

On Cycles, Attackers and Supporters – A Contribution to The Investigation of Dynamics in Abstract Argumentation

Ringo Baumann and Markus Ulbricht

Department of Computer Science, Leipzig University
{baumann,mulbricht}@informatik.uni-leipzig.de

Abstract


Abstract argumentation as defined by Dung in his seminal 1995 paper is by now a major research area in knowledge representation and reasoning. Dynamics of abstract argumentation frameworks (AFs) as well as syntactical consequences of semantical facts of them are the central issues of this paper. The first main part is engaged with the systematical study of the influence of attackers and supporters regarding the acceptability status of whole sets and/or single arguments. In particular, we investigate the impact of addition or removal of arguments, a line of research that has been around for more than a decade. Apart from entirely new results, we revisit, generalize and sum up similar results from the literature. To gain a comprehensive formal and intuitive understanding of the behavior of AFs we put special effort in comparing different kinds of semantics. We concentrate on classical admissibility-based semantics as well as semantics based on naivity. In the second main part we show how to infer syntactical information from semantical one. For instance, it is well-known that if a finite AF possesses no stable extension, then it has to contain an odd-cycle. In this paper, we even present a characterization of this issue. Moreover, we show that the change of the number of extensions if adding or removing an argument allows to conclude the existence of certain even or odd cycles in the considered AF without having further information.

1 Introduction

Abstract argumentation as defined by Dung in his seminal paper [Dung, 1995] is by now a major research area in knowledge representation and reasoning. Around the beginning of the 2010s several problems regarding *dynamics* of abstract argumentation have been addressed in the literature [Boella *et al.*, 2009; Cayrol *et al.*, 2010; Bisquert *et al.*, 2011; Liao *et al.*, 2011]. Within argumentation dynamics, a major strand of research is concerned with the so-called *enforcing problem* dealing with the question of how to modify a given AF such that a desired set of arguments becomes an extension [Baumann and Brewka, 2010; Baumann, 2012;

Doutre and Mailly, 2018; Wallner, 2020; Wallner *et al.*, 2017]. In this paper we continue another line of research, namely the formal study of syntactic manipulations preserving the acceptability status of former extensions and/or single arguments. Consider the following example.

Example 1.1 (Preserving acceptability). The graph below represents two AFs F and F_x , s.t. x and its associated attacks are contained in F but not in F_x . The two sets $E_1 = \{a, b_1\}$ and $E_2 = \{a, b_2\}$ are so-called preferred extensions of F (cf. background section). It is interesting to see that these sets semantically survive if we syntactically remove x . Note that x attacks both extensions. Moreover, deleting b , which also questions both sets would result in the same monotonic behavior. We will see that this is no coincidence since this kind of acceptability preserving can be shown formally.



In the second main part we consider how to infer syntactical information from semantical information. This line of research was already initiated by Dung. For instance, it is well-known that if a finite AF possesses no stable extension, then it has to contain an odd-cycle. We will strengthen this result in by presenting a first characterization of this issue. Further results are in the following spirit: If the number of stable extensions changes (semantical information) if moving from F to F_x , then there has to be a cycle in F (existence of cycles), s.t. x possesses a path to it (existence of a path).

The paper is organized as follows:

- We provide the formal background in Section 2.
- We investigate the impact of removing (Section 3) and adding (Section 4) arguments on the given AF.
- In Section 5, we focus on syntactical information that can be inferred about the AF at hand.

Throughout the paper we distinguish semantics based on admissibility and naivity [van der Torre and Vesic, 2018], with a special focus on the former. In the last part of the paper we briefly discuss and compare our results and present pointers for future work. Due to the limited space we present particularly interesting proofs only.

2 Formal Background

Syntax. We fix a non-finite background set \mathcal{U} . An *argumentation framework* (AF) [Dung, 1995] is just a directed graph $F = (A, R)$ with a set of arguments $A \subseteq \mathcal{U}$ and the attack relation $R \subseteq A \times A$. If not stated otherwise we assume any AF to be finite and we use \mathcal{F} for the set of all these graphs. If $(a, b) \in R$ we say that a *attacks* b as well as a *attacks* (the set) E given that $b \in E \subseteq A$. This situation is denoted as $a \rightarrow b$ or $a \rightarrow E$, respectively. The meaning of $E \rightarrow a$ is analogous.

We frequently use the so-called *range* of a set E defined as $E^\oplus = E \cup E^+$ where $E^+ = \{a \in A \mid E \text{ attacks } a\}$. We also let $E^- = \{a \in A \mid a \text{ attacks } E\}$. The *E-reduct* of F is the AF $F^E = (E^*, R \cap (E^* \times E^*))$ where $E^* = A \setminus E^\oplus$. This means, F^E is the subframework of F obtained by removing the range of E . For an argument x , we use F_x for the induced AF without x , i.e. $F_x = (A \setminus \{x\}, R \cap (A \setminus \{x\} \times A \setminus \{x\}))$.

A sequence a_0, a_1, \dots, a_n with $a_i \rightarrow a_{i+1}$ for $0 \leq i < n$ and $a_i \neq a_j$ for $0 \leq i < j < n$ is called a *path*. If n is odd, then it is an *attack path*, otherwise we call it *support path*. If additionally $a_n = a_0$ we call the path a *cycle*. A cycle is odd if $|\{a_0, a_1, \dots, a_n\}|$ is odd, otherwise even. We say a *purely attacks* (purely supports) b if each path from a to b is an attack (support) path. Moreover, we say a is *controversial* w.r.t. b if there are both attack and support paths from a to b . We mention that it is intended that by the above definitions an argument a with no path at all to b is both a pure supporter and a pure attacker of it, but not controversial.

Semantics. A *semantics* is a function $\sigma : \mathcal{F} \rightarrow 2^{2^A}$ with $F = (A, R) \mapsto \sigma(F) \subseteq 2^A$. This means, a semantics returns a set of subsets of A , so-called σ -*extensions*. We say that an argument $a \in A$ is *credulously accepted* if $a \in \bigcup \sigma(F)$. Similarly, a is considered as *skeptically accepted* if $|\sigma(F)| \geq 1$ and $a \in \bigcap \sigma(F)$. In case of uniquely defined semantics, i.e. $|\sigma(F)| = 1$ for any F we may simply speak of *accepted* arguments as both notions coincide.

In this paper we consider so-called *naive*, *stage*, *admissible*, *complete*, *preferred*, *grounded* and *stable* semantics (abbr. *na*, *stg*, *ad*, *co*, *pr*, *gr*, *stb*) [Dung, 1995; Verheij, 1996]. All mentioned semantics satisfy conflict-freeness. A set $E \subseteq A$ is conflict-free in F (for short, $E \in cf(F)$) iff for no $a, b \in E$, $a \rightarrow b$. For the present paper it will be convenient to utilize the so-called *characteristic function* Γ_F : For $E \subseteq A$ we have $\Gamma_F(E) = \{a \in A \mid E \text{ defends } a\}$. A set E *defends* a if any attacker of a is attacked by some argument of E .

Definition 2.1. Let $F = (A, R)$ be an AF and $E \in cf(F)$.

1. $E \in na(F)$ iff E is \subseteq -maximal in $cf(F)$,
2. $E \in stg(F)$ iff $E \in na(F)$ and E^\oplus is \subseteq -maximal in $\{I^\oplus \mid I \in na(F)\}$,
3. $E \in ad(F)$ iff $E \subseteq \Gamma_F(E)$,
4. $E \in co(F)$ iff $E = \Gamma_F(E)$,
5. $E \in pr(F)$ iff E is \subseteq -maximal in $co(F)$,
6. $E \in gr(F)$ iff $E = \bigcup_{i \in \mathbb{N}} \Gamma_F^i(\emptyset)$,
7. $E \in stb(F)$ iff E attacks any $a \in A \setminus E$.

3 Acceptability and Removal of Arguments

In this section we investigate the impact of removing arguments from a given AF F on the acceptability status of former extensions or single arguments.

Preserving Extensions

Let us start with the following question:

Given some $E \in \sigma(F)$, under which conditions can $E \in \sigma(F_x)$ be guaranteed?


Admissibility-based Semantics

Our first results is regarding admissible and stable extensions. On the one hand, they are quite different: Stable extensions are among the most restricted ones in the literature, while admissibility is such a basic concept that a whole family of semantics exists on top of the admissibility requirement. However, they have something essential in common: Both require E to be conflict-free and that E^+ is sufficiently large. This common ground yields the following result.

Proposition 3.1. *Given an AF F , a semantics $\sigma \in \{ad, stb\}$ and $E \in \sigma(F)$. If $x \notin E$, then $E \in \sigma(F_x)$.*

Does the above assertion conveys to complete, preferred and grounded semantics? The following example provides us with a negative answer.

Example 3.2. Let F be as depicted below.



We have $gr(F) = co(F) = pr(F) = \{\emptyset\}$ and obviously $x \notin \emptyset$. The fact $gr(F_x) = co(F_x) = pr(F_x) = \{\{a\}\}$ proves that the empty set does not survive changing F to F_x .

However, if we require that “ x is an immediate attacker of E ” we obtain the persistence of E again. Note that this condition is indeed a strengthening of “not being contained in E ” since the latter is implied by the conflict-freeness of the considered semantics.

Theorem 3.3. *Given an AF F , $\sigma \in \{ad, stb, co, pr, gr\}$ and $E \in \sigma(F)$. If $x \rightarrow E$, then $E \in \sigma(F_x)$.*


Naivity-based Semantics.

We start by stating the counterpart to Proposition 3.1 in case of naive semantics.

Proposition 3.4. *Let F be an AF. If $E \in na(F)$ and $x \notin E$, then $E \in na(F_x)$.*

For stage semantics these assumptions are too weak as the following example shows.

Example 3.5. Consider the following AF F :



We have $stg(F) = \{\{a\}, \{b\}\}$ since the ranges $\{a\}^\oplus$ and $\{b\}^\oplus$ are incomparable, but $stg(F_x) = \{\{b\}\}$.

Indeed, we require stronger assumptions for stage semantics. The intuitive reason is that we mainly take the range of E into consideration and only check conflict-freeness of the extension itself. Hence if we make sure that removing x does not undermine E^\oplus , then E survives the transition to F_x .

Proposition 3.6. *Let F be an AF. If $E \in \text{stg}(F)$ and $x \notin E^\oplus$, then $E \in \text{stg}(F_x)$.*

The reader is presumably not surprised about the fact that we did not make any assumption about incoming attacks (cf. Propositions 3.4 and 3.6). The intuitive reason is that naivety-based semantics do not tend to take incoming attacks seriously.

Credulous Acceptance

Let us now focus on acceptability of single arguments, rather than whole extensions. More concisely:


Given some credulously accepted argument a in F , under which conditions is a still accepted in F_x ?

From an intuitive point of view, a support path from x to a should be beneficial for the acceptance of a whereas an attack path should be harmful. This means, deleting beneficial paths may cause non-acceptance and deleting harmful ones should not change the former acceptance as already proven for direct attackers in Theorem 3.3.

Admissibility-based Semantics

We start with showing that there is no analogy to Theorem 3.3 in case of pure attackers because they themselves could be part of a given extension E . Remember that for pure attackers each path to arguments in the set has to be odd.

Example 3.7. We find $E = \{x, f, a\} \in \text{pr}(F)$. Moreover, x purely attacks E .



However, obviously $E \notin \text{pr}(F_x)$ since x is not even contained in F_x . Please observe that a and f are still credulously accepted in F_x since $\{c, f, a\} \in \text{pr}(F_x)$.

The last observation can be formally proven for several admissibility-based semantics. This means, although whole extensions do not necessarily survive we may at least ensure credulous acceptance.

Proposition 3.8. *Given an AF F , $\sigma \in \{ad, co, pr\}$ as well as two arguments $a \in \bigcup \sigma(F)$ and $x \neq a$. If x purely attacks a , then $a \in \bigcup \sigma(F_x)$.*

Proof. It suffices to show the assertion for admissible sets since $\bigcup ad(G) = \bigcup pr(G) = \bigcup co(G)$ for any AF G .

Given $F = (A, R)$ and let $a \in E' \in ad(F)$. If $x \notin E'$, then Proposition 3.1 ensures the credulous acceptance of a in F_x since $E' \in ad(F_x)$. Consider now $x \in E'$. We inductively construct a set $E \subseteq E'$ with $E \in ad(F_x)$ and $a \in E$. We set:

- $E_0 = \{a\}$,
- $S_0 = \{d \in A \setminus E_0^+ \mid d \rightarrow E_0\}$ and
- $E_1 = (S_0^- \cap E') \setminus E_0$

Note that since E' is admissible in F we always find arguments in E' counter-attacking all arguments in S_0 . This means, $E_0 \dot{\cup} E_1$ defends a . In general, assume we are given an integer n as well as E_0, \dots, E_n . We proceed as follows:

- $S_n = \{d \in A \setminus (E_0 \dot{\cup} \dots \dot{\cup} E_n)^+ \mid d \rightarrow E_n\}$
- $E_{n+1} = (S_n^- \cap E') \setminus E_n$.

This way we obtain a sequence E_1, \dots, E_{n+1} of pairwise disjoint sets. We show that for each n

- i) each argument in E_{n+1} possesses a support path to a ,
- ii) $E_0 \dot{\cup} \dots \dot{\cup} E_{n+1}$ defends E_n , and
- iii) for any $0 \leq i < j \leq n$, $S_i \cap S_j = \emptyset$.

(base case) i) We have $E_0 = \{a\}$. ii) $E_0 \dot{\cup} E_1 = S_0^- \cap E'$ defends E_0 since E' is admissible. iii) Since $S_0 \subseteq (E_0 \dot{\cup} E_1)^+$, $S_1 \cap S_0 = \emptyset$ by construction.

For the inductive step let the assertion be true for some integer n .

i) If E_{n+1} is empty, we are done. Otherwise let $e \in E_{n+1}$. Then e attacks some $s \in S_n$ which in turn attacks some $e' \in E_n$. Since e' possesses some support path to a and all sets are disjoint, we found our support path from e to a .

ii) Due to admissibility of E' we have $S_{n+1}^- \cap E' = S_{n+1}^-$ and from $E_{n+2} = (S_{n+1}^- \cap E') \setminus E_{n+1}$ we hence infer that $E_1 \dot{\cup} \dots \dot{\cup} E_{n+2}$ defends E_{n+1} .

iii) Hence the set S_{n+1} of attackers of E_{n+1} not occurring in the range of $E_1 \dot{\cup} \dots \dot{\cup} E_{n+1}$ must be disjoint with S_i , $0 \leq i \leq n$.

This finishes our induction. Since x is a pure attacker of a we conclude $x \notin E_n$ for each n . We thus have for each integer n that 0) $\bigcup_{i=1}^n E_i \subseteq E' \in cf(F)$, 1) $x \notin \bigcup_{i=1}^n E_i$, and 2) $\bigcup_{i=1}^n E_i$ defends E_{n-1} . Due to finiteness, $\bigcup_{i \in \mathbb{N}} E_i$ is an admissible extension of F containing a , but not x . \square

Observe that the proof above is constructive. This means, it shows how to obtain a witnessing set for credulous acceptance as illustrated in the following example.

Example 3.9. In Example 3.7, $a \in \{x, f, a\} \in \text{pr}(F)$. Obviously, x purely attacks a , $a \in \bigcup ad(F)$ and $x \neq a$. So we find $E_0 = \{a\}$, $E_0^+ = \emptyset$ as well as $S_0 = \{b\}$. Since $S_0^- = \{c, f\}$ we get $E_1 = (S_0^- \cap \{x, f, a\}) \setminus E_0 = \{f\}$. Observe that the only attacker of E_1 is e and moreover, $e \in (E_0 \dot{\cup} E_1)^+ = \{a, f\}^+ = \{e, b\}$. This means we get $S_1 = \{b \in A \setminus (E_0 \dot{\cup} E_1)^+ \mid b \rightarrow E_1\} = \emptyset$ and $E_2 = \emptyset$, too. The procedure stops and outputs $E_0 \dot{\cup} E_1 = \{a, f\}$ and we verify that $a \in \{a, f\} \in ad(F_x)$ as required.

The attentive reader may have noted that so far, the case of grounded semantics is still missing.¹ Indeed, an analogous result can be shown here, but the proof needs to be adjusted since credulous acceptance for $\sigma = gr$ is more restrictive than in the previous cases.

Proposition 3.10. *Given an AF F as well as $a \in \bigcup gr(F)$ and $x \neq a$. If x purely attacks a , then $a \in \bigcup gr(F_x)$.*

¹Proposition 3.10 can also be inferred from [Cayrol *et al.*, 2010, Proposition 7]. However, we noted a sneaky error in the proof, which is about the existence of attack paths: If x is a pure attacker of a and y has support paths to a , then – although counter-intuitive at first glance – x is not necessarily a pure attacker of y . We thus decided to give a novel proof.

Proof. Let G be the grounded extension of F . If $x \notin G$, then the claim follows since we obtain $\Gamma_F^i(\emptyset) \subseteq \Gamma_{F_x}^i(\emptyset)$ for each integer $i \geq 1$.

If $x \in G$, we consider the procedure from the proof of Proposition 3.8. This way, we obtain an admissible extension E of F with $a \in E$ and $x \notin E$. Since $E \subseteq G$ and by definition of the grounded extension, E can be traced back to unattacked arguments, i.e. there is a sequence E_0, \dots, E_n s.t. i) $E = E_0 \cup \dots \cup E_n$, ii) E_0 is unattacked in F , and iii) E_{i-1} defends E_i for each $1 \leq i \leq n$. One can verify that these properties also hold in F_x and hence, $E \subseteq G_x$ for the grounded extension G_x of F_x . \square

Combining all results leads to the following final theorem regarding pure attackers.

Theorem 3.11. *Given an AF F , $\sigma \in \{ad, co, pr, gr\}$ as well as two arguments $a \in \bigcup \sigma(F)$ and $x \neq a$. If x purely attacks a , then $a \in \bigcup \sigma(F_x)$.*


This theorem in turn helps us to obtain the analogous result for pure supporters.

Theorem 3.12. *Given an AF F , $\sigma \in \{ad, co, pr, gr\}$ as well as two arguments $a \notin \bigcup \sigma(F)$ and $x \neq a$. If x purely supports a , then $a \notin \bigcup \sigma(F_x)$.*

Naivity-based Semantics

We start with mentioning that neither Theorem 3.11, nor Theorem 3.12 hold for stable semantics. This means, the deletion of pure attackers or pure supporters of an argument may indeed negate its former credulous acceptance or non-acceptance. The following example illustrates this assertion for pure attackers.

Example 3.13. Let F be as depicted below. We have that $stb(F) = \{\{a, x\}, \{b, e, x\}\}$ and x is a pure attacker of a .



However, $stb(F_x) = \{\{b, d\}\}$ showing that a is not credulously accepted anymore.

For stable semantics, it suffices to assume that a attacks x :

Proposition 3.14. *Given an AF F as well as $a \in \bigcup stb(F)$ and $x \neq a$. If $a \rightarrow x$, then $a \in \bigcup stb(F_x)$.*

For naive semantics, we find that a is accepted in F_x whenever x and a are different. As for stb this does not depend on the kind of paths from x to a in F .

Proposition 3.15. *Given an AF F as well as two arguments $x \neq a$. If $a \in \bigcup na(F)$, then $a \in \bigcup na(F_x)$.*

We also present some first result concerned with stage semantics. The assumptions we make are rather restrictive and closely related to Proposition 3.6. The reason is that it is hard to predict the structure of $stg(F)$ after removing a certain argument. Finding milder conditions ensuring an analogous result is left for future work.

Proposition 3.16. *Given an AF F and two arguments a and x . If there is some $E \in stg(F)$ with $a \in E$ and $x \notin E^\oplus$, then $a \in \bigcup \sigma(F_x)$.*


4 Acceptability and Addition of Arguments

Preserving Extensions

In the following we consider a subclass of *normal expansions* [Baumann and Brewka, 2010], so-called *x-expansions* firstly introduced in [Cayrol et al., 2010]. The former adds new arguments and attacks to F , but does not introduce novel attacks within F . The latter is a special case of normal expansions, where only one argument x is added to F .

Definition 4.1. The AF $G = (B, S)$ is a *x-expansion* of an AF $F = (A, R)$ whenever i) $B = A \cup \{x\}$ and ii) if $(a, b) \in S \setminus R$, then $a = x$ or $b = x$.

Example 4.2. Consider the x-expansion G of F . We have $ad(F) = \{\emptyset, \{c\}, \{c, a\}\}$. Hence, x does not attack any admissible set.



We observe that any admissible set survives since we obtain $ad(G) = \{\emptyset, \{c\}, \{c, a\}, \{x\}, \{x, c\}, \{x, c, a\}\}$.

The following proposition shows that the above observation holds in general.


Proposition 4.3. *Let G be an x-expansion of $F = (A, R)$ and $E \in ad(F)$. If $x \not\rightarrow E$, then $E \in ad(G)$.*

Example 4.2 reveals that the above assertion is not true for the remaining semantics. Indeed, $\{c, a\}$ is stable, preferred, complete and grounded in F , but not in G . Apart from grounded semantics we may state that the extension E survives if E attacks x .

Proposition 4.4. *Let G be an x-expansion of F , $E \in \sigma(F)$ with $\sigma \in \{ad, co, pr, stb\}$. If $E \rightarrow x$, then $E \in \sigma(G)$.*

This does not suffice for the grounded extension since x may ensure that G does not contain any unattacked arguments.

Example 4.5. The following AFs exemplify Proposition 4.4. However, for grounded semantics we have $gr(F) = \{\{c, a\}\}$ and $gr(G) = \{\emptyset\}$.



In fact, in order to ensure that the grounded extension survives the x-expansion we have to combine the requirements of Propositions 4.3, 4.4.

Proposition 4.6. *Let G be an x-expansion of $F = (A, R)$ and $E \in gr(F)$. If $x \not\rightarrow E$ and $E \rightarrow x$, then $E \in gr(G)$.*

Let us now turn to naivity-based semantics. If we have neither $x \rightarrow E$ nor $E \rightarrow x$ for some $E \in cf(F)$, then $E \cup \{x\}$ is conflict-free as well. Hence in this case E cannot be naive in the x-expansion G . Otherwise, we obtain the following:

Proposition 4.7. *Let G be an x-expansion of $F = (A, R)$ and $E \in na(F)$. If $x \rightarrow E$ or $E \rightarrow x$, then $E \in na(G)$.*

For stage semantics it might happen that F possesses no stable extension at all, and $stb(G) = stg(G) = \{x\}$. Hence it is clear that an analogous result for stage semantics can only be inferred under quite restrictive conditions; for example whenever x cannot be in any extension itself.

Proposition 4.8. Let G be an x -expansion of $F = (A, R)$ and $E \in \text{stg}(F)$. If $E \rightarrow x$ and $x \rightarrow x$, then $E \in \text{stg}(G)$.

Again, finding milder conditions ensuring an analogous result is left for future work.

Credulous Acceptance

In this section we will strongly benefit from previous results and show that the credulous (non-)acceptance of arguments is intimately linked to the addition of pure attackers and pure supporters. The decisive point is that by definition of an x -expansion G of F we have $G_x = F$.

Theorem 4.9. Given an AF $F = (A, R)$ and a semantics $\sigma \in \{ad, co, pr, gr\}$. Let $a \in A$ with $a \notin \bigcup \sigma(F)$. If G is an x -expansion of F , s.t. x purely attacks a , then $a \notin \bigcup \sigma(G)$.

Theorem 4.10. Given an AF $F = (A, R)$ and a semantics $\sigma \in \{ad, co, pr, gr\}$. Let $a \in A$ with $a \in \bigcup \sigma(F)$. If G is an x -expansion of F , s.t. x purely supports a , then $a \in \bigcup \sigma(G)$.

Regarding naivity-based semantics, the case $\sigma = na$ is based on the observation that any argument occurs in a naive extension iff it is no self-attacker.

Proposition 4.11. Given an AF $F = (A, R)$ and an argument $a \in A$ with $a \in \bigcup na(F)$. If G is an x -expansion of F , then $a \in \bigcup na(G)$.

Considering stage semantics, we require quite strong assumptions as it was the case in Proposition 4.6.

Proposition 4.12. Given an AF $F = (A, R)$ and $a \in A$ with $a \in \bigcup \text{stg}(F)$. If G is an x -expansion of F with $a \rightarrow x$ and $x \rightarrow x$, then $a \in \bigcup \text{stg}(G)$.

5 On Cycles, Circuits and Infinite Sequences

This section is dedicated to cycles in AFs.

Revisiting and Strengthening Former Results

In the argumentation community the following two results are widely known as well as frequently used in proofs.


Proposition 5.1. Any odd-cycle-free F possesses at least one stable extension.

Proposition 5.2. Any acyclic F has exactly one complete extension which is grounded, preferred and stable.

However, inspecting the seminal paper [Dung, 1995] reveals that neither of those assertions was presented by Dung. Even worse, both propositions are wrong if not restricted properly. For instance, Example 5.4 presents an acyclic (hence odd-cycle free) AF which does not possess any stable extension [Baumann and Spanring, 2015, Example 2]. Dung formulated Proposition 5.1 with “limited controversial” instead of “odd-cycle-free” [Dung, 1995, Corollary 36] and Proposition 5.2 with “well-founded” instead of “acyclic” [Dung, 1995, Theorem 30]. Both concepts are actually closer to the definition of a so-called *circuit* than the one of a cycle. The former generalizes the concept of a cycle by tolerating the repetition of vertices (not only the first/last vertex).

Definition 5.3. An AF F is *well-founded* if there is no infinite sequence a_0, a_1, \dots with $a_{i+1} \rightarrow a_i$ for each $i \geq 1$; F is *limited controversial* if there exists no infinite sequence a_0, a_1, \dots s.t. a_{i+1} is controversial wrt. a_i for each $i \geq 1$.

Example 5.4. Consider the following infinite AF F .



The AF is not well-founded since it possesses an infinite attack sequence, namely a_0, a_1, \dots . It is not limited controversial, either, because of the sequence a_0, a_2, a_4, \dots : We have that a_{2n+2} is controversial to a_{2n} witnessed by the support path $a_{2n+2}, a_{2n+1}, a_{2n}$ and the attack path a_{2n+2}, a_{2n} .

The following two Lemmata justify Propositions 5.1 and 5.2 as they allow to translate Dungs original versions into the usually quoted ones in the realm of finiteness.

Lemma 5.5. Let F be a finite AF. The following statements are equivalent: i) There is no cycle in F , ii) there is no circuit in F , iii) there is no even circuit in F , iv) F is well-founded.

Lemma 5.6. Let F be a finite AF. The following statements are equivalent: i) There is no odd cycle in F , ii) there is no odd circuit in F , iii) F is limited controversial.

We want to emphasize that Lemmata 5.5 and 5.6 do not hold for infinite AFs. Hence infinite AFs with no stable extension do not necessarily contain odd cycles. Indeed, the AF from Example 5.4 is acyclic, yet not limited controversial.

Dung’s controversiality notion is rarely used and most papers in the literature focus on finite AFs anyway. In this regard Dung’s proof seems unnecessarily complicated and we hence decided to provide a concise proof tailored for finite AFs only. The recently introduced *modularization* property plays the key role in this proof and underlines again its important role in the foundations of abstract argumentation [Baumann *et al.*, 2020a]. Besides this, our proof relies on the notion of directionality [Baroni and Giacomin, 2007] showing nicely how we can benefit from former foundational research here.

Proof of Proposition 5.1. The punchline is at follows: We show that if F is odd-cycle free, then it contains a non-empty admissible set. We may then argue as follows: There is a preferred extension $E \in pr(F)$. Assume that E is not stable in F . Thus, F^E is non-empty. Now, since the reduct contains a non-empty admissible extension E' , then $E \cup E' \in ad(F)$ by [Baumann *et al.*, 2020a, Proposition 3.4], contradiction.

We have thus left to prove the existence of such a non-empty admissible extension. This is equivalent to the existence of a non-empty complete one. Due to the directionality principle [Baroni and Giacomin, 2007, Section 3.3 and Table 1], it suffices to proof the claim for an AF F forming an SCC. Consider a circuit $C = (a_1, \dots, a_n)$ through all arguments in A and let $E = \bigcup a_{2n-1}$. Trivially, $E^\oplus = A$. Moreover, if E was conflicting we would immediately obtain an odd sub-circuit, yielding a contradiction via Lemma 5.6. Hence, $E \in \text{stb}(F) \subseteq co(F)$ concluding the proof. \square

The following theorem sharpens Dung’s result regarding the existence of stable extensions as it presents the very first characterization of this issue. This means, we provide a sufficient and necessary condition for the existence of a stable extension. For notational convenience we use $OC(F) = \{ \{a_0, a_1, \dots, a_n\} \mid a_0, a_1, \dots, a_n \text{ is an odd-cycle in } F \}$.

Theorem 5.7. *Given a finite AF F . $stb(F) \neq \emptyset$ if and only if $\exists E \in ad(F)$ s.t. $E^+ \cap C \neq \emptyset$ for each $C \in OC(F)$.*

Proof. (\Rightarrow) If $E \in stb(F)$, then $E^\oplus \cap C \neq \emptyset$ by definition for any $C \in OC(F)$. Moreover, by the definition of a cycle, $E^\oplus \cap C \neq \emptyset$ implies $E^+ \cap C \neq \emptyset$.

(\Leftarrow) Let $E \in ad(F)$ with $E^+ \cap C \neq \emptyset$ for each odd cycle C . Now the reduct F^E is odd-cycle free. There is thus a stable extension $E' \in stb(F^E)$ due to Proposition 5.1. Since each stable extension is trivially admissible, $E \cup E' \in ad(F)$ due to [Baumann *et al.*, 2020a, Proposition 3.4]. However, by assumption $(E \cup E')^\oplus = A$ and hence $E \cup E' \in stb(F)$. \square

We want to emphasize that the “ \Leftarrow ”-direction in Theorem 5.7 is the crucial one as it provides a sufficient condition for the existence of a stable extension. Roughly speaking, if we find an admissible extension which is capable of disrupting all odd cycles in F , then the AF possesses stable extensions.

Syntactical Implications

Derivable by the Number of Extensions

As recalled in the previous section, oftentimes semantical properties yield syntactical conditions a given AF must satisfy. In this section we return to our initial setting and investigate derivable connections between F and F_x . More specifically we address the following issue:

Given some connection between $|\sigma(F)|$ and $|\sigma(F_x)|$, which connections between x and cycles occurring in F can be guaranteed?

Most of the presented results utilize the following observation as important cornerstone (see also [Dunne and Bench-Capon, 2001, Theorem 8] and [Dvořák, 2012, Proposition 15]).

Lemma 5.8. *Let $\sigma \in \{co, pr, stb\}$ and $E, E' \in co(F)$. Then $E \neq E'$ iff there is an argument a contained in an even cycle with $a \in E$ and $a \notin E'$.*

We proceed with the main results of this section. Most of them utilize analogous proof techniques, with suitable adjustments tailored for the given situation. We thus decided to demonstrate the technique only once and omit the remaining proofs due to space restrictions.

Theorem 5.9. *Let $\sigma \in \{co, pr\}$. Let F be an AF with $|\sigma(F_x)| \neq |\sigma(F)|$. Then there is some path from x to an even cycle in F .*

Proof. Let $\sigma = co$ (the other case is similar). Assume there is no such path. We show $|co(F_x)| = |co(F)|$.

Given $E_x \in co(F_x)$, let us consider the whole AF F . We let $P \subseteq A$ be the arguments s.t. there is a path from a to P . We let C be the arguments contained in an even cycle. By assumption, all even cycles in F are contained in $A \setminus P$. Since there is no path from P to any argument in $A \setminus P$, $E_x \cap (A \setminus P)$ is admissible in F by the directionality principle [Baroni and Giacomin, 2007]. Hence there is some complete extension E of F with $E_x \cap (A \setminus P) \subseteq E$. Now E and E_x coincide on all arguments occurring in an even cycle. Since this uniquely determines $E \in co(F)$ (see Lemma 5.8) this induces a one-to-one mapping $co(F_x) \rightarrow co(F)$ by letting $E_x \mapsto E$; we get $|\sigma(F_x)| \leq |\sigma(F)|$. Starting the same line of reasoning with $E \in co(F)$ yields $|\sigma(F)| \leq |\sigma(F_x)|$. \square

While even cycles provide alternatives, odd ones oftentimes constrain the set of extensions. The following theorem formalizes that a connection to an odd cycle can be found, whenever x is responsible for rendering $\sigma(F)$ trivial.

Theorem 5.10. *If $ad(F) = \{\emptyset\}$ as well as $ad(F_x) \neq \{\emptyset\}$, then there is some path from an odd cycle to x in F .*

Observe that these results transfer to $\sigma \in \{co, pr\}$ since $ad(F) = \{\emptyset\}$ iff this is the case for $\sigma \in \{co, pr\}$.

Finally, stable extensions are quite sensitive to odd cycles as indicated by Proposition 5.1. The following can thus only guarantee the existence of any cycle, not necessarily even.

Theorem 5.11. *Let F be an AF with $|stb(F_x)| \neq |stb(F)|$. Then there is some path from x to a cycle in F .*

However, the result regarding degenerated AFs transfers to stb . Note that we move from $ad(F) = \{\emptyset\}$ to $stb(F) = \emptyset$.

Theorem 5.12. *If $stb(F) = \emptyset$ as well as $stb(F_x) \neq \{\emptyset\}$, then there is some path from x to an odd cycle in F .*

6 Summary and Related Work

In this paper, we systematically analyzed the impact of adding or removing attackers or supporters to the acceptance status of certain arguments in a given AF. In particular the part about adding arguments is closely related to [Cayrol *et al.*, 2010]. However, our investigation covers significantly more semantics as only grounded and preferred extensions are formally studied in [Cayrol *et al.*, 2010, Section 4]. The presented results concerning removal of arguments are similar to those shown in [Amgoud and Vesic, 2020] (see also [Bisquert *et al.*, 2011; Ulbricht and Baumann, 2019]). Our investigation is less general in the sense that fewer cases are considered (only credulous acceptance vs. no credulous acceptance), but more general in the sense that pure attackers are considered in contrast to immediate attackers only. It is hence an interesting future direction to study more cases of acceptance and rejection as done in [Amgoud and Vesic, 2020]. Some of our results from Section 3 are comparable to [Alfano *et al.*, 2017]. It presents a similar investigation in the realm of labeling-based semantics. It will be future work to compare these results in a comprehensive fashion. One of the central insights Sections 3 and 4 provide is the confirmation of the intuition that admissibility-based semantics tend to be sensible to the length of paths. In particular, it is essential whether it is even or odd. The same does not hold for naivity-based semantics confirmed by several formal results and provided counterexamples. Recently, so-called weak admissibility [Baumann *et al.*, 2020b] has been introduced. Weak admissibility establishes a third line of semantics beside admissibility-based and naivity-based ones. It would thus be exciting to see how the results from this paper transfer to these newly introduced semantics.

Acknowledgements

This work was supported by the German Research Foundation (DFG, BA 6170/2-1) and the German Federal Ministry of Education and Research (BMBF, 01/S18026A-F) by funding the competence center for Big Data and AI “ScaDS.AI” Dresden/Leipzig.

References

- [Alfano *et al.*, 2017] Gianvincenzo Alfano, Sergio Greco, and Francesco Parisi. Efficient computation of extensions for dynamic abstract argumentation frameworks: An incremental approach. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 49–55, 2017.
- [Amgoud and Vesic, 2020] Leila Amgoud and Srdjan Vesic. Dung’s semantics satisfy attack removal monotonicity. *CoRR*, abs/2007.04221, 2020.
- [Baroni and Giacomin, 2007] Pietro Baroni and Massimiliano Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence*, 171:675–700, 2007.
- [Baumann and Brewka, 2010] Ringo Baumann and Gerhard Brewka. Expanding argumentation frameworks: Enforcing and monotonicity results. In *Computational Models of Argument: Proceedings of COMMA 2010, Desenzano del Garda, Italy, September 8-10, 2010.*, pages 75–86, 2010.
- [Baumann and Spanring, 2015] Ringo Baumann and Christof Spanring. Infinite argumentation frameworks - on the existence and uniqueness of extensions. In *Essays Dedicated to Gerhard Brewka on the Occasion of His 60th Birthday*, volume 9060, pages 281–295. Springer, 2015.
- [Baumann *et al.*, 2020a] Ringo Baumann, Gerhard Brewka, and Markus Ulbricht. 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*, pages 79–88, 9 2020.
- [Baumann *et al.*, 2020b] Ringo Baumann, Gerhard Brewka, and Markus Ulbricht. 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*, pages 2742–2749. AAAI Press, 2020.
- [Baumann, 2012] Ringo Baumann. What does it take to enforce an argument? minimal change in abstract argumentation. In *ECAI 2012 - 20th European Conference on Artificial Intelligence. Including Prestigious Applications of Artificial Intelligence (PAIS-2012) System Demonstrations Track, Montpellier, France, August 27-31, 2012*, pages 127–132, 2012.
- [Bisquert *et al.*, 2011] Pierre Bisquert, Claudette Cayrol, Florence Dupin de Saint-Cyr, and Marie-Christine Lagasquie-Schiex. Change in argumentation systems: Exploring the interest of removing an argument. In *Scalable Uncertainty Management - 5th International Conference, SUM 2011, Dayton, OH, USA, October 10-13, 2011. Proceedings*, pages 275–288, 2011.
- [Boella *et al.*, 2009] Guido Boella, Souhila Kaci, and Leendert W. N. van der Torre. Dynamics in argumentation with single extensions: Abstraction principles and the grounded extension. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty, 10th European Conference, ECSQARU 2009, Verona, Italy, July 1-3, 2009. Proceedings*, pages 107–118, 2009.
- [Cayrol *et al.*, 2010] Claudette Cayrol, Florence Dupin de Saint-Cyr, and Marie-Christine Lagasquie-Schiex. Change in abstract argumentation frameworks: Adding an argument. *Journal of Artificial Intelligence Research*, 38:49–84, 2010.
- [Doutre and Mailly, 2018] Sylvie Doutre and Jean-Guy Mailly. Constraints and changes: A survey of abstract argumentation dynamics. *Argument & Computation*, 9(3):223–248, 2018.
- [Dung, 1995] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, 1995.
- [Dunne and Bench-Capon, 2001] Paul E Dunne and Trevor JM Bench-Capon. Complexity and combinatorial properties of argument systems. *University of Liverpool, Department of Computer Science (ULCS), Technical report*, 2001.
- [Dvořák, 2012] Wolfgang Dvořák. *Computational Aspects of Abstract Argumentation*. PhD thesis, Technische Universität Wien, 2012.
- [Liao *et al.*, 2011] B. Liao, L. Jin, and R.C. Koons. Dynamics of argumentation systems: A division-based method. *Artificial Intelligence*, pages 1790–1814, 2011.
- [Ulbricht and Baumann, 2019] Markus Ulbricht and Ringo Baumann. If nothing is accepted - repairing argumentation frameworks. *Journal of Artificial Intelligence Research*, 66:1099–1145, 2019.
- [van der Torre and Vesic, 2018] Leendert van der Torre and Srdjan Vesic. The principle-based approach to abstract argumentation semantics. In Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert van der Torre, editors, *Handbook of Formal Argumentation*. College Publications, February 2018.
- [Verheij, 1996] Bart Verheij. Two approaches to dialectical argumentation: Admissible sets and argumentation stages. In *Proceedings of the biannual International Conference on Formal and Applied Practical Reasoning (FAPR) workshop 1996*, pages 357–368. Universiteit, 1996.
- [Wallner *et al.*, 2017] Johannes Peter Wallner, Andreas Niskanen, and Matti Järvisalo. Complexity results and algorithms for extension enforcement in abstract argumentation. *Journal of Artificial Intelligence Research*, 60:1–40, 2017.
- [Wallner, 2020] Johannes Peter Wallner. Structural constraints for dynamic operators in abstract argumentation. *Argument & Computation*, 11(1-2):151–190, 2020.