\setminus M} T$ we are done. Otherwise, there is $(T \setminus E, e') \in R'$, and $E' \mapsto_{R'\setminus M'} T \setminus E$. This means there is $(X', t) \in R' \setminus M'$ with $t \in T$ and $X' \subseteq E'$, and consequently $(X, t) \in R \setminus M$ with $X \setminus E = X'$. Since $X \mapsto_{R\setminus M} S$ and $E \cup E' \supseteq X$, it also holds $E \cup E' \mapsto_{R\setminus M} S$, i.e., $E \cup E'$ defends itself against S in SF considering M. Hence, we have $E \cup E' \in adm(SF, C, M)$.

Algorithm 2: 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), sets $C \subseteq A$, $M \subseteq R$

Output: The admissible set Y_i (considering M) of credulously accepted arguments in $Y \cap C$

```
1 i := 0
2 Y_0 := Y ∩ C
3 R_0 := R
4 repeat
5 |  i := i + 1
6 |  Y_i := Y_{i-1} \ {y | y ∈ Y_{i-1}, there is some (Z', y) ∈ R_{i-1} with Z' ⊆ Z such that
  |     ∀z ∈ Z' |{(Y', z) | (Y', z) ∈ R_{i-1} \ M}| = 0}
7 |  R_i := R_{i-1} \ {(Y', z) | Y' ⊆ Y, z ∈ Z, Y' ⊈ Y_i}
8 until Y_i = Y_{i-1};
```

2. Assume $E \cup E' \in adm(SF, C, M)$. We see $E' \in cf(SF')$ as follows: if $(T', e') \in R'$ with $T' \subseteq E'$, $e' \in E'$, then there is some $(T, e') \in R$ with $T' = T \setminus E$. Hence, $E \cup E' \mapsto_R E'$, a contradiction.

Now assume $E' \notin adm(SF', C', M')$. This means there is $(T', e') \in R'$ with $e' \in E'$, but there is no non-mitigated counter-attack from E' towards T'. Then there is $(T, e') \in R$ with $T' = T \setminus E$ and $T \cap E^+_{R \setminus M} = \emptyset$. By admissibility we know $E \cup E' \mapsto_{R \setminus M} T$, say $(T^*, t) \in R \setminus M$ with $T^* \subseteq E \cup E'$, $t \in T$. Since $E \cup E'$ is conflict-free in SF, $T^* \cap E^+_R = \emptyset$, and thus we either have a) $T^* \subseteq E$, contradicting $T \cap E^+_{R \setminus M} = \emptyset$, or b) $(T^* \setminus E, t) \in R' \setminus M'$, contradicting the assumption that there is no counter-attack. Finally, note that $E' \subseteq C'$, as otherwise there is $(T, e') \in R \cap M$ with $T \subseteq E, e' \in E'$, contradicting conflict-freeness. In summary, we conclude $E' \in adm(SF', C', M')$. □

It remains to show that we can find non-empty admissible sets in this context in polynomial time. To this end, we slightly adapt Algorithm 1 to also account for a candidate set C and mitigated attacks M, see Algorithm 2. The only differences to the original algorithm are in step 2, where we consider C, and in step 6, where for possible counter-attacks from our constructed admissible set to attackers we only take non-admissible attacks into account. The proof of the correctness and completeness is analogous to Lemma 7.20.

Proposition 7.30. *Let $SF = (A, R)$ be a primal-bipartite SETAF and $C, E \subseteq A$, $M \subseteq R$. We can decide whether $E \in pref(SF, C, M)$ in polynomial time.*

Finally, we can see that this result generalizes to odd-cycle-freeness if we restrict ourselves to the case where we only have a single SCC.

Proposition 7.31. *Let $SF = (A, R)$ be an odd-cycle-free SETAF with $|SCCs(SF)| = 1$. Then SF is primal-bipartite.*

Clearly, this implies that if $|SCCs(SF)| = 1$ in an odd-cycle-free SETAF we can also verify preferred extensions in polynomial time. Note also that by removing (parts of) attacks from SF we cannot introduce an odd-cycle.

In the following we argue that we can efficiently compute the (unique) preferred extension if the SETAF in question is even-cycle-free. We utilize the fact that an even-cycle-free SETAF has only one preferred extension, namely the grounded—this also holds true considering the candidate set C and mitigated attack M.

Proposition 7.32 (cf. (Dvořák, 2012))**.** *Let $SF = (A, R)$ be an even-cycle-free SETAF and $C, E \subseteq A$, $M \subseteq R$. We can decide whether $E \in pref(SF, C, M)$ in polynomial time.*

Taking these results together, we obtain the following characterization. This generalizes the FPT-result from Theorem 7.4 and illustrates that we can utilize various different graph properties of SCCs at once.

Theorem 7.33. *Let SF be a SETAF where for all SCCs $S \in SCCs(SF)$ it holds either*

- *S is acyclic,*
- *S is even-cycle-free,*
- *S is primal-bipartite,*
- *S odd-cycle-free, or*
- *the size of S is bounded by a parameter k, i.e., $|S| \leq k$.*

Then we can verify a given preferred extensions in $O\big(2^k \cdot poly(|SF|)\big)$ for some polynomial function poly.

8. Conclusion

In this work, we systematically analyzed semantics for SETAFs using a principle-based approach (see Table 1 for an overview of the investigated properties). Moreover, we introduced and investigated novel principles that are trivial on AFs, but desirable non-trivial properties of SETAF semantics. We pointed out interesting concepts that help us to understand the principles more deeply: edge cases that for AFs are hidden behind simple syntactic notions have to be considered explicitly for SETAFs, revealing semantic peculiarities that are already there in the special case. To this end, we highlight the usefulness of the *reduct* in this context—many seemingly unrelated notions from various concepts boil down to formalizations closely related to the reduct. We particularly focused on computational properties like modularization and SCC-recursiveness. The emphasis on the computation of argumentation tasks lead us to our investigations of graph properties in the context of SCCs, during which we introduced and analyzed the computational complexity of reasoning tasks for these restricted cases. Finally, we applied these findings in the context of SCC-recursiveness, which allowed us to push the boundaries of tractability for argumentation tasks even further.

The latter concept has recently been investigated for Abstract Dialectical Frameworks (ADFs) in a different context by Gaggl, Rudolph, and Straß (2021). In that work, the acceptance conditions of the statements in an ADFs (that encode the attacks in case the ADF recasts a SETAF) are modified. Note that SETAFs can be seen as a special case of ADFs, where each acceptance condition is a formula in conjunctive normal form with only negative literals. The modification of this formula in the approach of Gaggl et al. is indeed closely related to our idea of the SETAF-reduct: attacks (T, h) where in a prior SCC we learn that one argument $t \in T$ is attacked (defeated) in an extension effectively become redundant and are removed. In case an argument $t \in T$ is in an extension, the acceptance condition is modified, in SETAF terms this would correspond to an attack $(T \setminus \{t\}, h)$. However, while we treat the undecided state of an argument that is neither in nor attacked by an extension via mitigated attacks, in the approach of Gaggl et al. self-attacks are introduced to model the resulting effects (akin to the idea of splitting (Baumann, 2011; Linsbichler, 2014)). While both approaches effectively yield the same results, we expect the introduction of new (self-)attacks to be computationally disadvantageous compared to our approach of labeling an attack as mitigated, as additional computational effort is to be expected.

Our findings regarding computational speedups utilizing SCC-structures are in line with similar recent considerations for SETAFs such as backdoors, treewidth, and cycle length (Dvořák et al., 2022b, 2022a, 2021b). The first two—backdoors and treewidth—were recently considered in combination to further improve the runtime (Dvořák, Hecher, König, Schidler, Szeider, & Woltran, 2022); in the future it will be interesting to investigate whether these results generalize to SETAFs. Moreover, the reduct for SETAFs and the generalization of the recursive scheme for SETAFs allow for generalizations of several abstract argumentation semantics that have not yet been studied in the context of collective attacks (Baumann et al., 2020b). An interesting direction for future works is investigating semantics *cf2* (Baroni et al., 2005) and *stage2* (Dvořák & Gaggl, 2016), as well as the family of semantics based on weak admissibility (Baumann et al., 2020b). Moreover, checking whether the results carry over to other decision problems is also an interesting direction for future research. In particular our results in Section 7 show potential computational benefits from the notions we establish for SETAFs. A natural task for future work would be to utilize these findings in an implementation and compare the performance to state-of-the-art solvers.

Acknowledgments

This research has been supported by the Vienna Science and Technology Fund (WWTF) through project ICT19-065, the Austrian Science Fund (FWF) through project P32830, and by the Federal Ministry of Education and Research of Germany and by Sächsische Staatsministerium für Wissenschaft, Kultur und Tourismus in the programme Center of Excellence for AI-research "Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig", project identification number: ScaDS.AI.

Appendix A. Proof Details of Section 4

Lemma 4.4. *Given a SETAF $SF = (A, R)$ and two disjoint sets $E, E' \subseteq A$. Let $SF^E = (A', R')$.*

1. *If there is no $S \subseteq A$ s.t. $S \mapsto_R E'$, then the same is true in SF^E.*

2. *Assume E does not attack $E' \in cf(SF)$. Then, E defends E' in SF iff there is no $S' \subseteq A'$ s.t. $S' \mapsto_{R'} E'$.*

3. *Let $E \in cf(SF)$. If $E \cup E'$ does not attack E in SF and $E' \subseteq A'$, with $E' \in cf\left(SF^E\right)$ then $E \cup E' \in cf(SF)$.*

4. *Let $E \cup E' \in cf(SF)$. If $E' \mapsto_{R'} a$, then $E \cup E' \mapsto_R a$.*

5. *If $E \cup E' \in cf(SF)$, then $SF^{E \cup E'} = \left(SF^E\right)^{E'}$.*

Proof.

1. This is clear since SF^E contains (strictly) less attacks than SF.

2. ($\Rightarrow$) Assume E defends E' in SF. Now suppose there is some attacker in the reduct, i.e. $S' \mapsto_{R'} e'$. By definition, there is some $T' \subseteq A \setminus E^{\oplus}$ with $(T', e') \in R'$. By the definition of SF^E, $T' = T \cap A'$ for some $(T, e') \in R$. Now consider an arbitrary $(T, e') \in R$. Since E defends E', $E \mapsto_R T$. Again by definition of SF^E, (T, e') is removed since $T \cap E_R^+ \neq \emptyset$. Hence in R' there is no attack of the form (T', e') with $T' \subseteq T$, contradiction.

   ($\Leftarrow$) Now suppose E does not defend E'. There is thus some $(T, e') \in R$ and E does not attack T, i.e. $T \cap E_R^+ = \emptyset$. Suppose $T \setminus E = \emptyset$. Then $T \subseteq E$ contradicting that E does not attack E'. Thus, $T \setminus E \neq \emptyset$. Since E does not attack E' and $E \cap E' = \emptyset$, we have $e' \in A'$ for each $e' \in E'$. Therefore, in R' we find the attack (T', e') with $T' = T \cap A' \neq \emptyset$, $e' \in A'$, and $T \cap E_R^+ = \emptyset$.

3. We have to show that $E \cup E'$ does not attack E'. Suppose the contrary, i.e. let $T \subseteq E \cup E'$ with $(T, e') \in R$ for some $e' \in E'$. Since E does not attack E or E', $T \cap E_R^+ = \emptyset$. The case $T \subseteq E$ is impossible. Thus, (T, e') induces some attack $(T \setminus E, e')$ in SF^E. We infer $T \setminus E \subseteq E'$ implying $E' \notin cf\left(SF^E\right)$, contradiction.

4. If $E' \mapsto_{R'} a$, then $(T', a) \in R'$ for some $T' \subseteq E'$. Hence $(T, a) \in R$ for some T with $T \setminus E = T'$. The claim follows due to $T \subseteq T' \cup E \subseteq E' \cup E$.

5. ($\subseteq$) Let $a \in SF^{E \cup E'}$. Then $a \notin E \cup E'$ and $E \cup E'$ does not attack a. We infer $a \in A\left(SF^E\right)$. Now if $E' \mapsto_{R'} a$, then $E \cup E' \mapsto_R a$ by item 4. Thus $a \in A\big((SF^E)^{E'}\big)$.

   ($\supseteq$) Let $a \in A\big((SF^E)^{E'}\big)$. Hence $a \notin E \cup E'$ and E' does not attack a in SF^E. Assume $(T, a) \in R$ with $T \subseteq E \cup E'$. Since E' does not attack a in SF^E, there cannot be an attack of the form $(T \setminus E, a) \in R'$ satisfying $T \setminus E \neq \emptyset$ and $T \cap E_R^+ = \emptyset$. However, T satisfies $T \cap E_R^+ = \emptyset$ since $E \cup E' \in cf(SF)$. We thus infer $T \setminus E = \emptyset$. This yields $E \mapsto_R a$ contradicting $a \in A\big((SF^E)^{E'}\big)$. □

Proposition 4.7. *Let $SF = (A, R)$ be a SETAF, $E \in cf(SF)$ and $SF^E = (A', R')$.*

1. *$E \in stb(SF)$ iff $SF^E = (\emptyset, \emptyset)$,*

2. *$E \in adm(SF)$ iff $S \mapsto_R E$ implies $S \setminus E \not\subseteq A'$,*

3. *$E \in pref(SF)$ iff $E \in adm(SF)$ and $adm(SF^E) = \{\emptyset\}$,*

4. *$E \in com(SF)$ iff $E \in adm(SF)$ and $grd(SF^E) = \{\emptyset\}$,*

5. *$E \in sem(SF)$ iff $E \in pref(SF)$ and there is no $E' \in pref(SF)$ s.t. $A(SF^{E'}) \subsetneq A(SF^E)$.*

Proof. The characterizations for *stb* and *adm* are straightforward and *pref* is due to the modularization property of *adm*. For *com*(SF) we apply Lemma 4.4, item 2, to each singleton E' occurring in SF^E: assume towards contradiction E is complete in SF and there is some $a \in A'$ such that a is unattacked in SF' (and, hence, a is in the grounded extension of SF^E). As $a \in A'$ we know $E \not\mapsto_R \{a\}$. But then we can apply Lemma 4.4, item 2, and get that E defends $\{a\}$ in SF, contradicting $E \in com(SF)$. For *sem* recall that range-maximal preferred extensions are semi-stable. □

Appendix B. Proof Details of Section 5

Lemma 5.10. *Let $SF = (A, R)$ be a SETAF. The the unique ideal extension S satisfies*

$$S = \bigcup_{E \in adm(SF): \forall P \in pref(SF): E \subseteq P} E$$

Proof. We let

$$adm_{\subseteq pref}(SF) = \{E \in adm(SF) \mid \forall P \in pref(SF) : E \subseteq P\}.$$

We need to show that (a) S is conflict-free in SF, (b) every argument $a \in S$ is acceptable w.r.t. S, and (c) there is no larger set $S' \supset S$ that satisfies (a) and (b) and is a subset of every preferred extensions of SF. (a) is clear, because if there was a conflict caused by an attack $(T, h) \in R$ with $T \cup \{h\} \subseteq S$, this would mean two sets $E, E' \in adm_{\subseteq pref}(SF)$ attacked each other, which would mean a preferred extension is conflicting, a contradiction. (b) follows from the fact that for all $a \in S$ there is an $E \subseteq S$ with $a \in E, E \in adm(SF)$. (c) is clear since (a) and (b) characterize admissibility—if there was such a larger admissible set $S' \supset S$ with $S' \in adm_{\subseteq pref}(SF)$ by definition we would have $S' \subseteq S$, a contradiction. □

Appendix C. Proof Details of Section 6

Lemma 6.8. *Let $SF = (A, R)$ be a SETAF and let $E, S \subseteq A$. Then $SF\Downarrow_S^{(E\setminus S)^+} = SF^{(E\setminus S)}\downarrow_S$.*

Proof. We have $A(SF\Downarrow_S^{(E\setminus S)^+}) = A(SF^{(E\setminus S)}\downarrow_S)$ because

$$\begin{aligned}(A \cap S) \setminus (E \setminus S)^+ &= (A \setminus (E \setminus S)^+) \cap S \\ &= (A \setminus ((E \setminus S)^+ \cup (E \setminus S))) \cap S \\ &= (A \setminus (E \setminus S)^\oplus) \cap S.\end{aligned}$$

Then it holds that $R(SF\Downarrow_S^{(E\setminus S)^+}) = R(SF^{(E\setminus S)}\downarrow_S)$, as for some $(T, h) \in R(SF\Downarrow_S^{(E\setminus S)^+})$ with $h \in A'$ and $T \cap (E \setminus S)^+ = \emptyset$ we have $T \cap A' = \emptyset$ if and only if $T \not\subseteq (E \setminus S)$. The claim follows then from $(T \cap A', h) = (T \setminus (E \setminus S), h)$. □

Lemma 6.11. *Let SF be a SETAF and $E \in stb(SF)$, then for all $S \in SCCs(SF)$ it holds $P_{SF}(S, E) = \emptyset$.*

Proof. Assume towards contradiction that for some SCC S there is an argument $a \in P_{SF}(S, E)$. Then, by definition there is an attack $(T, a) \in R(SF)$ such that $T \subseteq A(SF) \setminus S$ and $T \cap E^+ = \emptyset$. Moreover, $a \notin D_{SF}(S, E)$ by definition, i.e. $T \not\subseteq E$. But then there is some $t \in T$ such that neither $t \in E^+$ nor $t \in E$, which is a contradiction to the assumption that E is stable. □

Proposition 6.12. *Let $SF = (A, R)$ be a SETAF and let $E \subseteq A$, then $E \in stb(SF)$ if and only if $\forall S \in SCCs(SF)$ it holds $(E \cap S) \in stb\left(SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E)}\right)$.*

Proof. Let $SF' = SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E)}$ for an arbitrary SCC S. We start by assuming $E \in stb(SF)$. We need to show that $(E \cap S) \in stb(SF')$, i.e.:

1. $(E \cap S) \subseteq UP_{SF}(S, E)$,
2. $(E \cap S)$ is conflict-free in SF', and
3. $\forall a \in UP_{SF}(S, E)$ if $a \notin (E \cap S)$ then $(E \cap S)$ attacks a in SF'.

For condition 1. note that $(E \cap X) \cap D_{SF}(X, E) = \emptyset$ holds for any $X \subseteq A$, as otherwise E would not be conflict-free in SF. For condition 2., assume towards contradiction that there is some $(T, h) \in R(SF')$ such that $T \cup \{h\} \subseteq (E \cap S)$. This means there is some $(T', h) \in R$ with $T' \supseteq T$. But by construction we would have $T' \setminus T \subseteq E$, and therefore $T' \cup \{h\} \subseteq E$, a contradiction to conflict-freeness of E. For condition 3. we consider an arbitrary argument $a \in UP_{SF}(S, E) \setminus (E \cap S)$. Since $a \notin E$ and E is stable, there is an attack $(T, a) \in R$ with $T \subseteq E$. Moreover, as $a \in UP_{SF}(S, E)$, it holds $a \notin D_{SF}(S, E)$, i.e. in particular $T \not\subseteq (E \setminus S)$, or in other words $T \cap S \neq \emptyset$. This means by the definition of the restriction and since $T \cap E^+_R = \emptyset$ (otherwise E would not be conflict-free in SF) there is an attack $(T \cap S, a) \in R(SF')$ with $(T \cap S) \subseteq E$.

Now assume $\forall S \in SCCs(SF)$ it holds $(E \cap S) \in stb(SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E)})$. We need to show $E \in stb(SF)$, i.e.

1. E is conflict-free in SF, and
2. E attacks all $a \in A \setminus E$ in SF.

For 1. assume towards contradiction there is some $(T, a) \in R$ such that $T \cup \{a\} \subseteq E$. Let S be the SCC containing a. Clearly $T \cup \{a\} \not\subseteq S$, as this violates our assumed conflict-freeness in $UP_{SF}(S, E)$. Moreover, we do not have $T \subseteq (A \setminus S)$, as this would mean $a \in D_{SF}(S, E)$. Hence, there is an attack $(T \cap S, a) \in R\left(SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E)}\right)$ such that $(T \cap S) \cup \{a\} \subseteq E \cap S$, a contradiction. For condition 2. let us consider an arbitrary argument $a \in A \setminus E$ and let S be the SCC containing a. Then either (i) $a \in D_{SF}(S, E)$ or (ii) $a \in UP_{SF}(S, E)$. In case (i) we immediately get E attacks a. For case (ii), we have $a \notin (E \cap S)$, and by assumption a is attacked in S, i.e. there is an attack $(T, a) \in R\left(SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E)}\right)$. By construction of the restriction, this means there is an attack $(T', a) \in R$ s.t. $T' \supseteq T$ and $T' \setminus T \subseteq E$. Hence, $T \subseteq E$, i.e. E attacks a in SF. □

Proposition 6.17. *Let $SF = (A, R)$ be a SETAF, and let $E, C \subseteq A$ and $M \subseteq R$. Let $E \subseteq C$ be a set of arguments s.t. each $a \in E$ is acceptable considering M w.r.t. E. Then $E \in cf(SF, C, M)$ if and only if $E \in cf(SF)$.*

Proof. The ($\Leftarrow$) direction is clear since $E \in cf(SF)$ is a stricter notion. For ($\Rightarrow$) suppose $E \in cf(SF, C, M)$. We have to show that even for mitigated attacks $(T, h) \in M$ it holds that $T \cup \{h\} \nsubseteq E$. Striving for a contradicting suppose otherwise. Then we have in particular that $h \in E$. Since h is acceptable w.r.t. E by assumption, there is some non-mitigated attack $(X, t) \in R \setminus M$ with $X \subseteq E$ and $t \in T$. Since $T \subseteq E$, it follows $t \in E$. Hence, the attack (X, t) causes a conflict (not making use of mitigated attacks), contradiction. □

Theorem 6.18. *Let SF be a SETAF, and let $C \subseteq A$ and $M \subseteq R$. Then,*

1. *$F^M_{SF,C}$ is monotonic,*
2. *the fundamental lemma holds, i.e. if $E \in adm(SF, C, M)$ and $a \in A \cap C$ is acceptable w.r.t. E considering M, then $E \cup \{a\} \in adm(SF, C, M)$,*
3. *$E \in grd(SF, C, M)$ is the least set in $com(SF, C, M)$ w.r.t. $\subseteq$, and*
4. *$E \in pref(SF, C, M)$ are the maximal sets in $com(SF, C, M)$ w.r.t. $\subseteq$.*

Proof.

1. Monotonicity of the mapping

$$F^M_{SF,C}(E) = \{a \in C \mid a \text{ is acceptable considering } M \text{ w.r.t. } E\}$$

holds by definition of defense.

2. Let $a \in A \setminus C$ be acceptable w.r.t. $E \in adm(SF, C, M)$ considering M. By monotonicity of defense, each argument in $E \cup \{a\}$ is acceptable w.r.t. $E \cup \{a\}$ considering M. As our notion of defense w.r.t. M implies the usual defense, we can apply the standard fundamental lemma for SETAFs (Nielsen & Parsons, 2006) and obtain $E \cup \{a\} \in cf(SF)$. Therefore, the conditions for applying Proposition 6.17 are met and we deduce $E \cup \{a\} \in cf(SF, C, M)$. Hence $E \cup \{a\} \in adm(SF, C, M)$ follows.

3. Setting $G = \bigcup_{i \geq 1} (F^M_{SF,C})^i(\emptyset)$ we claim that G is the least set in $com(SF, C, M)$. Due to the fundamental lemma (see 2.) admissibility of $\emptyset$ implies inductively

$$\forall n \in \mathbb{N} : \bigcup_{1 \leq i \leq n} \left(F^M_{SF,C}\right)^i(\emptyset) \in adm(SF, C, M).$$

Since SF is finite and by monotonicity, there is some n s.t.

$$\bigcup_{1 \leq i \leq n} \left(F^M_{SF,C}\right)^i(\emptyset) = \bigcup_{1 \leq i} \left(F^M_{SF,C}\right)^i(\emptyset) = G$$

Thus, G is complete. Now let $E \in com(SF, C, M)$. By monotonicity of $F^M_{SF,C}$ we get $F^M_{SF,C}(\emptyset) \subseteq F^M_{SF,C}(E)$. By induction, $(F^M_{SF,C})^i(\emptyset) \subseteq (F^M_{SF,C})^i(E)$ therefore also holds for any integer $i \geq 1$. Since E is complete, $E = (F^M_{SF,C})^i(E)$ holds for each integer i, i.e. the right-hand side is actually constant. We conclude for each n

$$G = \bigcup_{1 \leq i} \left(F^M_{SF,C}\right)^i(\emptyset) = \bigcup_{1 \leq i \leq n} \left(F^M_{SF,C}\right)^i(\emptyset) \subseteq \bigcup_{1 \leq i \leq n} \left(F^M_{SF,C}\right)^i(E) = E,$$

thus it follows that $G \subseteq E$.

4. By definition $E \in \mathit{pref}(SF, C, M)$ is maximal in $\mathit{adm}(SF, C, M)$. So we show E is maximal in $\mathit{adm}(SF, C, M)$ iff E is maximal in $\mathit{com}(SF, C, M)$.

   ($\Rightarrow$) Suppose $E \in \mathit{pref}(SF, C, M)$ is not maximal in $\mathit{com}(SF, C, M)$. Then there is a proper complete superset E' of E; since E' is in particular admissible, E is not maximal in $\mathit{adm}(SF, C, M)$.

   ($\Leftarrow$) Now suppose E is not maximal in $\mathit{adm}(SF, C, M)$. Take $E' \in \mathit{adm}(SF, C, M)$ with $E \subsetneq E'$. By the fundamental lemma and monotonicity of $F^M_{SF,C}$, we find that $\bigcup_{1 \leq i}(F^M_{SF,C})^i(E')$ is a complete proper superset of E (analogous to (ii)). Hence E is not maximal in $\mathit{com}(SF, C, M)$. □

Lemma 6.23. *Let $SF = (A, R)$ be a SETAF, let $M \subseteq R$, $C \subseteq A$, and let $E \in \mathit{adm}(SF, C, M)$ be an admissible set of arguments and let $a \in A \cap C$ be acceptable w.r.t. E considering M in SF, where $a \in S$ for some SCC S. Then*

1. *we have $a \in U_{SF}(S, E, M)$ and a is acceptable w.r.t. $(E \cap S)$ in $SF \Downarrow^{(E \setminus S)^+}_{UP_{SF}(S,E,M)}$ considering $M_{SF}(S, E, M)$;*

2. *it holds that $(E \cap S)$ is conflict-free in $SF \Downarrow^{(E \setminus S)^+}_{UP_{SF}(S,E,M)}$.*

Proof. Set $SF \Downarrow^{(E \setminus S)^+}_{UP_{SF}(S,E,M)} = SF' = (A', R')$.

1. By Theorem 6.18, item 2, we get that $E \cup \{a\} \in \mathit{adm}(SF, C, M)$, i.e. $a \notin D_{SF}(S, E, M)$ by conflict-freeness of E and $a \notin P_{SF}(S, E, M)$ by defense. Consequently, we infer $a \in U_{SF}(S, E, M)$. Likewise, we get $(E \cap S) \subseteq U_{SF}(S, E, M)$, and therefore $(E \cap S) \subseteq A'$.

   To show that a is acceptable in this context we have to consider attacks towards a, i.e. $(T, a) \in R'$, and establish $T \not\subseteq E$ (conflict-freeness) and $(E \cap A')$ attacks T in SF' via non-mitigated attacks (defense). As E is admissible in SF, there is a (non-mitigated) counter-attack $(X, t) \in R \setminus M$ with $t \in T$ and $X \subseteq E$. In particular, this means that $t \notin E$, as (X, t) would otherwise contradict the conflict-freeness of E. Hence, $T \not\subseteq E$. Moreover, because $(T, a) \in R'$, it must be that $X \cap S \neq \emptyset$, as otherwise the attack (T, a) would be deleted when we construct the appropriate restriction to the SCC S. Let $X' = X \cap S$, i.e. there is an attack $(X', t) \in R'$. In other words, $E \cap S$ defends a in SF'. Finally, $X \subseteq E$ and $(X, t) \notin M$ implies $(X', t) \notin M_{SF}(S, E, M)$.

2. Towards contradiction assume there is an attack $(T, h) \in R'$ with $T \cup \{h\} \subseteq (E \cap S)$. This means there is an attack $(T', h) \in R$ with $T \subseteq T'$. As E is admissible in SF there is a counter-attack $(X, t) \in R \setminus M$ with $X \subseteq E$ for some $t \in T'$. If $t \notin S$ then $(T, h) \notin R'$, a contradiction. Therefore $t \in S$ and by assumption $t \in E$. However, this means $X \cup \{t\} \subseteq E$, a contradiction to the conflict-freeness of E in SF. □

Lemma 6.24. *Let $SF = (A, R)$ be a SETAF, let $M \subseteq R$, let $E \subseteq A$ such that*

$$(E \cap S) \in adm\left(SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S, E, M), M_{SF}(S, E, M)\right)$$

for all $S \in SCCs(SF)$. Moreover, let $S' \in SCCs(SF)$ and let $a \in U_{SF}(S', E, M)$ be acceptable w.r.t. $(E \cap S')$ in $SF\Downarrow^{(E\setminus S')^+}_{UP_{SF}(S',E,M)}$ considering $M_{SF}(S', E, M)$. Then a is acceptable w.r.t. E in SF considering M.

Proof. We have to show for each $(T, a) \in R$ that E attacks T in SF with non-mitigated attacks. As before we set $SF\Downarrow^{(E\setminus S')^+}_{UP_{SF}(S',E,M)} = SF' = (A', R')$. We distinguish the following three cases:

1. $(T \subseteq S')$ If $T \cap D_{SF}(S', E, M) \neq \emptyset$ we are done, because this means by Definition 6.20 there is an attack in $R \setminus M$ with $T \subseteq E$. Otherwise, all $t \in T$ are in $UP_{SF}(S', E, M)$. Then, $(T, a) \in R'$ and there must be a (not mitigated) counter-attack (X, t) with $t \in T$ and $X \subseteq E \cap S'$ within SF', as we assumed a is acceptable w.r.t. $E \cap S'$ in SF' considering $M_{SF}(S', E, M)$. This means there is an attack $(X', t) \in R \setminus M$ with $X \subseteq X'$, and as (X, t) is not mitigated in SF' we know $X' \subseteq E$. In summary, a is acceptable w.r.t. E in SF considering M.

2. $(T \subseteq A \setminus S')$ Then $T \cap E^+ \neq \emptyset$ by $a \in U_{SF}(S', E, M)$ (otherwise, if E would not attack T in SF, if $(T, a) \in M$ then $a \in P_{SF}(S', E, M)$, and if $(T, a) \notin M$ then $a \in D_{SF}(S', E, M)$).

3. $(T \cap S' \neq \emptyset$ and $T \cap (A \setminus S') \neq \emptyset)$ Assume towards contradiction there is no non-mitigated attack from E to T in SF. Then there is an attack $(T', a) \in R'$ with $T' \subseteq T$ and $(T', a) \notin M_{SF}(S', E, M)$. Now the reasoning proceeds as in case (1).

As we established that there are counter-attacks in all cases (1)-(3), the desired property holds. □

Proposition 6.25. *Let $SF = (A, R)$ be a SETAF and let $E \subseteq A$ be a set of arguments. Then for each $C \subseteq A$ and $M \subseteq R$ it holds $E \in adm(SF, C, M)$ if and only if $\forall S \in SCCs(SF)$ it holds $(E \cap S) \in adm(SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S, E, M) \cap C, M_{SF}(S, E, M))$.*

Proof. Let $SF' = SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}$.

(⇒) Since $E \subseteq C$ and all $a \in E$ are acceptable w.r.t. E in SF considering M, we can apply Lemma 6.23 and get that every $a \in (E \cap S)$ are in $U_{SF}(S, E, M) \cap C$ for any given SCC S. Moreover, we get that a is acceptable w.r.t. $(E \cap S)$ in SF' considering $M_{SF}(S, E, M)$ and that $(E \cap S)$ is conflict-free in SF'. Hence, $(E \cap S)$ is admissible in SF' considering $M_{SF}(E, S, M)$.

(⇐) As for all SCCs S we assume $(E \cap S) \subseteq (S \cap C)$ we know $E \subseteq C$, i.e. we only need to show admissibility in SF considering M. Towards contradiction assume E is not conflict-free in SF, i.e. there is an attack $(T, h) \in R$ with $T \cup \{h\} \subseteq E$. Let S' be the SCC containing h.

1. We cannot have (1) $T \subseteq S'$, as this would contradict the assumption that $E \cap S'$ is conflict-free in $SF{\Downarrow}^{(E\setminus S')^+}_{UP_{SF}(S',E)}$ (note that in this case the attack (T,h) is also necessarily in SF').

2. Moreover, it cannot be that (2) $T \subseteq A \setminus S'$, because then $h \in D_{SF}(S',E,M)$ (or $h \in P_{SF}(S',E,M)$ if $(T,h) \in M$) while we assumed $h \in U_{SF}(S',E,M)$.

3. Finally, consider the case (3) where $T \cap S' \neq \emptyset$ and $T \cap (A \setminus S') \neq \emptyset$. Then there is a non-mitigated attack from $(E \setminus S')$ to T, as otherwise there would be $(T \cap S', h) \in SF{\Downarrow}^{(E\setminus S')^+}_{UP_{SF}(S',E,M)}$, contradicting our assumption of local conflict-freeness. Call this attack $(X,t) \in R \setminus M$ with $X \subseteq (E \setminus S')$ and $t \in T \setminus S'$. Let S'' be the SCC t is in. As before, we cannot have (1) $X \subseteq S''$ or (2) $X \subseteq A \setminus S''$. The only remaining case is again (3) $X \cap S'' \neq \emptyset$ and $X \cap (A \setminus S'') \neq \emptyset$—as this step (3) always takes us to a prior SCC and we assume SF finite, eventually this recursion will stop in case (1) or (2). Now, by induction we get a contradiction for the initial case.

It remains to show that every $a \in E$ is acceptable w.r.t. E in SF considering M. Let S^* be the SCC a is in and let $SF^* = SF{\Downarrow}^{(E\setminus S^*)^+}_{UP_{SF}(S^*,E,M)}$. By assumption, $(E \cap S^*) \in$ *adm*$(SF^*, U_{SF}(S^*,E,M), M_{SF}(S^*,E,M))$, i.e. a is acceptable w.r.t. $E \cap S$ in SF^* considering $M_{SF}(S^*,E,M)$. Since we also have $a \in U_{SF}(S^*,E,M)$, we can apply Lemma 6.24 and get that a is acceptable w.r.t. E in SF considering M. □

Proposition 6.28. *Let $SF = (A,R)$ be a SETAF, let $M \subseteq R$, and let $E \subseteq A$ be a set of arguments. Then $\forall C \subseteq A$ it holds $E \in com(SF,C,M)$ if and only if $\forall S \in SCCs(SF)$ it holds $(E \cap S) \in com(SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S,E,M) \cap C, M_{SF}(S,E,M))$.*

Proof. ($\Rightarrow$) If $E \in com(SF,C,M)$, then in particular $E \in adm(SF,C,M)$. Hence by Proposition 6.25 we get

$$\forall S \in SCCs(SF): \quad (E \cap S) \in adm(SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}, UP_{SF}(S,E,M) \cap C, M_{SF}(S,E,M)).$$

For an arbitrary SCC $S' \in SCCs(SF)$, let $a \in U_{SF}(S',E,M)$ be an argument acceptable w.r.t. $(E \cap S')$ in $SF{\Downarrow}^{(E\setminus S')^+}_{UP_{SF}(S',E,M)}$ considering $M_{SF}(S',E,M)$. By Lemma 6.24, a is acceptable w.r.t. E in SF considering M, and, hence, $a \in E$ and $a \in E \cap S'$ by completeness.

($\Leftarrow$) We get $E \in adm(SF,C,M)$ by Proposition 6.25. For an arbitrary $a \in C$, let S' be the SCC a is in. If a is acceptable w.r.t. E in SF considering M, by Lemma 6.23 we get that a is acceptable w.r.t. $(E \cap S')$ in $SF{\Downarrow}^{(E\setminus S')^+}_{UP_{SF}(S',E,M)}$ considering $M_{SF}(S',E,M)$. As $(E \cap S')$ is locally complete, we get $a \in E$. □

Lemma 6.30. *Let $SF = (A,R)$, let $M \subseteq R$, and let $E \in adm(SF,A,M)$, let $S \in SCCs(SF)$ be an SCC. Moreover, let $E' \subseteq A$ be a set of arguments such that $(E \cap S) \subseteq E' \subseteq U_{SF}(S,E,M)$, and $E' \in adm(SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S,E,M), M_{SF}(S,E,M))$. Then $E \cup E'$ is admissible in SF considering M.*

Proof. We first show that $(E \cup E')$ is conflict-free in SF. Again, let $SF' = SF{\Downarrow}^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}$. Assume towards contradiction there is $(T,h) \in R$ with $T \cup \{h\} \subseteq (E \cup E')$. Then we have either (1) $h \in E$ or (2) $h \in E' \setminus E$.

(1) Since $E \in adm(SF, A, M)$ in (1) we have $E \mapsto_R T$ via a non-mitigated attack. We have $E \in cf(SF)$, this means $E \mapsto_R T'$ where $T' = T \setminus E = T \cap E' \neq \emptyset$. But this means $E' \cap D_{SF}(S, E, M) \neq \emptyset$, contradicting our assumption $E' \subseteq U_{SF}(S, E, M)$.

(2) Regarding (2), if $T \subseteq E$, then if $(T, h) \notin M$ we have $h \in D_{SF}(S, E, M)$ (or $h \in P_{SF}(S, E, M)$ if $(T, h) \in M$), a contradiction. Hence, $T \cap (E' \setminus E) \neq \emptyset$. It follows there is $(T', h) \in R(SF')$ with $T' \subseteq T$. However, since we assume E' is conflict-free in SF' it holds $T' \cup \{h\} \not\subseteq E'$, a contradiction because $E \cap S \subseteq E'$.

As both possibilities lead to contradictions, we conclude $(E \cup E') \in cf(SF)$.

It remains to show defense in SF considering M, i.e. for all $(T, h) \in R$ with $h \in E \cup E'$ we show $E \cup E' \mapsto_R T$ via non-mitigated attacks. If $h \in E$, this follows from $E \in adm(SF, A, M)$. For $h \in E' \setminus E$, we distinguish 4 cases:

(1) E attacks T via non-mitigated attacks, then we are done.

(2) $T \subseteq A \setminus S$. But then either $E \mapsto_R T$ via non-mitigated attacks—see (1)—or all attacks from E to T are mitigated, or $E \not\mapsto_R T$, and therefore $h \in P_{SF}(S, E, M)$, a contradiction to $h \in E' \subseteq U_{SF}(S, E, M)$.

(3) $T \subseteq A(SF')$. But this means $(T, h) \in R(SF')$ and therefore since E' is admissible in this context there is a non-mitigated counter-attack $(X, t) \in R(SF')$ with $X \subseteq E'$ and $t \in T$. As (X, t) is not mitigated, there is a non-mitigated "original" attack $(X', t) \in R$ with $X' \supseteq X$ and $X' \setminus X \subseteq E$, i.e. $X' \subseteq (E \cup E')$, contradicting the earlier established conflict-freeness.

(4) $T \cap A(SF') \neq \emptyset$ and $T \cap (A \setminus A(SF')) \neq \emptyset$. If we assume we are not in case (1) then there is an attack $(T', h) \in R(SF')$, and we proceed as in (3). In any case, there is a defense against the attack (T, h), therefore, $(E \cup E') \in adm(SF, A, M)$. □

Proposition 6.31. *Let $SF = (A, R)$ be a SETAF, let $M \subseteq R$ and let $E \subseteq A$ be a set of arguments. Then $\forall C \subseteq A$ it holds $E \in pref(SF, C, M)$ if and only if $\forall S \in SCCs(SF)$ it holds $(E \cap S) \in pref(SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S, E, M) \cap C, M_{SF}(S, E, M))$.*

Proof. ($\Rightarrow$) We assume $E \in pref(SF, C, M)$, and can apply Proposition 6.25 and obtain that

$$\forall S \in SCCs(SF): \quad (E \cap S) \in adm(SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S, E, M) \cap C, M_{SF}(S, E, M)).$$

Assume towards contradiction that for some $S' \in SCCs(SF)$ there is a set

$$E' \in adm(SF\Downarrow^{(E\setminus S')^+}_{UP_{SF}(S',E,M)}, U_{SF}(S', E, M) \cap C, M_{SF}(S', E, M))$$

with $E \cap S' \subsetneq E'$. However, by Lemma 6.30 this means the set $E \cup E'$ is in $adm(SF, C, M)$, but since $E \subsetneq E \cup E'$ this contradicts our assumption $E \in pref(SF, C, M)$.

($\Leftarrow$) From Proposition 6.25 we get $E \in adm(SF, C, M)$. Towards contradiction assume there is an $E' \in adm(SF, C; M)$ with $E' \supsetneq E$. This means there is some SCC $S \in SCCs(SF)$ such that $(E \cap S) \subsetneq (E' \cap S)$. W.l.o.g. we choose S such that no ancestor SCC of S has this property. This means that $U_{SF}(S, E, M) = U_{SF}(S, E', M)$

and $P_{SF}(S, E, M) = P_{SF}(S, E', M)$ for S and all of its ancestor SCCs. Consequently, $(E' \cap S) \subseteq U_{SF}(S, E, M) = U_{SF}(S, E', M)$, and by another application of Proposition 6.25 we get $E' \in adm(SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S, E, M) \cap C, M_{SF}(S, E, M))$. However, this contradicts our assumption $E \in pref(SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S, E, M) \cap C, M_{SF}(S, E, M))$. □

Proposition 6.33. *Let $SF = (A, R)$ be a SETAF, $M \subseteq R$, and let $E \subseteq A$ be a set of arguments. Then $\forall C \subseteq A$ it holds $E \in grd(SF, C, M)$ if and only if $\forall S \in SCCs(SF)$ it holds $(E \cap S) \in grd(SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S, E, M) \cap C, M_{SF}(S, E, M))$.*

Proof. ($\Rightarrow$) We assume $E \in grd(SF, C)$, and can apply Proposition 6.28 and obtain that

$$\forall S \in SCCs(SF) : \quad (E \cap S) \in com(SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S, E, M) \cap C, M_{SF}(S, E, M)).$$

Assume towards contradiction that for some SCC S' the set $(E \cap S')$ is not minimal among the locally complete extensions. W.l.o.g. we choose S' such that no ancestor SCC of S' has this property. Let $E' \in grd(SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}, U_{SF}(S, E, M) \cap C, M_{SF}(S, E, M))$. We can construct E'' such that for the ancestor SCCs of S' the new set E'' coincides with E, for S' it coincides with E', and for the remaining SCCs S is determined by $grd(SF\Downarrow^{(E'\setminus S)^+}_{U_{SF}(S,E',M)}, UP_{SF}(S, E', M) \cap C, M_{SF}(S, E', M))$ (see Section 7 for details). But then $E'' \in com(SF, C, M)$ by Proposition 6.28 and $E \not\subseteq E''$, a contradiction to our assumption $E \in grd(SF, C, M)$.

($\Leftarrow$) We get $E \in com(SF, C, M)$ by Proposition 6.28. Towards contradiction assume there is some $E' \subsetneq E$ with $E' \in grd(SF, C, M)$. This means there is an SCC S where $(E' \cap S) \subsetneq (E \cap S)$. W.l.o.g., we choose S such that no ancestor SCC of S has this property. This means that $U_{SF}(S, E, M) = U_{SF}(S, E', M)$ and $P_{SF}(S, E, M) = P_{SF}(S, E', M)$ for S and its ancestor SCCs. Consequently,

$$(E' \cap S) \in com(SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}, UP_{SF}(S, E, M) \cap C, M_{SF}(S, E, M)).$$

However, this contradicts our assumption

$$E \in grd(SF\Downarrow^{(E\setminus S)^+}_{UP_{SF}(S,E,M)}, UP_{SF}(S, E, M) \cap C, M_{SF}(S, E, M)),$$

since $(E' \cap S) \subsetneq (E \cap S)$. □

Appendix D. Proof Details of Section 7

Proposition 7.7. *Let $SF = (A, R)$ be a SETAF. If $|com(SF)| \geq 2$ then SF contains an even-cycle.*

Proof. (cf. (Dvořák, 2012; König, 2020)) Let G be the grounded extension of SF and E a complete extension such that $E \neq G$. We thus have $E \supset G$. This means there is some $x_0 \in E \setminus G$ with $(Y_0, x_0) \in R$ such that $G \not\mapsto Y_0$. Since E is conflict-free and we have $x_0 \in E$, we also have $Y_0 \not\subseteq E$, i.e. there is an argument $y \in Y_0$ such that $y \notin E$. In order to defend x_0 in E there is some set $X_1 \subseteq E$ such that $(X_1, y_0) \in R$ for some $y_0 \in Y_0 \setminus E$.

Now the same reasoning holds for some $x_1 \in X_1 \setminus G$, and inductively we get an infinite sequence $x_0, y_0, x_1, y_1, \dots$ such that $x_i \in E \setminus G$ and $(Y_i, x_i), (X_{i+1}, y_i) \in R$ with $y_i \in Y_i$ and $x_i \in X_i$ for $i \geq 1$. Since $E \setminus G$ is finite we get that $x_i = x_j$ for some $i \neq j$, then $(x_j, (X_j, y_{j-1}), y_{j-1}, (Y_{j-1}, x_{j-1}), x_{j-1}, \dots, (Y_i, x_i), x_i)$ is an even-cycle. □

Proof of the correctness of the translation illustrated in Example 7.13:

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

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

Lemma D.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. □

Theorem D.2 ((König, 2020))**.** *Let $\sigma \in \{cf, adm, stb, pref, stage, sem\}$. Then Tr_1 is an acceptance-preserving*[6] *translation for $\sigma \Rightarrow \sigma$ such that for every self-attack-free SETAF SF its translation $SF' = Tr_1(SF) = (A', R')$ is δ-symmetric.*

Proof. We will show two statements for each 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 $E' \cap A = E$ ("$\Rightarrow$"). Secondly we will show that for each extension $E' \in \sigma(SF')$ the corresponding extension $E = E' \cap A$ is an extension $E \in \sigma(SF)$ ("$\Leftarrow$").

1. For $\sigma = cf$:
"$\Rightarrow$": 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.
"$\Leftarrow$": 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 $\sigma = adm$:
"$\Rightarrow$": 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

6. Let σ, σ' be semantics, then a (SETAF-)translation Tr is called *acceptance-preserving* for $\sigma \Rightarrow \sigma'$ if for every SETAF $SF = (A, R)$ we have $\sigma(SF) = \{E \cap A \mid E \in \sigma'(Tr(SF))\}$ (König, 2020).

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.
"$\Leftarrow$": 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 $\sigma = stb$:
"$\Rightarrow$": 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 D.1 we have $A \subseteq E'^{\oplus}_{R'}$, and by construction we have $A' \setminus A \subseteq E'^{\oplus}_{R'}$.
"$\Leftarrow$": 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 D.1 we have $A \subseteq E^{\oplus}_{R}$.

4. For $\sigma = pref$:
"$\Rightarrow$": 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 $E \in pref(SF)$. 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 $S \in adm(SF)$ and $S \supset E$ again we have a contradiction to the assumption that $E \in pref(SF)$.
"$\Leftarrow$": Let $E' \in pref(SF')$ and let $E = E' \cap A$. We know $E \in adm(SF)$. Assume towards contradiction there is a set $S \in adm(SF)$ such that $S \supset E$. Let $S' = S \cup (E' \setminus E)$. By construction we have $S' \supset E'$. Moreover we have $S' \in adm(SF')$: 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 $S \in adm(SF)$. 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 $E' \in adm(SF')$ 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 $S' \in adm(SF')$, and since $S' \supset E'$ we have a contradiction to the assumption $E' \in pref(SF')$.

5. For $\sigma = stage$:
"$\Rightarrow$": Let $E \in stage(SF)$ 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. Assume towards contradiction there is a set $S' \in cf(SF')$ such that $S'^{\oplus}_{R'} \supset E'^{\oplus}_{R'}$. Let $S = S' \cap A$. We know $S \in cf(SF)$. Moreover we have $S^{\oplus}_{R} \supseteq E^{\oplus}_{R}$ by Lemma D.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 $E \in stage(SF)$.
"$\Leftarrow$": Let $E' \in stage(SF')$ and let $E = E' \cap A$. We know $E \in cf(SF)$. Assume towards contradiction there is a set $S \in cf(SF)$ 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 $S' \in cf(SF')$ by construction. As before we have $A' \setminus A \subseteq S'^{\oplus}_{R'}$. Moreover, by Lemma D.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 $E' \in stage(SF')$.

6. For $\sigma = sem$:
"$\Rightarrow$": Let $E \in sem(SF)$ and let $E' = \{a^1_{r,t} \mid r = (T,h), t \in T, E \mapsto_R t\}$. We already know $E' \in adm(SF')$. Assume towards contradiction there is a set $S' \in adm(SF')$ such that $S'^{\oplus}_{R'} \supset E'^{\oplus}_{R'}$. Let $S = S' \cap A$. We know $S \in adm(SF)$. Moreover by Lemma D.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 D.1 we get $S^{\oplus}_{R} \supset E^{\oplus}_{R}$, which is a contradiction to our assumption $E \in sem(SF)$. 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 $E \in sem(SF)$.
"$\Leftarrow$": Let $E' \in sem(SF')$ and let $E = E' \cap A$. We know $E \in adm(SF)$. Assume towards contradiction there is a set $S \in adm(SF)$ such that $S^{\oplus}_{R} \supset E^{\oplus}_{R}$. Let $S' = \{a^1_{r,t} \mid r = (T,h), t \in T, S \mapsto_R t\}$. By construction we have $S' \in adm(SF')$. By Lemma D.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 $a = a^1_{r,t}$ (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 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 $t \in E^{\oplus}_{R}$, so by assumption $t \in S^{\oplus}_{R}$, and then again by construction $a \in S'^{\oplus}_{R'}$ (either because $t \in S'$ or because $S \mapsto_R t$). □

Proposition 7.16. *For fully-symmetric SETAFs the problems $Cred_\sigma$ for $\sigma \in \{adm, stb, pref, com\}$ are* NP*-complete.*

Proof. Let φ be a CNF-formula with clauses C over atoms Y. We construct SF_φ according to Reduction 7.15.

Note that no argument $c \in C$ can be in an admissible set, as the argument φ cannot be defeated. As with the standard reduction, φ is acceptable if and only if the assignment corresponding to the arguments $X \cup \bar{X}$ satisfies the original formula, thus concluding the reduction from the boolean satisfiability problem. □

Lemma 7.20. *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. "$\Rightarrow$": We will show inductively that for every iteration of the algorithm the arguments that are removed in step 6 cannot be defended 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.

"$\Leftarrow$": 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 $x \in S$. □

Proof for stage semantics of Theorem 7.24 (König, 2020):

Reduction D.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*

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

We have (as we will show in Lemma D.5) 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

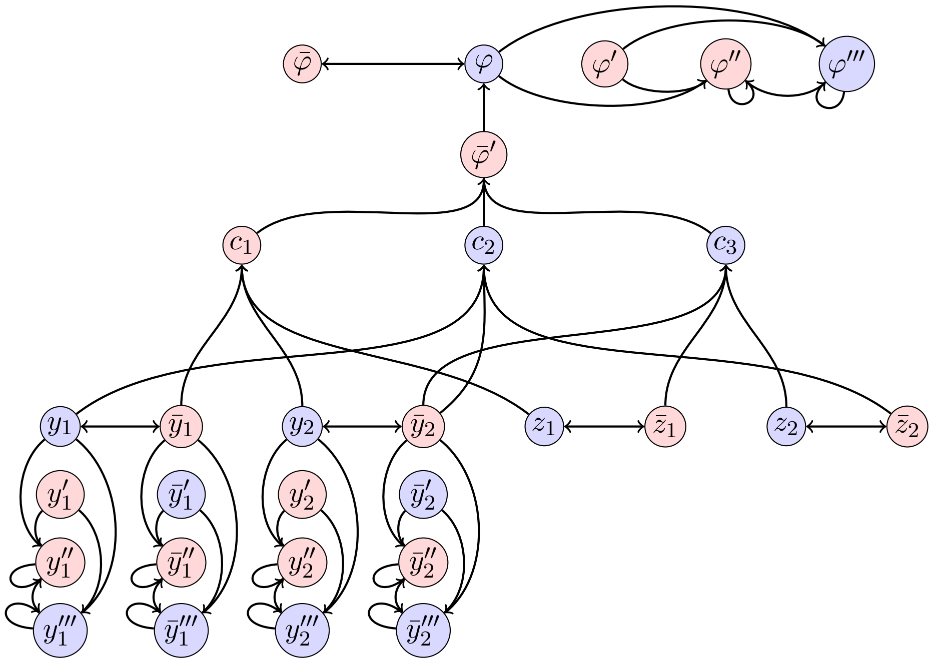

Figure 3: Illustration of SF_3^Φ for $\Phi = \forall Y \exists Z \varphi(Y, Z)$ with $Y = \{y_1, y_2\}$, $Z = \{z_1, z_2\}$, and $\varphi = \{\{y_1, \bar{y}_2, \bar{z}_1\}, \{\bar{y}_1, y_2, z_2)\}, \{y_2, z_1, \bar{z}_2\}\}$. The coloring of the arguments corresponds to a possible partitioning that shows the 2-colorability of SF_3^Φ, i.e. we have that no attack is monochromatic.

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 3 corresponds to such a partitioning).
The following proof follows the structure of (Dvořák, 2012, p. 52-55).

Lemma D.4. *Let Φ be a $QBF_\forall^2$ formula and let $SF_3^\Phi = (A, R)$, then for every extension $E \in stage(SF_3^\Phi)$ 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 $E_R'^{\oplus} \supset E_R^{\oplus}$. 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 $E_R'^{\oplus} \supset E_R^{\oplus}$. By conflict-freeness we also have that at most one of x and $\bar{x}$ is in E. □

Lemma D.5. *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 $E \in stage(SF_3^\Phi)$.*

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

Lemma D.6. *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 $E \in \mathit{stage}(SF_3^\Phi)$ 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 D.4 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, $E \in \mathit{stage}(SF_3^\Phi)$ 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 $E \in \mathit{stage}(SF_3^\Phi)$. □

From this the desired result immediately follows.

Proposition D.7. *For 2-colorable SETAFs the complexity of $\mathit{Cred}_{\mathit{stage}}$ and $\mathit{Skept}_{\mathit{stage}}$ coincides with the complexity of the general problem.*

Proposition 7.26. *The problem $\mathit{Ver}_{\mathit{pref}}$ remains* coNP*-complete even for self-attack-free SETAFs SF where all SCCs $S \in \mathit{SCCs}(SF)$ are fully-symmetric, i.e., $SF{\downarrow_S}$ is fully symmetric.*

Proof. Let $SF = (A, R)$ be an arbitrary SETAF and let $x, y \notin A$ be new arguments. We define $SF' = (A \cup \{x, y\}, R')$, with

$$R' \;=\; R \cup \{(\{y, z\}, t) \mid (t, z) \in \mathit{primal}(SF), (z, t) \notin \mathit{primal}(SF)\}.$$

In the resulting framework SF' clearly all SCCs are primal-symmetric (note that $\{a\}, \{b\}$ are SCCs of SF', as well as all loosely connected components of SF). Clearly $x \in G$ and $y \in G^+$ for $G \in \mathit{grd}(SF')$, and hence, $x \in E$ and $y \in E^+$ for all $E \in \mathit{pref}(SF')$. By construction we can apply the reduct on $\{x\}$ and obtain $SF'^{\{x\}} = SF$, and by modularization we hence get $\mathit{pref}(SF') = \{E \cup \{x\} \mid E \in \mathit{pref}(SF)\}$, which means $E \in \mathit{pref}(SF)$ if and only if $E \cup \{x\} \in \mathit{pref}(SF')$. □

Proposition 7.30. *Let $SF = (A, R)$ be a primal-bipartite SETAF and $C, E \subseteq A$, $M \subseteq R$. We can decide whether $E \in \mathit{pref}(SF, C, M)$ in polynomial time.*

Proof. We can check in polynomial time whether $E \in adm(SF, C, M)$, and compute SF^E_M. By Proposition 7.29 it suffices to check whether there is a non-empty set $E' \in adm(SF', C', M')$, which we can find out by running Algorithm 2 for both both partitions. □

Proposition 7.31. *Let $SF = (A, R)$ be an odd-cycle-free SETAF with $|SCCs(SF)| = 1$. Then SF is primal-bipartite.*

Proof. Let x be an arbitrary argument from SF, we say *x reaches argument $y \in A$ in n steps* if there are attacks $(X_1, x_2), (X_2, x_3), \ldots, (X_n, y) \in R$ with $x \in X_1$, $x_i \in X_i$ for $1 \leq i \leq n$. Let $S = \{a \in A \mid a$ can be reached from x in an even number of steps $\}$. Then $(S, A \setminus S)$ is a partitioning for SF: Clearly, $A \setminus S$ is the set of arguments that can be reached from x in an *odd* number of steps, and because SF is strongly connected at the same time the set of argument from which x can be reached in a odd number of steps. Assume towards contradiction there are two arguments $a, b \in (A \setminus S)$ s.t. a reaches b in 1 step. However, this introduces an odd-length primal-cycle, as x reaches a in an odd number of steps, a reaches b in 1 step, and b reaches x in an odd number of steps, a contradiction. Likewise, there can be no pair a, b of arguments in S where a reaches b in 1 step. □

Theorem 7.33. *Let SF be a SETAF where for all SCCs $S \in SCCs(SF)$ it holds either*

- *S is acyclic,*
- *S is even-cycle-free,*
- *S is primal-bipartite,*
- *S odd-cycle-free, or*
- *the size of S is bounded by a parameter k, i.e., $|S| \leq k$.*

Then we can verify a given preferred extensions in $O\big(2^k \cdot poly(|SF|)\big)$ for some polynomial function poly.

Proof. Follows from Proposition 7.30, Proposition 7.31, and Proposition 7.32 and the fact that acyclic SETAFs are also even- and odd-cycle-free. □

References

Baroni, P., Caminada, M., & Giacomin, M. (2011). An introduction to argumentation semantics. *Knowledge Eng. Review*, *26*(4), 365–410.

Baroni, P., & Giacomin, M. (2007). On principle-based evaluation of extension-based argumentation semantics. *Artif. Intell.*, *171*(10-15), 675–700.

Baroni, P., Giacomin, M., & Guida, G. (2005). SCC-recursiveness: a general schema for argumentation semantics. *Artif. Intell.*, *168*(1-2), 162–210.

Baroni, P., Giacomin, M., & Liao, B. (2014). On topology-related properties of abstract argumentation semantics. A correction and extension to dynamics of argumentation systems: A division-based method. *Artif. Intell.*, *212*, 104–115.

Baroni, P., Rago, A., & Toni, F. (2019). From fine-grained properties to broad principles for gradual argumentation: A principled spectrum. *Int. J. Approx. Reason.*, *105*, 252–286.

Baumann, R. (2011). Splitting an argumentation framework. In *Proceedings of the 11th International Conference on Logic Programming and Nonmonotonic Reasoning, LPNMR 2011*, Vol. 6645 of *Lecture Notes in Computer Science*, pp. 40–53. Springer.

Baumann, R., Brewka, G., & Ulbricht, M. (2020a). Comparing weak admissibility semantics to their dung-style counterparts - reduct, modularization, and strong equivalence in abstract argumentation. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR 2020*, pp. 79–88.

Baumann, R., Brewka, G., & Ulbricht, M. (2020b). Revisiting the foundations of abstract argumentation - semantics based on weak admissibility and weak defense. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*, pp. 2742–2749. AAAI Press.

Bikakis, A., Cohen, A., Dvořák, W., Flouris, G., & Parsons, S. (2021). Joint attacks and accrual in argumentation frameworks. *FLAP*, *8*(6), 1437–1501.

Bonzon, E., Delobelle, J., Konieczny, S., & Maudet, N. (2017). A parametrized ranking-based semantics for persuasion. In *Proceedings of the 11th International Conference on Scalable Uncertainty Management, SUM 2017*, Vol. 10564 of *Lecture Notes in Computer Science*, pp. 237–251. Springer.

Caminada, M. (2014). Strong admissibility revisited. In *Proceedings of the 5th International Conference on Computational Models of Argument, COMMA 2014*, Vol. 266 of *Frontiers in Artificial Intelligence and Applications*, pp. 197–208. IOS Press.

Caminada, M., Carnielli, W. A., & Dunne, P. E. (2012). Semi-stable semantics. *J. Log. Comput.*, *22*, 1207–1254.

Cerutti, F., Giacomin, M., Vallati, M., & Zanella, M. (2014). An SCC recursive meta-algorithm for computing preferred labellings in abstract argumentation. In *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning, KR 2014*. AAAI Press.

Coste-Marquis, S., Devred, C., & Marquis, P. (2005). Symmetric argumentation frameworks. In *Proceedings of the 8th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, ECSQARU 2005*, Vol. 3571 of *Lecture Notes in Computer Science*, pp. 317–328. Springer.

Diller, M., Keshavarzi Zafarghandi, A., Linsbichler, T., & Woltran, S. (2020). Investigating subclasses of abstract dialectical frameworks. *Argument & Computation*, *11*, 191–219.

Dimopoulos, Y., & Torres, A. (1996). Graph theoretical structures in logic programs and default theories. *Theor. Comput. Sci.*, *170*(1-2), 209–244.

Dung, P. M. (1995). On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.*, *77*(2), 321–358.

Dunne, P. E. (2007). Computational properties of argument systems satisfying graph-theoretic constraints. *Artif. Intell.*, *171*(10-15), 701–729.

Dunne, P. E. (2009). The computational complexity of ideal semantics. *Artif. Intell.*, *173*(18), 1559–1591.

Dunne, P. E., & Bench-Capon, T. J. M. (2002). Coherence in finite argument systems. *Artif. Intell.*, *141*(1/2), 187–203.

Dunne, P. E., Dvořák, W., Linsbichler, T., & Woltran, S. (2015). Characteristics of multiple viewpoints in abstract argumentation. *Artif. Intell.*, *228*, 153–178.

Dvořák, W. (2012). *Computational Aspects of Abstract Argumentation*. Ph.D. thesis, Vienna University of Technology, Institute of Information Systems.

Dvořák, W., & Dunne, P. E. (2017). Computational problems in formal argumentation and their complexity. *FLAP*, *4*(8), 2557–2622.

Dvořák, W., Fandinno, J., & Woltran, S. (2019). On the expressive power of collective attacks. *Argument Comput.*, *10*(2), 191–230.

Dvořák, W., & Gaggl, S. A. (2016). Stage semantics and the SCC-recursive schema for argumentation semantics. *J. Log. Comput.*, *26*(4), 1149–1202.

Dvořák, W., Greßler, A., & Woltran, S. (2018). Evaluating SETAFs via answer-set programming. In *Proceedings of the Second International Workshop on Systems and Algorithms for Formal Argumentation, SAFA 2018*, Vol. 2171 of *CEUR Workshop Proceedings*, pp. 10–21. CEUR-WS.org.

Dvořák, W., Hecher, M., König, M., Schidler, A., Szeider, S., & Woltran, S. (2022). Tractable abstract argumentation via backdoor-treewidth. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence, AAAI 2022*, pp. 5608–5615. AAAI Press.

Dvořák, W., Järvisalo, M., Wallner, J. P., & Woltran, S. (2014). Complexity-sensitive decision procedures for abstract argumentation. *Artif. Intell.*, *206*(0), 53 – 78.

Dvořák, W., König, M., Ulbricht, M., & Woltran, S. (2021). A reduct-driven study of argumentation frameworks with collective attacks. In *Proceedings of the 19th International Workshop on Non-Monotonic Reasoning, NMR 2021*, pp. 285–294.

Dvořák, W., König, M., Ulbricht, M., & Woltran, S. (2022). Rediscovering argumentation principles utilizing collective attacks. In *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning, KR 2022*, pp. 122–131.

Dvořák, W., König, M., & Woltran, S. (2021a). Graph-classes of argumentation frameworks with collective attacks. In *Proceedings of the 17th European Conference on Logics in Artificial Intelligence, JELIA 2021*, Vol. 12678 of *Lecture Notes in Computer Science*, pp. 3–17. Springer.

Dvořák, W., König, M., & Woltran, S. (2021b). On the complexity of preferred semantics in argumentation frameworks with bounded cycle length. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR 2021*, pp. 671–675.

Dvořák, W., König, M., & Woltran, S. (2022a). Deletion-backdoors for argumentation frameworks with collective attacks. In *Proceedings of the 4th International Work-*

shop on Systems and Algorithms for Formal Argumentation, SAFA 2022, Vol. 3236 of *CEUR Workshop Proceedings*, pp. 98–110. CEUR-WS.org.

Dvořák, W., König, M., & Woltran, S. (2022b). Treewidth for argumentation frameworks with collective attacks. In *Proceedings of the 9th International Conference on Computational Models of Argument, COMMA 2022*, pp. 140–151. IOS Press.

Dvořák, W., Ordyniak, S., & Szeider, S. (2012). Augmenting tractable fragments of abstract argumentation. *Artificial Intelligence*, *186*(0), 157–173.

Dvořák, W., Rapberger, A., & Woltran, S. (2020). Argumentation semantics under a claim-centric view: Properties, expressiveness and relation to SETAFs. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR 2020*, pp. 341–350.

Dvořák, W., & Woltran, S. (2010). Complexity of semi-stable and stage semantics in argumentation frameworks. *Inf. Process. Lett.*, *110*(11), 425–430.

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

Flouris, G., & Bikakis, A. (2019). A comprehensive study of argumentation frameworks with sets of attacking arguments. *Int. J. Approx. Reason.*, *109*, 55–86.

Friese, T., & Ulbricht, M. (2021). On the relationship of modularity notions in abstract argumentation. In *Proceedings of the 19th International Workshop on Non-Monotonic Reasoning, NMR 2021*, pp. 51–60.

Gaggl, S. A., Rudolph, S., & Straß, H. (2021). On the decomposition of abstract dialectical frameworks and the complexity of naive-based semantics. *J. Artif. Intell. Res.*, *70*, 1–64.

Kaci, S., van der Torre, L. W. N., & Villata, S. (2018). Preference in abstract argumentation. In *Proceedings of the 7th International Conference on Computational Models of Argument, COMMA 2018*, Vol. 305 of *Frontiers in Artificial Intelligence and Applications*, pp. 405–412. IOS Press.

Keshavarzi Zafarghandi, A., Verbrugge, R., & Verheij, B. (2022). Strong admissibility for abstract dialectical frameworks. *Argument Comput.*, *13*(3), 249–289.

König, M. (2020). Graph-classes of argumentation frameworks with collective attacks. Master's thesis, TU Wien.

König, M., Rapberger, A., & Ulbricht, M. (2022). Just a matter of perspective. In *Proceedings of the 9th International Conference on Computational Models of Argument, COMMA 2022*, Vol. 353 of *Frontiers in Artificial Intelligence and Applications*, pp. 212–223. IOS Press.

Liao, B. S., Jin, L., & Koons, R. C. (2011). Dynamics of argumentation systems: A division-based method. *Artif. Intell.*, *175*(11), 1790–1814.

Linsbichler, T. (2014). Splitting abstract dialectical frameworks. In *Proceedings of the 5th International Conference on Computational Models of Argument, COMMA 2014*, Vol. 266 of *Frontiers in Artificial Intelligence and Applications*, pp. 357–368. IOS Press.

Linsbichler, T., Pührer, J., & Strass, H. (2016). A uniform account of realizability in abstract argumentation. In *Proceedings of the 22nd European Conference on Artificial Intelligence, ECAI 2016*, Vol. 285 of *Frontiers in Artificial Intelligence and Applications*, pp. 252–260. IOS Press.

Nielsen, S. H., & Parsons, S. (2006). A generalization of Dung's abstract framework for argumentation: Arguing with sets of attacking arguments. In *Proceedings of the 3rd International Workshop on Argumentation in Multi-Agent Systems, ArgMAS 2006*, Vol. 4766 of *Lecture Notes in Computer Science*, pp. 54–73. Springer.

Polberg, S. (2016). Understanding the abstract dialectical framework. In *Proceedings of the 15th European Conference on Logics in Artificial Intelligence, JELIA 2016*, Vol. 10021 of *Lecture Notes in Computer Science*, pp. 430–446.

Polberg, S. (2017). *Developing the Abstract Dialectical Framework*. Ph.D. thesis, Vienna University of Technology, Institute of Information Systems.

Rossit, J., Mailly, J., Dimopoulos, Y., & Moraitis, P. (2021). United we stand: Accruals in strength-based argumentation. *Argument Comput.*, *12*(1), 87–113.

van der Torre, L., & Vesic, S. (2017). The principle-based approach to abstract argumentation semantics. *FLAP*, *4*(8), 2735–2778.

Yu, L., Chen, D., Qiao, L., Shen, Y., & van der Torre, L. (2021). A principle-based analysis of abstract agent argumentation semantics. In *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning, KR 2021*, pp. 629–640.

Yun, B., Vesic, S., & Croitoru, M. (2020). Sets of attacking arguments for inconsistent datalog knowledge bases. In *Proceedings of the 8th International Conference on Computational Models of Argument, COMMA 2020*, Vol. 326 of *Frontiers in Artificial Intelligence and Applications*, pp. 419–430. IOS Press.