Expanding Argumentation Frameworks: Enforcing and Monotonicity Results

Ringo Baumann and Gerhard Brewka
Universität Leipzig, Augustusplatz 10-11, 04109 Leipzig, Germany
$\{$*baumann, brewka*$\}$*@informatik.uni-leipzig.de*

Abstract. This paper addresses the problem of revising a Dung-style argumentation framework by adding finitely many *new* arguments which may interact with *old* ones. We study the behavior of the extensions of an argumentation framework if we extend it (new information) and/or change the underlying semantics (change of proof standards). We show both possibility and impossibility results related to the problem of enforcing a desired set of arguments. Furthermore we will prove some monotonicity results for a special class of expansions with respect to the cardinality of the set of extensions and the justification state.

Keywords. argumentation theory, belief revision, dynamics of argumentation

1. Introduction

Argumentation theory has become a popular research area in Artificial Intelligence (a very good overview is given in [1]). Argumentation frameworks (AFs), as introduced in Dung's seminal paper [2], are set-theoretically just directed graphs whose nodes are arguments and whose edges represent conflicts between them. There is a rich variety of semantics, which define the acceptable sets of arguments, so-called *extensions*. The motivations of these semantics range from the desired treatment of specific examples to fulfilling a number of abstract principles. A number of papers are engaged with the investigation of the properties and interrelations between the different semantics. Baroni/Giacomin [3] introduced several general criteria for comparing and evaluating semantics.

Dung's argumentation frameworks are static: they specify sets of acceptable arguments given a fixed set of arguments and attacks among them. Since argumentation is a dynamic process, it is natural to investigate the dynamic behavior of AFs. More precisely, we are interested in how extensions of an AF change when new arguments and/or attack relations are added. This task can be subsumed under the catchphrase *belief revision*[1] . In particular, we want to investigate whether, and if so how, it is possible to modify a given AF in such a way that a desired set of arguments becomes an extension. This question is certainly of great interest in a multi-agent scenario where an agent wants another agent to accept a particular set of arguments.

[1] Boella et. al. [4] called this task *argument-attack refinement.*

The modifications we allow include expansions of the AF which add arguments and possibly new attack relations, but we also allow changing the semantics. In this respect our approach differs from existing work (see Conclusion). Including also deletions of arguments and attacks is not overly interesting as this trivializes the problem (one could then just delete everything and add the wanted arguments without any attacks).

Our contributions to this topic are:

- impossibility results w.r.t. enforcing[2] a desired set of arguments
- possibility results for enforcing desired sets for specific semantics
- monotonicity results for weak expansions[3] concerning the number of extensions and the justification state of arguments

The paper is organized as follows: Section 2 reviews the necessary definitions at work in argumentation systems and introduces several kinds of expansions and enforcings. Section 3, the main part of this paper, contains the (im)possibility and monotonicity results. The last section discusses related results and conclusions.

2. Background

2.1. Argumentation Frameworks

Definition 1. An *argumentation framework* $\mathcal{A}$ is a pair (A, R), where A is a non-empty finite set whose elements are called *arguments* and $R \subseteq A \times A$ a binary relation, called *attack relation*.

If $(a, b) \in R$ holds we say that a attacks b, or b is attacked by a. In the following we consider a fixed countably infinite set $\mathcal{U}$ of arguments, called *universe*. Quantified formulae refer to this universe and all denoted sets are finite subsets of $\mathcal{U}$ or $\mathcal{U} \times \mathcal{U}$ respectively. We introduce some useful abbreviations.

Definition 2. Let $\mathcal{A} = (A, R)$ be an AF, B and B' subsets of A and $a \in A$. Then

1. $(B, B') \,\bar{\in}\, R \Leftrightarrow_{def} \exists b \exists b' : b \in B \wedge b' \in B' \wedge (b, b') \in R$,
2. B is unattacked in $\mathcal{A} \Leftrightarrow_{def} (A \backslash B, B) \,\bar{\notin}\, R$,
3. a is defended by B in $\mathcal{A} \Leftrightarrow_{def} \forall a' : a' \in A \wedge (a', a) \in R \rightarrow (B, \{a'\}) \,\bar{\in}\, R$,
4. B is conflict-free in $\mathcal{A} \Leftrightarrow_{def} (B, B) \,\bar{\notin}\, R$,
5. $\mathcal{A}_{\downarrow B}$ is the restriction of $\mathcal{A}$ to $B \Leftrightarrow_{def} \mathcal{A}_{\downarrow B} = (B, R \bigcap (B \times B))$.

2.2. Extension-Based Semantics

Semantics of argumentation frameworks specify certain conditions for selecting subsets of a given AF $\mathcal{A}$. The selected subsets are called *extensions*. The set of all extensions of $\mathcal{A}$ under semantics $\mathcal{S}$ is denoted by $\mathcal{E}_{\mathcal{S}}(A)$. We consider the classical (stable, preferred, complete, grounded [2]) and the ideal semantics [5].

[2] Enforcing means modifying an AF and/or changing the semantics with the result that a desired set becomes an extension (see Definition 6).

[3] A weak expansion adds further arguments which do not attack previous arguments (see Definition 5).

Definition 3. Let $\mathcal{A} = (A, R)$ be an AF and $E \subseteq A$. E is a

1. *stable extension* ($E \in \mathcal{E}_{st}(\mathcal{A})$) iff
   $(E, E) \notin R$ and for every $a \in A \backslash E$ holds: $(E, \{a\}) \bar{\in} R$;
2. *admissible extension*[4] ($E \in \mathcal{E}_{ad}(\mathcal{A})$) iff
   $(E, E) \notin R$ and each $a \in E$ is defended by E in $\mathcal{A}$,
3. *preferred extension* ($E \in \mathcal{E}_{pr}(\mathcal{A})$) iff
   $E \in \mathcal{E}_{ad}(\mathcal{A})$ and for each $E' \in \mathcal{E}_{ad}(\mathcal{A})$ holds: $E \not\subset E'$,
4. *complete extension* ($E \in \mathcal{E}_{co}(\mathcal{A})$) iff
   $E \in \mathcal{E}_{ad}(\mathcal{A})$ and for each $a \in A$ defended by E in $\mathcal{A}$ holds: $a \in E$,
5. *grounded extension* ($E \in \mathcal{E}_{gr}(\mathcal{A})$) iff
   $E \in \mathcal{E}_{co}(\mathcal{A})$ and for each $E' \in \mathcal{E}_{co}(\mathcal{A})$ holds: $E' \not\subset E$,
6. *ideal extension* of $\mathcal{A}$ ($E \in \mathcal{E}_{id}(\mathcal{A})$) iff
   $E \in \mathcal{E}_{ad}(\mathcal{A})$, $E \subseteq \bigcap_{P \in \mathcal{E}_{pr}(\mathcal{A})} P$ and for each $A \in \mathcal{E}_{ad}(\mathcal{A})$ w.t.p. $A \subseteq \bigcap_{P \in \mathcal{E}_{pr}(\mathcal{A})} P$ holds $E \not\subset A$

The defined semantics are only a subset of all existing semantics. There are several interrelations between them. Two important ones are $\mathcal{E}_{st}(\mathcal{A}) \subseteq \mathcal{E}_{pr}(\mathcal{A}) \subseteq \mathcal{E}_{co}(\mathcal{A}) \subseteq \mathcal{E}_{ad}(\mathcal{A})$ and $|\mathcal{E}_{gr}(\mathcal{A})| = |\mathcal{E}_{id}(\mathcal{A})| = 1$, i.e grounded and ideal semantics follow the unique-status approach [6]. The definitions of the semantics are guided by general principles [3]. Some of them will be introduced in the following. The set of all AFs $\mathcal{A}$ with at least one extension under semantics $\mathcal{S}$ and the set of all unattacked subsets in $\mathcal{A}$ are denoted as $\mathcal{D}_{\mathcal{S}}$ or $\mathcal{US}(\mathcal{A})$ respectively.

Definition 4. (abstract principles) A semantics $\mathcal{S}$ satisfies

1. *admissibility*,
2. *reinstatement*,
3. *conflict-freeness*,
4. *the directionality principle*

if and only if for any argumentation framework $\mathcal{A} \in \mathcal{D}_{\mathcal{S}}$, any extension $E \in \mathcal{E}_{\mathcal{S}}(\mathcal{A})$ and any unattacked set $U \in \mathcal{US}(\mathcal{A})$ it holds that:

1. $\forall a \ (a \in E \rightarrow \forall b \ (b \in A \wedge (b, a) \in R \rightarrow (E, \{b\}) \bar{\in} R))$,
2. $\forall a \ (\forall b \ (b \in A \wedge (b, a) \in R \rightarrow (E, \{b\}) \bar{\in} R) \rightarrow a \in E)$,
3. $(E, E) \notin R$,
4. $\mathcal{E}_{\mathcal{S}}(\mathcal{A}_{\downarrow U}) = \{(E \cap U) \,|\, E \in \mathcal{E}_{\mathcal{S}}(\mathcal{A})\}$.

Stable semantics does not fulfill directionality and the admissible semantics does not satisfy the reinstatement principle. Apart from that all considered semantics satisfy all mentioned principles. A good overview about these results is given in [7].

2.3. Expansions of AFs

Argumentation is a dynamic process. We focus in this paper on expansions $\mathcal{A}^*$ of an AF $\mathcal{A} = (A, R)$ where new arguments and attacks may be added but the attacks among the old arguments remain unchanged. We thus assume that - before adding new arguments

[4] Note that it is more common to speak about admissible sets instead of the admissible semantics. For reasons of unified notation we used the uncommon version.

- the attack relationships among arguments put forward earlier have been fully clarified and there is no further dispute concerning these relations. The following definition takes this idea into account.

Definition 5. An AF $\mathcal{A}^*$ is an *expansion* of AF $\mathcal{A} = (A, R)$ iff $\mathcal{A}^* = (A \cup A^*, R \cup R^*)$ for some nonempty A^* disjoint from A. An expansion is

1. *normal* ($\mathcal{A} \prec^N \mathcal{A}^*$) iff $\forall ab\ ((a,b) \in R^* \to a \in A^* \vee b \in A^*)$,
2. *strong* ($\mathcal{A} \prec^N_S \mathcal{A}^*$) iff $\mathcal{A} \prec^N \mathcal{A}^*$ and $\forall ab\ ((a,b) \in R^* \to \neg(a \in A \wedge b \in A^*))$,
3. *weak* ($\mathcal{A} \prec^N_W \mathcal{A}^*$) iff $\mathcal{A} \prec^N \mathcal{A}^*$ and $\forall ab\ ((a,b) \in R^* \to \neg(a \in A^* \wedge b \in A))$.

Normal expansions contain new arguments and possibly new attack relations. The latter have to involve at least one new argument. Strong and weak expansions[5] restrict the possible attacks between A and A^* to a single direction. The following figure illustrates a normal expansion. The dashed arrows are those that can be added.

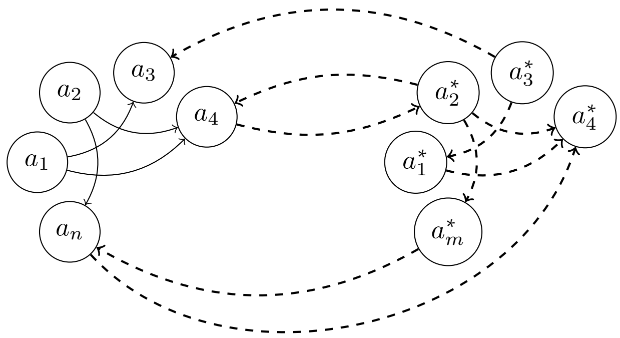

Figure 1. Normal Expansion of $\mathcal{A}$

2.4. Enforcing Extensions

Enforcing an extension E^* means modifying an argumentation framework in such a way that E^* becomes one of its extensions. The modifications we are interested in here are (a) normal expansions and (b) changes in the semantics. The former correspond to additional arguments that are brought into play together with the attack relations among them and older arguments. The latter can be viewed as a switch in the applied proof standard, for instance a switch from stable semantics to the more cautious grounded semantics.

Definition 6. Let $\mathcal{A} = (A, R)$ be an AF, $\mathcal{S}$ a semantics. Furthermore, let E^* be a *desired* set which is not an extension, i.e. $E^* \notin \mathcal{E}_\mathcal{S}(\mathcal{A})$. An $(\mathcal{A}, \mathcal{S})$*-enforcement* of E^* is a pair $\mathcal{F} = (\mathcal{A}^*, \mathcal{S}^*)$ such that (1) $\mathcal{A}^* = \mathcal{A}$ or $\mathcal{A} \prec^N \mathcal{A}^*$, and (2) $E^* \in \mathcal{E}_{\mathcal{S}^*}(\mathcal{A}^*)$. $\mathcal{F}$ is called

1. *conservative* if $\mathcal{S} = \mathcal{S}^*$,
2. *conservative strong* if $\mathcal{A} \prec^N_S \mathcal{A}^*$ and $\mathcal{S} = \mathcal{S}^*$,

[5]The terms are inspired by the fact that added arguments never receive attacks from previous arguments (*strong* arguments) or attack previous arguments (*weak* arguments), respectively.

3. *conservative weak* if $\mathcal{A} \prec_W^N \mathcal{A}^*$ and $\mathcal{S} = \mathcal{S}^*$,
4. *liberal* if $\mathcal{S} \neq \mathcal{S}^*$,
5. *liberal strong* if $\mathcal{A} \prec_S^N \mathcal{A}^*$ and $\mathcal{S} \neq \mathcal{S}^*$,
6. *liberal weak* if $\mathcal{A} \prec_W^N \mathcal{A}^*$ and $\mathcal{S} \neq \mathcal{S}^*$.

Whenever $\mathcal{A}$ and $\mathcal{S}$ are clear from context we simply speak of enforcements of E^*. We call the first three types *conservative* because the considered semantics remains constant. Almost all existing papers dealing with belief revision consider a fixed semantics. Liberal enforcements change the semantics. As mentioned earlier, this may be interpreted as a change of proof standard or paradigm shift. Imagine a judicial proceeding. It is vitally important whether you are accused on the base of criminal or civil law. The required evidence is different and hence the acceptable sets of arguments differ.

Note that more general modifications like leaving out previous attack relations or adding further attacks between previous arguments are excluded by our definition. If these types of manipulation are allowed the problem becomes trivial because one may add or delete arguments and attack relations at will.

To familiarize the reader with enforcements we give two examples.

Example 1. Let $\mathcal{A}$ be the following AF:

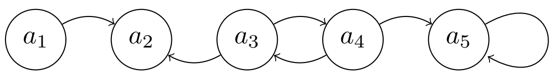

Let $\mathcal{S}$ be stable semantics and $E^* = \{a_1, a_3\}$ the *desired* set of arguments. Obviously we have $E^* \notin \mathcal{E}_{\mathcal{S}}(\mathcal{A})$. How to enforce E^*? Define a liberal enforcement with $\mathcal{A} = \mathcal{A}^*$ and $\mathcal{S}^* = pr$. One can check that E^* is indeed a preferred extension of $\mathcal{A}^*$.

Example 2. Let $\mathcal{A} = (A, R) = (\{a_1, a_2, a_3\}, \{(a_1, a_2), (a_2, a_1), (a_2, a_3)\})$ be an AF:

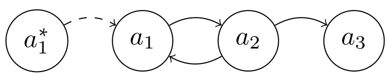

Let $\mathcal{S}$ be grounded semantics and $E^* = \{a_1^*, a_2\}$ the *desired* set of arguments. Obviously $E^* \notin \mathcal{E}_{\mathcal{S}}(\mathcal{A})$ since $a_1^* \notin A$. How to enforce E^*? Define a conservative enforcement with $\mathcal{A}^* = (A \cup \{a_1^*\}, R \cup \{(a_1^*, a_1)\})$. E^* is the grounded extension of $\mathcal{A}^*$.

3. Manipulating Argumentation Scenarios

The examples above illustrate the possibility to enforce extensions. This is not the case in general, i.e. not all *desired* sets are enforceable[6]. In the next two subsections we want to study necessary and sufficient conditions for the enforcement problem.

[6] Note that we consider conservative and liberal enforcements (compare definition 7). For arbitrary modifications (elimination of arguments, further attack relations between *old* arguments) of an AF $\mathcal{A}$ it is always possible.

3.1. *Impossibility Results (Enforcing of Extensions)*

We now show some useful interrelations between subsets of an AF, its normal expansions and abstract principles of semantics. The following properties are pretty obvious and hardly need any proof. Being aware of this fact, we still present them in the form of a proposition to be able to refer to them later on.

Proposition 1. *(simple properties) Let $\mathcal{A} = (A, R)$ be an AF and $\mathcal{S}$ a semantics.*

1. *If $\mathcal{S}$ satisfies admissibility and $A' \subseteq A$ does not defend all its elements in $\mathcal{A}$, then there is no conservative enforcement of A'.*
2. *If $\mathcal{S}$ satisfies reinstatement and $A' \subseteq A$ does not contain all defended elements in $\mathcal{A}$, then there is no conservative weak enforcement of A'.*
3. *If $\mathcal{S}$ satisfies conflict-freeness and $A' \subseteq A$ is conflicting, i.e. $(A', A') \bar{\in} R$, then there is no conservative enforcement of A'', if $A' \subseteq A''$ holds.*
4. *If $A' \subseteq A$ is unattacked in $\mathcal{A}$, i.e. $A' \in \mathcal{US}(\mathcal{A})$, then A' is unattacked in each weak expansion $\mathcal{A}^*$ of $\mathcal{A}$.*

Proof: Remember that normal expansions never delete previous arguments and never add or delete attacks between them. The equality-case ($\mathcal{A} = \mathcal{A}^*$) is obvious for all propositions.
1. given that A' does not defend all its elements we conclude $\exists ab\ (a \in A' \wedge b \in A \wedge (b, a) \in R \wedge (A', \{b\}) \bar{\notin} R)$; using $A' \subseteq A$, $b \in A$ and $\mathcal{A} \prec^N \mathcal{A}^*$, i.e. further attack relations have to involve at least one new argument, we conclude $(A', \{b\}) \bar{\notin} R \cup R^*$, hence $A' \notin \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)$;
2. given that A' does not contain all defended arguments in $\mathcal{A}$ we conclude $\exists a \in A \backslash A'$: $\forall b\ (b \in A \wedge (b, a) \in R \rightarrow (A', \{b\}) \bar{\in} R)$; considering that $\mathcal{A} \prec^N_W \mathcal{A}^*$ it follows that all possible attackers of a are already elements of A, hence A' does not contain all defended arguments in $\mathcal{A}^*$, i.e. $A' \notin \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)$;
3. obvious, supersets of conflicting sets are conflicting;
4. obvious, unattacked previous arguments stay unattacked if all additional attacks go to *new* elements. □

It is important to emphasize that these results are strict in the following sense: The conclusions of proposition 1.1, 1.2 and 1.4 are not necessarily true (or false) for supersets and subsets of A'. The same holds for subsets of conflicting sets. The following example illustrates the subset-case for proposition 1.1:

Example 3. Consider the following normal expansion $\mathcal{A}^*$ of $\mathcal{A} = (\{a_1, a_2\}, \{(a_1, a_2)\})$:

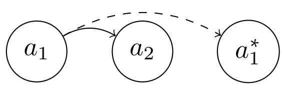

The set $A' = \{a_1, a_2\}$ does not defend a_2 in $\mathcal{A}$. Consider now the normal expansion $\mathcal{A}^* = (A \cup \{a_1^*\}, R \cup \{(a_1, a_1^*)\})$ of $\mathcal{A}$. On the one hand the set $\{a_1\} \subseteq A'$ is a stable[7] extension of $\mathcal{A}^*$, i.e. $\{a_1\} \in \mathcal{E}_{st}(\mathcal{A}^*)$ (existence of a conservative enforcement) and on the other hand the set $\{a_2\} \subseteq A'$ is itself a set, which does not defend all its elements in $\mathcal{A}$ (i.e., prop. 1.1 is applicable), hence there is no conservative enforcement of $\{a_2\}$.

[7] Stable semantics satisfies the admissibility principle.

Propositions 1.1, 1.2 and 1.3 present obvious necessary conditions for the enforcement problem[8]. The following impossibility theorems are more sophisticated than the mentioned proposition. The first one shows limitations for exchanging believed with unattacking arguments.

Theorem 2. *(conservative (liberal) enforcing [exchanging arguments]) Given an AF $\mathcal{A} = (A, R)$ and*

1. *a semantics $\mathcal{S}$ satisfying reinstatement,*
2. *a semantics $\mathcal{S}^*$, satisfying admissibility and conflict-freeness,*
3. *a set E such that $E \in \mathcal{E}_{\mathcal{S}}(\mathcal{A})$ and*
4. *two sets E', C such that $E' \subseteq E$, $C \subseteq A \backslash E$, $C \neq \emptyset$, $(C, A \backslash \{E' \cup C\}) \not\in R$ and $E^* := E' \cup C \notin \mathcal{E}_{\mathcal{S}}(\mathcal{A})$*

For all normal expansions $\mathcal{A}^*$ of $\mathcal{A}$ we have $E^* \notin \mathcal{E}_{\mathcal{S}^*}(\mathcal{A}^*)$.

Proof (reduction to the absurd) given $\mathcal{A}^* = (A \cup A^*, R \cup R^*)$ such that $\mathcal{A} \prec^N \mathcal{A}^*$ and $E^* \in \mathcal{E}_{\mathcal{S}^*}(\mathcal{A}^*)$; C is a nonempty subset of $A \backslash E$, hence $\exists c : c \in C \wedge c \notin E$; using the reinstatement property of $\mathcal{S}$ and $E \in \mathcal{E}_{\mathcal{S}}(\mathcal{A})$ we deduce $\exists b : b \in A \wedge (b, c) \in R \wedge (E, \{b\}) \not\in R$ (*); it obviously holds that $b \in A \cup A^* \wedge (b, c) \in R \cup R^*$, taking admissibility of $\mathcal{S}^*$ and $E^* \in \mathcal{E}_{\mathcal{S}^*}(\mathcal{A}^*)$ into account we conclude $(E^*, \{b\}) \in R \cup R^*$; this formula is fulfilled if $(E^*, \{b\}) \in R \vee (E^*, \{b\}) \in R^*$ holds; the latter disjunct is false because new attack relations have to involve at least one new argument; the former disjunct is equivalent to $(E', \{b\}) \in R \vee (C, \{b\}) \in R$; $(E', \{b\}) \in R$ contradicts $(E, \{b\}) \not\in R$ (compare (*)) because $E' \subseteq E$ holds; we deduce $b \notin E^*$ because $\mathcal{S}^*$ satisfies conflict-freeness and the assumption $E^* \in \mathcal{E}_{\mathcal{S}^*}(\mathcal{A}^*)$; hence $b \in A \backslash \{E' \cup C\}$, i.e. $(C, \{b\}) \not\in R$ because of the fourth premise $(C, A \backslash \{E' \cup C\}) \not\in R$. $\square$

Intuitively, the theorem says the following: if the involved semantics satisfy the specified properties, then it is impossible to find a normal expansion which possesses an extension E^* composed of a subset of an old extension E and some formerly unaccepted arguments C, given no element of C attacks some element which is not in the new extension.

The second impossibility theorem demonstrates limitations for eliminating arguments of existing extensions.

Theorem 3. *(conservative (liberal) weak enforcing [eliminating arguments]) Given an AF $\mathcal{A} = (A, R)$ and*

1. *a semantics $\mathcal{S}$, satisfying admissibility and conflict-freeness,*
2. *a semantics $\mathcal{S}^*$, satisfying reinstatement,*
3. *a set E such that $E \in \mathcal{E}_{\mathcal{S}}(\mathcal{A})$ and*
4. *a set C such that $C \subsetneq E$, $(C, A \backslash E) \not\in R$ and $E^* := E \backslash C \notin \mathcal{E}_{\mathcal{S}}(\mathcal{A})$*

For all weak expansions $\mathcal{A}^*$ of $\mathcal{A}$ we have $E^* \notin \mathcal{E}_{\mathcal{S}^*}(\mathcal{A}^*)$.

Proof (reduction to the absurd) Given $\mathcal{A}^* = (A \cup A^*, R \cup R^*)$ such that $\mathcal{A} \prec_W^N \mathcal{A}^*$ and $E^* \in \mathcal{E}_{\mathcal{S}^*}(\mathcal{A}^*)$; C is not empty because $E \in \mathcal{E}_{\mathcal{S}}(\mathcal{A})$ and $E^* \notin \mathcal{E}_{\mathcal{S}}(\mathcal{A})$ hold, hence we conclude $\exists c : c \in C \wedge c \notin E^*$; using the reinstatement property of $\mathcal{S}^*$ and $E^* \in \mathcal{E}_{\mathcal{S}^*}(\mathcal{A}^*)$

[8] Proposition 1.4 is not a conclusion about the set of extensions. It is essential to prove the monotonicity result (compare Theorem 4).

we deduce $\exists b : b \in A \cup A^* \wedge (b,c) \in R \cup R^* \wedge (E^*, \{b\}) \not\in R \cup R^*$; it obviously holds $(E^*, \{b\}) \not\in R$ (*) and due to the fact that $\mathcal{A}^*$ is a weak expansion we have $(b,c) \in R \wedge b \in A$; $c \in E$, admissibility of $\mathcal{S}$ and $E \in \mathcal{E}_{\mathcal{S}}(\mathcal{A})$ yield $(E, \{b\}) \in R$, hence $b \in A \backslash E$ (**) because of the conflict-freeness of $\mathcal{S}$; $(E, \{b\}) \in R$ is equivalent to $(E^*, \{b\}) \in R \vee (C, \{b\}) \in R$; $(E^*, \{b\}) \in R$ contradicts (*) and $(C, \{b\}) \in R$ contradicts (**), given the fourth premise $(C, A\backslash E) \not\in R$. □

Intuitively, the theorem says the following: if the involved semantics satisfy the specified properties, then it is impossible to find a weak expansion possessing an extension which is a proper subset of an old extension and not already an extension of the original AF, unless one of the arguments left out in the new extension attacks an element which was not in the old extension.

Due to the use of abstract principles, these theorems establish a connection between the set of extensions of the AF $\mathcal{A}$ under the semantics $\mathcal{S}$ and the set of extensions of the normal or weak expansion $\mathcal{A}^*$ under the semantics $\mathcal{S}^*$. They represent sufficient conditions for the impossibility to enforce a desired set.

3.2. Possibility Result (Enforcing of Extensions)

Consider the following simple AF $\mathcal{A}$:

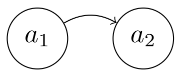

Proposition 1 (as well as Theorem 2) implies that $\{a_2\}$ never can be an extension of a normal expansion $\mathcal{A}^*$ of $\mathcal{A}$ under a semantics $\mathcal{S}$ satisfying admissibility. What about the following weaker claim: Is there a normal expansion $\mathcal{A}^*$ of $\mathcal{A}$ with an extension E^* such that a_2 is an element of E^*? In this subsection we want to prove the positive answer for all defined semantics. Furthermore we will prove the uniqueness or superset-property of this extension.

Theorem 4. *(conservative strong enforcing [$\sigma \in \{st, ad, pr, co, gr, id\}$]) Let $\mathcal{A} = (A, R)$ be an AF and $C \subseteq A$ a conflict-free set such that $C \notin \mathcal{E}_\sigma(\mathcal{A})$. There is a strong expansion $\mathcal{A}^*$ of $\mathcal{A}$ such that $C \subset E^*$ for some $E^* \in \mathcal{E}_\sigma(\mathcal{A}^*)$. Moreover, it is possible to construct $\mathcal{A}^*$ in such a way that (1) $|E^* \backslash C| = 1$ and (2) E^* is the unique extension of $\mathcal{A}^*$ (for $\sigma \in \{st, pr, co, gr, id\}$) or the set-inclusion maximal extension of $\mathcal{A}^*$ (for $\sigma = ad$).*

Proof: Without loss of generality we assume $A = \{a_1, ..., a_n\}$ and $C = \{a_1, ..., a_i\}$. Consider the strong expansion $\mathcal{A}^* = (A \cup \{a_1^*\}, R \cup \{(a_1^*, a_{i+1}), ..., (a_1^*, a_n)\})$ of $\mathcal{A}$ and let $E^* = C \cup \{a_1^*\}$. We show $E^* \in \mathcal{E}_\sigma(\mathcal{A}^*)$.
Given $C \notin \mathcal{E}_\sigma(\mathcal{A})$ we conclude $i < n$ (for all σ); furthermore $C \cup \{a_1^*\}$ is conflict-free in $\mathcal{A}^*$ by construction.
1. ($\sigma = st$) for every $a_j \in \{A \cup \{a_1^*\}\} \backslash \{C \cup \{a_1^*\}\}$ holds $(a_1^*, a_j) \in R^*$; **(uniqueness)** given $E^{**} \in \mathcal{E}_{st}(\mathcal{A}^*)$, therefore $a_1^* \in E^{**}$ holds because it is unattacked; hence for all $a_j \in \{a_{i+1}, ..., a_n\}$ holds $a_j \notin E^{**}$ (conflict); furthermore C has to be a subset of E^{**} because no element is attacked by a_1^* and they are pairwise conflict-free, i.e. $E^{**} = E^*$.
2. ($\sigma = ad$) each $a_j \in E^*$ is defended by E^* in $\mathcal{A}^*$ because every possible attacker in the complement $\{A \cup \{a_1^*\}\} \backslash \{C \cup \{a_1^*\}\}$ is counterattacked by a_1^*; **(super-set-property)**

given an extension $E^{**} \in \mathcal{E}_{ad}(\mathcal{A}^*)$; for every $a_j \in \{a_{i+1}, ..., a_n\}$ we have $a_j \notin E^{**}$ because there is no counterattack to a_1^*, hence $E^{**} \subseteq E^*$.
3. ($\sigma = pr$) admissibility is clear (compare 2.); maximality follows by the super-set-property of E^*; **(uniqueness)** given $E^{**} \in \mathcal{E}_{pr}(A^*)$, hence $E^{**} \in \mathcal{E}_{ad}(A^*)$; we conclude $E^{**} \subseteq E^*$ (super-set-property) and finally $E^{**} = E^*$ because it has to be maximal.
4. ($\sigma = co$) admissibility is clear (compare 2.); furthermore, for all $a_j \in \{a_{i+1}, ..., a_n\}$ (i.e. $a_j \notin E^*$) holds: a_j is not defended by $C \cup \{a_1^*\}$ in $\mathcal{A}^*$ because a_1^* is unattacked; **(uniqueness)** given $E^{**} \in \mathcal{E}_{pr}(A^*)$, hence $E^{**} \in \mathcal{E}_{ad}(A^*)$; we conclude $E^{**} \subseteq E^*$, assuming that $E^{**} \neq E^*$ yields a contradiction because a_1^* has to be in E^{**} (unattacked, thus defended by E^{**}) and all elements in C are defended by a_1^*, consequently by supersets of a_1^* (like E^{**}) too.
5. ($\sigma = gr$) $E^* \in \mathcal{E}_{co}(A^*)$ follows by the first part and minimality by the second part (uniqueness) in 4.; **(uniqueness)** obvious, because it is a unique-status approach.
6. ($\sigma = id$) admissibility is clear (compare 2.); furthermore, $E^* = \mathcal{E}_{pr}(A^*)$ (compare 3.), hence $E^* \subseteq \bigcap_{P \in \mathcal{E}_{pr}(\mathcal{A}^*)} P$ holds; in addition E^* is the maximal admissible set (super-set-property compare 2.), hence for each $A \in \mathcal{E}_{ad}(\mathcal{A}^*)$ such that $A \subseteq \bigcap_{P \in \mathcal{E}_{pr}(\mathcal{A}^*)} P$ holds $E^* \not\subset A$; **(uniqueness)** obvious, because it is a unique-status approach. □

The theorem shows that whenever a set C is conflict-free we may add one additional argument a and certain attacks so that the union of C and a is the unique extension of the constructed AF. It is important to emphasize that in special cases the enforcement of the desired set may be reached with less additional attack relations. Our construction shows the potential possibility only. To exemplify the standard construction consider the following AF.

Example 4. Let $\mathcal{A}^* = (A \cup \{a_1^*\}, R \cup \{(a_1^*, a_1), (a_1^*, a_3), (a_1^*, a_5)\})$ be the following strong expansion of $\mathcal{A} = (A, R)$:

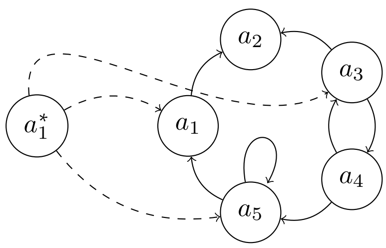

$C = \{a_2, a_4\}$ is the desired set of arguments. C is conflict-free and $C \notin \mathcal{E}_\sigma(\mathcal{A})$ holds for all $\sigma \in \{st, ad, pr, co, gr, id\}$, i.e. Theorem 3 is applicable. Hence, the standard-construction has the property that the set $E^* = \{a_2, a_4, a_1^*\}$ is the only extension of $\mathcal{A}^*$ **for all** semantics $\sigma \in \{st, pr, co, gr, id\}$[9].

3.3. Monotonic Relations between Sets of Extensions

Adding new arguments and their associated interactions obviously may change the outcome of an AF in a nonmonotonic way: arguments accepted earlier may become unac-

[9] For $\sigma = ad$ the super-set-property holds (compare 2.).

cepted, others become accepted; the number of extensions may shrink - or increase, depending on the new arguments. In this section we will identify special cases where arguments accepted earlier survive, and where the number of extensions can only increase. We will show that this holds for the class of weak expansions of an AF $\mathcal{A}$ and a semantics $\mathcal{S}$ satisfying the directionality principle. Here, every old belief set is contained in a new one and furthermore every new belief set is the union of an old one and a (possibly empty) set of new arguments.

Theorem 5. *(monotonicity) Given an AF $\mathcal{A} = (A, R)$ and a semantics $\mathcal{S}$ satisfying directionality, then for all weak expansions $\mathcal{A}^*$ of $\mathcal{A}$ the following holds:*

1. $|\mathcal{E}_{\mathcal{S}}(\mathcal{A})| \leq |\mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)|$,
2. $\forall E \in \mathcal{E}_{\mathcal{S}}(\mathcal{A})\ \exists E^* \in \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*) : E \subseteq E^*$ *and*
3. $\forall E^* \in \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)\ \exists E_i \in \mathcal{E}_{\mathcal{S}}(\mathcal{A})\ \exists A_i^* \subseteq A^* : E^* = E_i \bigcup A_i^*$

Proof: Given $\mathcal{A}^* = (A \cup A^*, R \cup R^*)$ such that $\mathcal{A} \prec_W^N \mathcal{A}^*$ and a semantics $\mathcal{S}$ satisfying the directionality principle; obviously we have $A \in \mathcal{US}(\mathcal{A}^*)$[10], hence the directionality principle yields $\mathcal{E}_{\mathcal{S}}(\mathcal{A}^*{}_{\downarrow A}) = \{(E^* \bigcap A) | E^* \in \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)\}$ (*); using that $\forall ab \in R^* : a \in A^* \vee b \in A^*$ holds, we conclude $\mathcal{A}^*{}_{\downarrow A} = \mathcal{A}$, hence $\mathcal{E}_{\mathcal{S}}(\mathcal{A}^*{}_{\downarrow A}) = \mathcal{E}_{\mathcal{S}}(\mathcal{A})$; (this is used for the following proofs)
1. (reduction to the absurd) Let $n, m \in \mathbf{N}$ and assume that $m = |\mathcal{E}_{\mathcal{S}}(\mathcal{A})| > |\mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)| = n$ holds, consequently $|\mathcal{E}_{\mathcal{S}}(\mathcal{A}^*{}_{\downarrow A})| = m$; on the other hand with $|\mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)| = n$ we get $|\{(E^* \bigcap A) | E^* \in \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)\}| \leq n$, i.e. $|\mathcal{E}_{\mathcal{S}}(\mathcal{A}^*{}_{\downarrow A})| \neq |\{(E^* \bigcap A) | E^* \in \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)\}|$, hence $\mathcal{E}_{\mathcal{S}}(\mathcal{A}^*{}_{\downarrow A}) \neq \{(E^* \bigcap A) | E^* \in \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)\}$ (contradicts (*)).
2. (reduction to the absurd) Assume that $\exists E \in \mathcal{E}_{\mathcal{S}}(\mathcal{A})\ \forall E^* \in \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*) : E \nsubseteq E^*$; i.e. $E \in \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*{}_{\downarrow A})$; we infer $E \nsubseteq E^* \bigcap A$ for all $E^* \in \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)$ from $E \nsubseteq E^*$, i.e. $E \neq E^* \bigcap A\ \forall E^* \in \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)$; hence $E \notin \{(E^* \bigcap A) | E^* \in \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)\}$ and consequently $\mathcal{E}_{\mathcal{S}}(\mathcal{A}^*{}_{\downarrow A}) \neq \{(E^* \bigcap A) | E^* \in \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)\}$ (contradicts (*)).
3. (reduction to the absurd) $\exists E^* \in \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)\ \forall E_i \in \mathcal{E}_{\mathcal{S}}(\mathcal{A})\ \forall A_i^* \subseteq A^* : E^* \neq E_i \bigcup A_i^*$; we deduce $E^* \bigcap A \neq E_i$ for all $E_i \in \mathcal{E}_{\mathcal{S}}(\mathcal{A})$ because assuming the existence of an E_i such that $E^* \bigcap A = E_i$ we may find an $A_i^* \subseteq A^*$ (set $A_i^* = E^* \backslash A$), so that $E^* = E_i \bigcup A_i^*$; hence $E^* \bigcap A \notin \mathcal{E}_{\mathcal{S}}(\mathcal{A})$ and consequently $\mathcal{E}_{\mathcal{S}}(\mathcal{A}^*{}_{\downarrow A}) \neq \{(E^* \bigcap A) | E^* \in \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)\}$ (contradicts (*)). □

The second and the third parts of Theorem 5 imply nice properties with respect to the justification state[11] of an argument a.

Corollary 6. *Given the same assumptions as in Theorem 5, then*

1. $\bigcup_{E \in \mathcal{E}_{\mathcal{S}}(\mathcal{A})} E \subseteq \bigcup_{E^* \in \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)} E^*$ *(credulously justified args persist),*

2. $\bigcap_{E \in \mathcal{E}_{\mathcal{S}}(\mathcal{A})} E \subseteq \bigcap_{E^* \in \mathcal{E}_{\mathcal{S}}(\mathcal{A}^*)} E^*$ *(skeptically justified args persist).*

[10] The general interrelation between unattacked sets of $\mathcal{A}$ and a weak expansion $\mathcal{A}^*$ of $\mathcal{A}$ is the following: $\forall U\ (U \in \mathcal{US}(\mathcal{A}) \rightarrow U \in \mathcal{US}(\mathcal{A}^*))$. Compare proposition 1.4.

[11] A very good overview about justification states is given in [1] (chapter 2.4).

That means, a skeptically (credulously) justified argument in $\mathcal{A}$ is skeptically (credulously) justified in $\mathcal{A}^*$. Remember that the grounded, complete, preferred and ideal semantics as well as CF2 [8] and the prudent version of the grounded semantics [9] satisfy the directionality principle. A weak expansion of an AF considered under any of these semantics behaves in a monotonic fashion with respect to the cardinality of the set of extensions as well as the mentioned subset-properties. Note that we cannot strengthen the cardinality statement in the sense of excluding equality.

Example 5. Consider the following two weak expansions $\mathcal{A}_1^*$, $\mathcal{A}_2^*$ of $\mathcal{A} = (\{a_1\}, \emptyset)$:

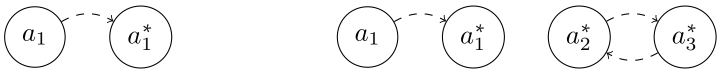

Let $\mathcal{S}$ be preferred semantics, hence $\mathcal{E}_{\mathcal{S}}(\mathcal{A}) = \{\{a_1\}\}$, i.e. $|\mathcal{E}_{\mathcal{S}}(\mathcal{A})| = 1$. The set of extensions of the first weak expansion stays equal ($\mathcal{E}_{\mathcal{S}}(\mathcal{A}_1^*) = \{\{a_1\}\}$) in contrast to the second weak expansion ($\mathcal{E}_{\mathcal{S}}(\mathcal{A}^*) = \{\{a_1, a_2^*\}, \{a_1, a_3^*\}\}$).

The monotonicity result is of interest as it provides a simplified method for checking whether an argument a is in some, respectively all extensions of an AF $\mathcal{A}$. First we introduce a new concept.

Definition 7. Let $C = \langle \mathcal{A}_0, \ldots, \mathcal{A}_n \rangle$ be a sequence of AFs, $\mathcal{A}$ an AF. C is called *expansion chain of* $\mathcal{A}$ iff

1. $\mathcal{A} = \mathcal{A}_n$ and
2. $\mathcal{A}_i \prec^N \mathcal{A}_{i+1}$ ($\mathcal{A}_{i+1}$ is a normal expansion of $\mathcal{A}_i$) for all i: $0 \leq i \leq n-1$.

C is called *weak* (resp. *strong*) if all expansions in the chain are weak (resp. strong).

Corollary 7. *Let $C = \langle \mathcal{A}_0, \ldots, \mathcal{A}_n \rangle$ be a weak expansion chain of $\mathcal{A}$, and let i be the smallest integer such that $\mathcal{A}$ covers a.*[12] *Given the same assumptions as in Theorem 5, we get: a is in some/all extensions of $\mathcal{A}$ iff a is in some/all extensions of $\mathcal{A}_i$.*

The corollary shows that it is sufficient to check the acceptability of an argument a in the chain-member $\mathcal{A}_i$ which is the first AF in which a appears. The following example illustrates this method.

Example 6. Acceptability of a_5 in the shaded AF decides its acceptability in the whole AF.

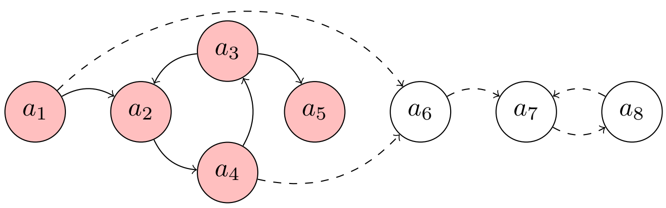

[12] $\mathcal{A} = (A, R)$ covers a whenever $a \in A$.

4. Related Work and Conclusions

We investigated a particular type of dynamics in Dung-style AFs. We provided theoretical insights about the impact of further arguments and attack relations. In particular, we investigated conditions for the (im)possibility to enforce a *desired* set of arguments. Moreover, we showed that the class of weak expansions behaves monotonically w.r.t. the cardinality of extensions and justification state of arguments if the considered semantics satisfies directionality. The use of abstract principles makes our results general enough to cover semantics which may be defined in the future.

There are several papers dealing with adding new information in argumentation. The closest to ours is [10]. Cayrol et al. proposed a typology of revisions (one new argument, one new interaction). Furthermore they proved sufficient conditions for being of a certain revision type. Other papers like [11] as well as [12] include additional knowledge from which arguments and attack relations can be justified. A further mentionable work in this context is [13]. Rotstein et al. define a warrant-prioritized revision operation which adds one additional argument that must be accepted afterwards. None of these papers establishes relationships between abstract properties of semantics and the (im)possibility of enforcements, as we do.

Future work may involve further abstract principles like I-maximality or several kinds of scepticism-adequacy [3]. Another direction is the investigation of similar properties for AFs if preferences are integrated.

References

[1] I. Rahwan and G. Simari (editors), "Argumentation in Artificial Intelligence" Springer 2009, Berlin, Germany

[2] P.M.Dung, "On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and $n-$person games", Artificial Intelligence 77 (1995), p. 321-357

[3] Baroni, P. and Giacomin, M. 2006. "Evaluation and comparison criteria for extension-based argumentation semantics", Proceedings of COMMA 2006 P. E. Dunne and T. J. Bench-Capon, Eds. Frontiers in Artificial Intelligence and Applications, vol. 144. IOS Press, Amsterdam, The Netherlands, p. 157-168

[4] Boella, G., Kaci, S. and van der Torre, L. 2009b. "Dynamics in argumentation with single extensions: attack refinement and the grounded extension" In Proc. AAMAS 2009, p. 1213-1214

[5] P.M. Dung, P. Mancarella, F. Toni, "A dialectic procedure for sceptical, assumption-based argumentation" in: P.E. Dunne, T. Bench-Capon (Eds.), Proceedings of the 1st COMMA 2006, IOS Press, Liverpool, UK, 2006, pp. 145-156

[6] H. Prakken, G.A.W. Vreeswijk, "Logics for defeasible argumentation" in: D.M. Gabbay, F. Guenthner (Eds.), Handbook of Philosophical Logic, second ed., Kluwer Academic Publishers, Dordrecht, 2001

[7] Baroni, P. and Giacomin, M. 2007, "On principle-based evaluation of extension-based argumentation semantics", Artif. Intell. 171, 10-15 (Jul. 2007), p. 675-700

[8] P. Baroni, M. Giacomin, "Solving semantic problems with odd-length cycles in argumentation" in: ECSQARU 2003, LNAI 2711, Springer-Verlag, Aalborg, Denmark, 2003, p. 440-451

[9] S. Coste-Marquis, C. Devred, P. Marquis, "Prudent semantics for argumentation frameworks, in: Proceedings of the 17th ICTAI 2005, IEEE Computer Society, Hong Kong, China, 2005, p. 568-572

[10] Cayrol, C., Dupin de Saint-Cyr, F. and Lagasquie-Schiex, M.- C. 2008. "Revision of an argumentation system", In Proc. KR 2008, p. 124-134

[11] Wassermann, R. 1999, "Full acceptance through argumentation - a preliminary report", in Proc. of IJCAI Workshop Practical Reasoning and Rationality

[12] Falappa, M., Garcia, A. and Simari, G. 2004, "Belief dynamics and defeasible argumentation in rational agents", in Proc. of NMR, p. 164-170

[13] Rotstein, N. D., Moguillansky, M. O., Falappa, M. A., Garcia, A. J. and Simari, G. R. 2008a, "Argument theory change: Revision upon warrant", in Proc. COMMA 2008, p. 336-347