Splitting Abstract Dialectical Frameworks

Thomas LINSBICHLER
Vienna University of Technology, Institute of Information Systems, Austria

Abstract. Among the abundance of generalizations of abstract argumentation frameworks, the formalism of abstract dialectical frameworks (ADFs) proved to be powerful in modelling various argumentation problems. Implementations of reasoning tasks that come within ADFs struggle with their high computational complexity. Thus methods simplifying the evaluation process are required. One such method is splitting, which was shown to be an effective optimization technique in other nonmonotonic formalisms. We apply this approach to ADFs by providing suitable techniques for directional splitting (allowing links only from the first to the second part of the splitting) under all the standard semantics of ADFs as well as preliminary results on general splitting.

Keywords. argumentation, abstract dialectical frameworks, splitting, optimization

1. Introduction

Among the study of argumentation in artificial intelligence [1,2], *abstract argumentation frameworks* (AFs) [3] have proven suitable to model and solve argumentation problems. Their limitations in terms of expressibility have been overcome by several enriched formalisms (e.g. [4,5,6]). In particular, *abstract dialectical frameworks* (ADFs) [7,8] constitute a very powerful generalization of AFs. In ADFs a propositional formula is assigned to each argument as acceptance condition instead of restricting the relation between arguments to simple attacks. This allows to capture several methods of preferential reasoning with ADFs such as [9,10]. Recently several new and alternative semantics for ADFs have been proposed [11,12].

Reasoning tasks within ADFs such as determining the results of applying a specific semantics are computationally hard (for an extensive complexity analysis we refer to [13]). This already suggests investigation of optimization techniques. In this work we follow the idea of *splitting* [14]: First, a given ADF is divided into several partitions. Then the results of the first part are computed followed by a transformation of the second part according to these results. After an iteration of this procedure through all parts the composition of all partial solutions gives the solution of the whole ADF. This divides one, possibly extensive, reasoning task into more but smaller subtasks and therefore reduces the computational effort of the overall process. This approach has already been studied for AFs [14,15] and logic programs [16], and turned out to achieve considerable optimization for AFs in an empirical evaluation [17]. Similar approaches were presented in [18] and [19]. An investigation on the application of splitting techniques under the ADF-semantics presented in [8] is important in various aspects:

- First, any implementation, such as [20], can make use of splitting results as a possible means to limit the search space of computation.

- Second, results on splitting ADFs are significant when dynamic aspects of argumentation come into play. Since argumentation is a dynamic process there is the need to handle situations where additional information is added or existing knowledge is revised. With splitting results at hand one has the possibility to reuse already computed (partial) results and only deal with the changed part of the framework.
- Finally, our results give some insights about how "local" semantics of ADFs are, that is whether the acceptance of an argument depends only on directly linked arguments or not. For AFs this is called directionality [21] and has been thoroughly investigated together with other basic principles of argumentation semantics [22].

The main contributions of the paper are presented as follows:

- Section 2 recalls the formal foundations of ADFs and completes the picture of relations between the semantics from [8].
- In Section 3 we provide positive results for *directional splitting* for all of these semantics. Directional splitting allows partitions of ADFs only in such a way that links appear only from one part to the other. Our results show that for all semantics directional splitting allows an incremental computation of the results of the semantics.
- Finally, Section 4 contains preliminary results on *general splitting*, allowing arbitrary partitioning of a given ADF. We show that under two-valued models, any ADF can be transformed to an equivalent ADF which then allows for directional splitting. For admissible interpretations we provide a procedure for general splitting.

2. Preliminaries

An *argumentation framework* (AF) is a directed graph where nodes represent arguments and a directed edge stands for an attack from the predecessor node to the successor node. The intended meaning of these attacks includes that an argument can only be accepted if all attacking arguments are not accepted. The strong limitation caused by this modeling choice is overcome by *abstract dialectical frameworks* (ADFs) by providing each argument with an acceptance condition. The acceptance of argument a is still based on the parents of a (denoted a_D^- in some ADF D). Now with ADFs, the acceptance condition of an argument a is a total function assigning to each subset of a_D^- one of the truth values $\mathbf{t}$ and $\mathbf{f}$. Given a subset $B \subseteq a_D^-$, the intended reading of $C_a(B) = \mathbf{t}$ (resp. $C_a(B) = \mathbf{f}$) is that a should be accepted (resp. should not be accepted) given that each argument in B is accepted and each argument in $(a_D^- \setminus B)$ is not. In the remainder of this section we recall the formal foundations of ADFs and its semantics, for more details we refer to [7,8].

Definition 1. An *abstract dialectical framework* is a tuple $D = (A, L, C)$ where

- A is a set of arguments (statements, positions),
- $L \subseteq A \times A$ is a set of links, and
- $C = \{C_a \mid a \in A\}$ is a set of total functions $C_a : 2^{a_D^-} \mapsto \{\mathbf{t}, \mathbf{f}\}$,

with $a_D^- = \{b \in A \mid (b, a) \in L\}$. C_a is called acceptance condition of a.

It seems natural to represent the acceptance condition of an argument a by a propositional formula over a_D^-. For the sake of convenience we will mostly use this propositional representation of ADFs.

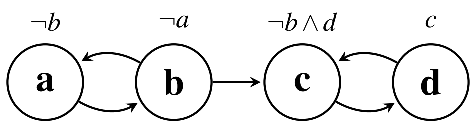

Figure 1. ADF D addressed in Examples 1 and 2.

Definition 2. In an ADF $D = (A, L, C)$, the propositional acceptance condition of argument $a \in A$, denoted as φ_a^D (or just φ_a if no ambiguity arises), is a formula of the form

$$\theta := x \in a_D^- \mid \top \mid \bot \mid \neg\theta \mid (\theta \wedge \theta) \mid (\theta \vee \theta) \mid (\theta \leftrightarrow \theta),$$

where each argument in a_D^- appears as atom in φ_a^D. By $\varphi_a^D[b/\theta : c(b)]$ we denote the replacement of each occurrence of atom b fulfilling condition $c(b)$ by a formula θ in φ_a^D.

Note that for any ADF $D = (A, L, C)$, the set of links L is immediate by the acceptance conditions $C = \{\varphi_a \mid a \in A\}$. The set of ingoing links a_D^- for each argument a coincides with the atoms in φ_a. Therefore we will frequently identify an ADF D by a set of pairs, where one pair $\langle a, \varphi_a \rangle$ stands for an argument together with its propositional acceptance condition. A_D is the set of all arguments of D and L_D will denote the set of links of D implicitly given by the pairs, i.e. $L_D = \{(x, a) \mid a \in A_D, x \text{ is an atom in } \varphi_a\}$.

Semantics of ADFs provide rules for the acceptance of arguments. Possible candidates for the outcome of applying a semantics to an ADF are interpretations.

Definition 3. Given an ADF $D = (A, L, C)$, a *two-valued interpretation* is a mapping $v : A \mapsto \{\mathbf{t}, \mathbf{f}\}$. For an interpretation v, the set $v^{\mathbf{t}} = \{a \in A \mid v(a) = \mathbf{t}\}$ denotes the unique *extension* associated with v.

Given an argument a, its propositional acceptance condition φ_a, and a two-valued interpretation v, $v(\varphi_a)$ is the result of the evaluation of φ_a under standard semantics of propositional logic, with the truth values of the propositional atoms given by v.

The semantics to be presented in Definitions 4 and 7 have been shown to be proper generalizations of AF-semantics [7,8].

Definition 4. Let D be an ADF. A two-valued interpretation is

- *conflict-free* iff for all $a \in v^{\mathbf{t}}$ it holds that $v(a) = v(\varphi_a)$ and
- a *two-valued model* of D iff for all $a \in A_D$ it holds that $v(a) = v(\varphi_a)$.

In the remainder of this paper we will use $cf(D)$ and $val_2(D)$ to denote the sets of conflict-free interpretations and two-valued models of ADF D, respectively.

Example 1. The ADF $D = \{\langle a, \neg b\rangle, \langle b, \neg a\rangle, \langle c, \neg b \wedge d\rangle, \langle d, c\rangle\}$ is depicted in Figure 1. The implicit links L_D are represented by directed edges. The extensions associated with the conflict-free interpretations of D are $\emptyset$, $\{a\}$, $\{b\}$, $\{c, d\}$, and $\{a, c, d\}$ and the extensions associated with the two-valued models of D are $\{a\}$, $\{b\}$, and $\{a, c, d\}$.

Most of the remaining semantics we consider in this paper abandon the need of a definite decision between acceptance and non-acceptance ($\mathbf{t}$ and $\mathbf{f}$) of an argument by being based on Kleene's strong three-valued logic [23].

Definition 5. Given an ADF D, a *(three-valued) interpretation* is a mapping $v : A_D \mapsto \{\mathbf{t},\mathbf{f},\mathbf{u}\}$. We denote $v^x = \{a \in A_D \mid v(a) = x\}$ for $x \in \{\mathbf{t},\mathbf{f},\mathbf{u}\}$.

A three-valued interpretation assigns a truth value true (**t**), false (**f**), or unknown (**u**) to each argument. The three truth values are partially ordered by $\leq_i$ according to their information content with the only pairs in $<_i$ being $\mathbf{u} <_i \mathbf{t}$ and $\mathbf{u} <_i \mathbf{f}$, following the intuition that true and false contain more information than unknown. Further we define the meet operation $\sqcap$ as $\mathbf{t} \sqcap \mathbf{t} = \mathbf{t}$, $\mathbf{f} \sqcap \mathbf{f} = \mathbf{f}$ and **u** otherwise, which can be read as the consensus of truth values.

This ordering can easily be transferred to interpretations. For (three-valued) interpretations v_1 and v_2 of some ADF D it holds that $v_1 \leq_i v_2$, i.e. v_2 contains at least as much information as v_1, if and only if for all $a \in A_D$, $v_1(a) \leq_i v_2(a)$. The meet operation $\sqcap$ is given by $(v_1 \sqcap v_2)(a) = v_1(a) \sqcap v_2(a)$ for all $a \in A_D$.

It is easy to see that each two-valued interpretation v is also a three-valued interpretation. They are the maximal elements with respect to $\leq_i$, while a three-valued interpretation mapping **u** to each argument is $\leq_i$-minimal. An interpretation w extends another interpretation v iff $v \leq_i w$. We denote the set of all two-valued interpretations extending some three-valued interpretation v by $[v]_2$. An element $w \in [v]_2$ maps each argument which is mapped to **u** by v to **t** or **f**.

The evaluation of propositional acceptance conditions under a three-valued interpretation follows the truth tables of the connectives given in Figure 2.

$\neg$	
t	f
f	t
u	u

$\vee$	t	u	f
t	t	t	t
u	t	u	u
f	t	u	f

$\wedge$	t	u	f
t	t	u	f
u	u	u	f
f	f	f	f

$\leftrightarrow$	t	u	f
t	t	u	f
u	u	u	u
f	f	u	t

Figure 2. Truth tables for the three-valued logic of Kleene.

Lemma 1. *Given a propositional acceptance condition θ and three-valued interpretations v_1 and v_2 thereof, it holds that if $v_1 \leq_i v_2$ then $v_1(\theta) \leq_i v_2(\theta)$.*

Proof. Let θ be a propositional acceptance condition and v_1, v_2 three-valued interpretations with $v_1 \leq_i v_2$, i.e. for each atom x in θ it holds that $v_1(x) \leq_i v_2(x)$. In the following let x and y be arbitrary atoms. If $\theta = y$ or $\theta = \neg y$, it is clear that $v_1(\theta) \leq_i v_2(\theta)$. Let $\theta = x \vee y$. If $v_1(x) = \mathbf{u}$ and $v_1(y) \in \{\mathbf{u},\mathbf{f}\}$, then $v_1(\theta) = \mathbf{u} \leq_i v_2(\theta)$. If $v_1(x) = \mathbf{u}$ and $v_1(y) = \mathbf{t}$, then $v_1(y) = v_2(y)$, hence $v_1(\theta) = v_2(\theta) = \mathbf{t}$. If $v_1(x), v_1(y) \in \{\mathbf{t},\mathbf{f}\}$, then necessarily $v_1 = v_2$ for x and y. The other cases follow by symmetry. Conjunction is dual to disjunction, hence $v_1(\theta) \leq_i v_2(\theta)$ for $\theta = x \wedge y$ follows. The result for general propositional acceptance conditions easily follows by induction. □

In order to introduce the remaining semantics we define the characteristic operator Γ_D they are based on:

Definition 6. Given an ADF D, the operator $\Gamma_D : (A_D \mapsto \{\mathbf{t},\mathbf{f},\mathbf{u}\}) \mapsto (A_D \mapsto \{\mathbf{t},\mathbf{f},\mathbf{u}\})$ maps three-valued interpretations to three-valued interpretations such that

$$\Gamma_D(v)(a) = \bigsqcap_{w \in [v]_2} w(\varphi_a).$$

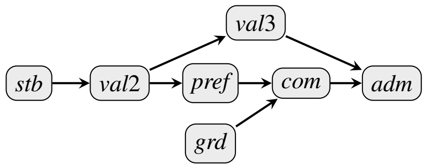

Figure 3. Relations between semantics. An arrow from σ to τ means that in any ADF D, $\sigma(D) \subseteq \tau(D)$.

The operator returns, for each argument a, the consensus truth value of the evaluation of the acceptance formula φ_a with each two-valued interpretation extending v. Moreover, Γ_D was shown to be $\leq_i$-monotone in [7]:

Lemma 2. *Given an* ADF *D and three-valued interpretations v_1 and v_2 thereof, it holds that if $v_1 \leq_i v_2$ then $\Gamma_D(v_1) \leq_i \Gamma_D(v_2)$.*

Definition 7. Given an ADF D, a three-valued interpretation v is

- a *three-valued model* of D iff for all $a \in A_D$, $v(a) \neq \mathbf{u}$ implies $v(a) = v(\varphi_a)$,
- an *admissible interpretation* of D iff $v \leq_i \Gamma_D(v)$,
- a *preferred interpretation* of D iff v is a $\leq_i$-maximal admissible interpretation of D,
- a *complete interpretation* of D iff $v = \Gamma_D(v)$,
- the *grounded interpretation* of D iff v is the least fixpoint of Γ_D wrt. $\leq_i$,
- a *stable model* of D iff v is a two-valued model of D and $v^{\mathbf{t}} = w^{\mathbf{t}}$, with w being the grounded interpretation of $D^{v-} = \{\langle a, \varphi_a[x/\bot : v(x) = \mathbf{f}]\rangle \mid a \in v^{\mathbf{t}}\}$.

We will denote the three-valued models, admissible, preferred, complete, grounded interpretations, and stable models of an ADF D as $val_3(D)$, $adm(D)$, $pref(D)$, $com(D)$, $grd(D)$, and $stb(D)$, respectively.

Example 2. Again consider the ADF D depicted in Figure 1. The admissible interpretations $\{v_1, \dots, v_8\}$ coincide with the three-valued models: $v_1 = \{a \mapsto \mathbf{u}, b \mapsto \mathbf{u}, c \mapsto \mathbf{u}, d \mapsto \mathbf{u}\}$, $v_2 = \{a \mapsto \mathbf{t}, b \mapsto \mathbf{f}, c \mapsto \mathbf{u}, d \mapsto \mathbf{u}\}$, $v_3 = \{a \mapsto \mathbf{t}, b \mapsto \mathbf{f}, c \mapsto \mathbf{f}, d \mapsto \mathbf{f}\}$, $v_4 = \{a \mapsto \mathbf{t}, b \mapsto \mathbf{f}, c \mapsto \mathbf{t}, d \mapsto \mathbf{t}\}$, $v_5 = \{a \mapsto \mathbf{f}, b \mapsto \mathbf{t}, c \mapsto \mathbf{u}, d \mapsto \mathbf{u}\}$, $v_6 = \{a \mapsto \mathbf{f}, b \mapsto \mathbf{t}, c \mapsto \mathbf{f}, d \mapsto \mathbf{u}\}$, $v_7 = \{a \mapsto \mathbf{f}, b \mapsto \mathbf{t}, c \mapsto \mathbf{f}, d \mapsto \mathbf{f}\}$, $v_8 = \{a \mapsto \mathbf{u}, b \mapsto \mathbf{u}, c \mapsto \mathbf{f}, d \mapsto \mathbf{f}\}$. Further observe that $com(D) = \{v_1, v_2, v_3, v_4, v_7, v_8\}$, $pref(D) = \{v_3, v_4, v_7\}$, $grd(D) = \{v_1\}$, and $stb(D) = \{v_3, v_7\}$.

Proposition 3. *The relations between semantics depicted in Figure 3 hold.*

Proof. We show that, for any ADF each three-valued model is an admissible interpretation. The other relations were shown in [8]. To this end consider an arbitrary ADF D and let $v \in val_3(D)$ and $a \in A_D$. If $v(a) = \mathbf{u}$ obviously $v(a) \leq_i \Gamma_D(v)(a)$. So let $v(a) \neq \mathbf{u}$. By assumption, $v(a) = v(\varphi_a)$. Now note that for each $w \in [v]_2$ it holds that $v \leq_i w$. Hence we know by Lemma 1 that $v(\varphi_a) \leq_i w(\varphi_a)$. Therefore $v(\varphi_a) \leq_i \bigsqcap_{w \in [v]_2} w(\varphi_a)$ and consequently $v(a) \leq_i \Gamma_D(v)(a)$, showing that v is an admissible interpretation. □

In order to present our results we will make use of the following technical concepts: Given an ADF D and a set of arguments $B \subseteq A_D$, the restriction of D to B is defined as $D|_B = \{\langle a, \varphi_a \rangle \in D \mid a \in B\}$. Note that $D|_B$ is an ADF iff $\forall b \in B : b_D^- \subseteq B$. Further for interpretations v_1 over A_1 and v_2 over A_2, $v_1|_B$ is the restriction of v_1 to the arguments $B \cap A_1$ and the union $v_1 \cup v_2$ is an interpretation over $(A_1 \cup A_2)$ iff $(A_1 \cap A_2) = \emptyset$.

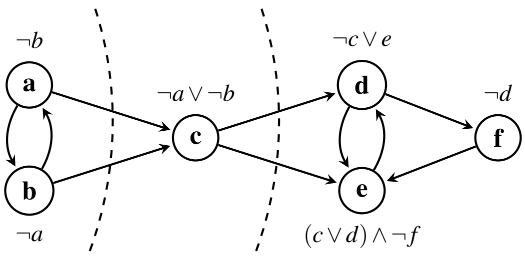

Figure 4. Possible directional splittings of the depicted ADF.

3. Directional Splitting

The definition of the various semantics of ADFs already suggests that not every decomposition of an ADF can be treated equivalently. In this section we focus on *directional splitting*, that is, given an ADF D, a partition of the graph underlying D into two disjoint subgraphs G_1 and G_2 such that there are no links from G_2 to G_1 in D. In other words it is splitting "along the lines" of the strongly connected components of an ADF.

Definition 8. Let $G_1 = (A_1, L_1)$ and $G_2 = (A_2, L_2)$ be directed graphs such that $A_1 \cap A_2 = \emptyset$. Moreover let $L_3 \subseteq A_1 \times A_2$. We call the tuple (G_1, G_2, L_3) a directional splitting of an ADF $D = (A_1 \cup A_2, L_1 \cup L_2 \cup L_3, C)$.

Figure 4 illustrates the two possible directional splittings of an exemplary ADF. Note that any other splitting of this frameworks contains links in both directions between the subgraphs and is therefore not directional.

Definition 9. Let $G_1 = (A_1, L_1)$ and $G_2 = (A_2, L_2)$ be directed graphs such that (G_1, G_2, L_3) is a directional splitting of the ADF D. Further let v be an interpretation of $D|_{A_1}$. The *v-reduct* of D is defined as

$$D^v = \{\langle a, \varphi_a[b/v(b) : b \in (v^{\mathbf{t}} \cup v^{\mathbf{f}})][c/x_c : c \in v^{\mathbf{u}}]\rangle \mid a \in A_2\} \ \cup \ \{\langle x_c, \neg x_c\rangle \mid c \in v^{\mathbf{u}}\},$$

where x_c is a newly introduced argument for each $c \in v^{\mathbf{u}}$.

The idea of directional splitting is to propagate truth values assigned to arguments by an interpretation of the first part along the links to the second part. An example is illustrated in Figure 5. Here the dotted borderline on the left suggests the splitting (G_1, G_2, L_3) of ADF D with $L_3 = \{(b,d),(c,d)\}$. The right part depicts the resulting ADFs D_1 and D^{v_1}, where $v_1 = \{a \mapsto \mathbf{t}, b \mapsto \mathbf{u}, c \mapsto \mathbf{f}\}$ is a possible (complete) interpretation of D_1. In the acceptance condition φ_d the atom c ($v_1(c) = \mathbf{f}$) is replaced by the propositional constant $\bot$ and b ($v_1(b) = \mathbf{u}$) by the newly introduced argument x_b.

With this at hand, we are ready to formulate our first results on directional splitting.

Theorem 4. *Let $\sigma \in \{cf, val_2\}$, and $G_1 = (A_1, L_1)$ and $G_2 = (A_2, L_2)$ be directed graphs such that (G_1, G_2, L_3) is a directional splitting of the* ADF *D. Further let $D_1 = D|_{A_1}$. It holds that*

1. $v_1 \in \sigma(D_1) \wedge v_2 \in \sigma(D^{v_1}) \Rightarrow (v_1 \cup v_2) \in \sigma(D)$,
2. $v \in \sigma(D) \Rightarrow v|_{A_1} \in \sigma(D_1) \wedge v|_{A_2} \in \sigma(D^{v|_{A_1}})$.

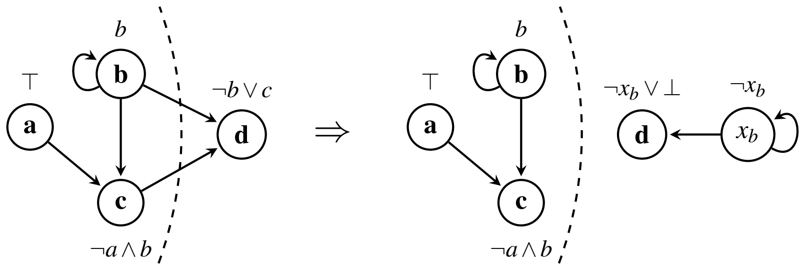

Figure 5. Directional splitting of the ADF D on the left into the ADFs $D|_{A_1}$ and D^{v_1}, where $A_1 = \{a,b,c\}$ and $v_1 = \{a \mapsto \mathbf{t}, b \mapsto \mathbf{u}, c \mapsto \mathbf{f}\}$ is a complete interpretation of D_1.

Proof. Note that for a two-valued interpretation v the v-reduct does not contain additional arguments but is built by solely replacing atoms by truth values in the acceptance conditions, i.e. $D^v = \{\langle a, \varphi_a[b/v(b) : b \in A_1]\rangle \mid a \in A_2\}$. Hence the result follows directly from the fact that for any acceptance condition φ_a over atoms B and any two-valued interpretation v, it holds that $v(\varphi_a) = v|_{B'}(\varphi_a[b/v(b) : b \in (B \setminus B')]$, where $B' \subseteq B$. □

The following lemma will be useful to prove splitting results for the other semantics.

Lemma 5. *Let $G_1 = (A_1, L_1)$ and $G_2 = (A_2, L_2)$ be directed graphs such that (G_1, G_2, L_3) is a directional splitting of the* ADF *D, $A = A_1 \cup A_2$ and $D_1 = D|_{A_1}$. The following holds:*

1. *If v_1 is an interpretation of D_1 and v_2 is an admissible interpretation of D^{v_1}, then*
$$(\Gamma_{D_1}(v_1) \cup \Gamma_{D^{v_1}}(v_2))|_A = \Gamma_D\left((v_1 \cup v_2)|_A\right).$$

2. *If v is an interpretation of D, then*
$$\Gamma_D(v) = \Gamma_{D_1}(v|_{A_1}) \cup \left(\Gamma_{D^{v|_{A_1}}}(v')\right)\Big|_{A_2}$$
where $v' = v|_{A_2} \cup \{x_c \mapsto \mathbf{u} \mid c \in (v|_{A_1})^{\mathbf{u}}\}$.

Proof. 1) We need to show that for each $a \in A_1$, $\Gamma_{D_1}(v_1)(a) = \Gamma_D((v_1 \cup v_2)|_A)(a)$ and for each $a \in A_2$, $\Gamma_{D^{v_1}}(v_2)(a) = \Gamma_D((v_1 \cup v_2)|_A)(a)$. The former is trivial by the fact that for each $a \in A_1$, $\varphi_a^D = \varphi_a^{D_1}$ and $a_D^- \subseteq A_1$. The latter is by the following chain of equalities, letting $a \in A_2$ and $v_1^* = v_1|_{((v_1)^{\mathbf{t}} \cup (v_1)^{\mathbf{f}})}$:

$$\Gamma_{D^{v_1}}(v_2)(a) = \bigsqcap_{w \in [v_2]_2} w(\varphi_a[b/v_1(b) : b \in ((v_1)^{\mathbf{t}} \cup (v_1)^{\mathbf{f}})][c/x_c : c \in (v_1)^{\mathbf{u}}]) \tag{1}$$

$$= \bigsqcap_{w \in [v_2 \cup v_1^*]_2} w(\varphi_a[c/x_c : c \in (v_1)^{\mathbf{u}}]) \tag{2}$$

$$= \bigsqcap_{w \in [v_2 \cup v_1]_2} w(\varphi_a) \tag{3}$$

$$= \bigsqcap_{w \in [(v_1 \cup v_2)|_A]_2} w(\varphi_a) \tag{4}$$

$$= \Gamma_D((v_1 \cup v_2)|_A)(a) \tag{5}$$

Note that (2) = (3) is by the fact that since $\varphi_{x_c} = \neg x_c$ and $v_2 \in adm(D^{v_1})$ necessarily $v_2(x_c) = \mathbf{u}$ and by definition of D^{v_1}, $v_1(c) = \mathbf{u}$. (3) = (4) is by $a_D^- \subseteq A$ for any $a \in A_2$.
2) Again, for $a \in A_1$, $\Gamma_D(v)(a) = \Gamma_{D_1}(v|_{A_1})(a)$ is trivial. Let $a \in A_2$ and consider $v' = v|_{A_2} \cup \{x_c \mapsto \mathbf{u} \mid c \in (v|_{A_1})^{\mathbf{u}}\}$. We get $\Gamma_D(v)(a) = \bigsqcap_{w \in [v']_2} w(\varphi_a[b/v(b) : b \in (v|_{A_1})^{\mathbf{t}} \cup (v|_{A_1})^{\mathbf{f}}][c/x_c : c \in (v|_{A_1})^{\mathbf{u}}]) = \Gamma_{D^{v|_{A_1}}}(v')(a)$ by substituting **t**-/**f**-values in $v|_{A_1}$ directly into φ_a and replacing atoms with **u**-values by other atoms which are **u** in v' by definition. □

In the general case of three-valued interpretations the reduct involves the introduction of additional arguments, hence the following equalities only hold under projection.

Theorem 6. *Let $\sigma \in \{val_3, adm, pref, com, grd\}$, and $G_1 = (A_1, L_1)$ and $G_2 = (A_2, L_2)$ be directed graphs such that (G_1, G_2, L_3) is a directional splitting of the* ADF *D. Further let $D_1 = D|_{A_1}$. It holds that*

1. $v_1 \in \sigma(D_1) \wedge v_2 \in \sigma(D^{v_1}) \Rightarrow (v_1 \cup v_2)|_A \in \sigma(D)$,
2. $v \in \sigma(D) \Rightarrow v|_{A_1} \in \sigma(D_1) \wedge \exists v_2 \in \sigma(D^{v|_{A_1}}) : v_2|_{A_2} = v|_{A_2}$.

Proof. 1) val_3: Let $v_1 \in val_3(D_1)$ and $v_2 \in val_3(D^{v_1})$. We have to show $(v_1 \cup v_2)|_A(a) = (v_1 \cup v_2)|_A(\varphi_a)$ for all $a \in A_D$. If $a \in A_1$ this follows from $a_D^- \subseteq A_1$. Let $a \in A_2$ with $v_2(a) \neq \mathbf{u}$. We know that $v_2(a) = (v_1 \cup v_2)|_A(a) = v_2(\varphi_a[b/v_1(b) : b \in ((v_1)^{\mathbf{t}} \cup (v_1)^{\mathbf{f}})][c/x_c : c \in (v_1)^{\mathbf{u}}])$. Since $\varphi_{x_c}^{D^{v_1}} = \neg x_c$ for all $c \in (v_1)^{\mathbf{u}}$, necessarily $v_2(x_c) = \mathbf{u}$. Hence $(v_1 \cup v_2)|_A(a) = (v_1 \cup v_2)(\varphi_a^D)$.
$adm, pref, com, grd$: Since in any case, v_2 is an admissible interpretation of D^{v_1} (cf. Proposition 3), the result follows directly by Lemma 5.1.

2) val_3: Let $v \in val_3(D)$. $v|_{A_1} \in val_3(D_1)$ is clear. Now let $v_2 = v|_{A_2} \cup \{x_c \mapsto \mathbf{u} \mid c \in (v|_{A_1})^{\mathbf{u}}\}$. We have to show that $v_2 \in val_3(D^{v|_{A_1}})$. For all x_c where $c \in (v|_{A_1})^{\mathbf{u}}$, $v_2(x_c) = v_2(\varphi_{x_c}^{D^{v|_{A_1}}}) = \mathbf{u}$ holds since $\varphi_{x_c}^{D^{v|_{A_1}}} = \neg x_c$. Finally consider $a \in A_2$. It holds that $v(a) = v(\varphi_a^D)$. Since $v_2(x_c) = \mathbf{u}$ also $v(a) = v_2(\varphi_a^D[b/v|_{A_1}(b) : b \in ((v|_{A_1})^{\mathbf{t}} \cup (v|_{A_1})^{\mathbf{f}})][c/x_c : c \in (v|_{A_1})^{\mathbf{u}}])$, i.e. $v(a) = v_2(\varphi_a^{D^{v|_{A_1}}})$. As $v(a) = v|_{A_2}(a) = v_2(a)$ holds, $v_2 \in val_3(D^{v|_{A_1}})$.
$adm, pref, com, grd$: Follows directly by Lemma 5.2. □

Corollary 7. *Let $G_1 = (A_1, L_1)$ and $G_2 = (A_2, L_2)$ be directed graphs such that (G_1, G_2, L_3) is a directional splitting of the* ADF *D and $D_1 = D|_{A_1}$. It holds that*

1. $v_1 \in stb(D_1) \wedge v_2 \in stb(D^{v_1}) \Rightarrow (v_1 \cup v_2)|_A \in stb(D)$,
2. $v \in stb(D) \Rightarrow v|_{A_1} \in stb(D_1) \wedge v|_{A_2} \in stb(D^{v|_{A_1}})$.

Proof. Follows directly by Theorem 4 (two-valued model) and Theorem 6 (grounded interpretation). □

Example 3. The ADF D on the left of Figure 5 has $com(D) = \{\{a \mapsto \mathbf{t}, b \mapsto \mathbf{t}, c \mapsto \mathbf{f}, d \mapsto \mathbf{f}\}, \{a \mapsto \mathbf{t}, b \mapsto \mathbf{f}, c \mapsto \mathbf{f}, d \mapsto \mathbf{t}\}, \{a \mapsto \mathbf{t}, b \mapsto \mathbf{u}, c \mapsto \mathbf{f}, d \mapsto \mathbf{u}\}\}$. Now consider the splitting $((\{a,b,c\}, \{(a,c),(b,c),(b,b)\}), (\{d\}, \emptyset), \{(b,d),(c,d)\})$ of D. The interpretation $v_1 = \{a \mapsto \mathbf{t}, b \mapsto \mathbf{u}, c \mapsto \mathbf{f}\}$ is complete for D_1. The v_1-reduct D^{v_1} is depicted on the very right. The only complete interpretation of D^{v_1} is $v_2 = \{d \mapsto \mathbf{u}, x_b \mapsto \mathbf{u}\}$ and indeed $(v_1 \cup v_2)|_{A_D} \in com(D)$. It is easy to verify that this also holds for the reduct based on any other complete interpretation of D_1.

4. General Splitting

So far we have only dealt with splittings of ADFs along the lines of strongly connected components. However, the graph induced by an ADF may not be sparse enough to be suitable for directional splitting. In the following we give some preliminary results on *general splitting* of ADFs. This is identified just by a subset of the arguments of an ADF, which shall represent the first part of the split ADF.

Definition 10. Given an ADF D we call a set $S \subseteq A_D$ a general splitting of D.

First we consider two-valued models. Here we can clear the way for directional splitting by transforming a given ADF while preserving equality with respect to two-valued models.

Definition 11. Given an ADF D, let $L \subseteq L_D$ be a set of links in D. We define $L^- = \{b \mid (b,a) \in L\}$. Moreover, the *L-elimination* of D is defined as

$$D^L = \{\langle a, \varphi_a[b/x_b : b \in L^-]\rangle \mid a \in A_D\} \cup$$
$$\{\langle x_b, x_b\rangle \mid b \in L^-\} \cup \{\langle \omega(D^L), \neg\left(\bigwedge_{b \in L^-} b \leftrightarrow x_b\right) \wedge \neg\omega(D^L)\rangle\},$$

where $\omega(D^L)$ and x_b for each $b \in L^-$ are newly introduced arguments.

Lemma 8. *Let D be an* ADF *and $L \subseteq L_D$. For any two-valued model $v \in val_2(D^L)$ it holds for any $b \in L^-$ that $v(b) = v(x_b)$.*

Proof. Let $v \in val_2(D^L)$ and assume towards a contradiction that $v(b) \neq v(x_b)$. If $v(\omega(D^L)) = \mathbf{t}$ then $v(\varphi_{\omega(D^L)}) = \mathbf{f}$ and if $v(\omega(D^L)) = \mathbf{f}$ then $v(\varphi_{\omega(D^L)}) = \mathbf{t}$, hence v cannot be a two-valued model of D^L. □

The L-elimination indeed preserves equality of two-valued models under projection:

Proposition 9. *Given an* ADF *D, for any $L \subseteq L_D$ it holds that*

1. $v \in val_2(D^L) \Rightarrow v|_{A_D} \in val_2(D)$,
2. $v \in val_2(D) \Rightarrow \exists v' \in val_2(D^L) : v = v'|_{A_D}$.

Proof. 1) Let $v \in val_2(D^L)$ and $a \in A_D$. We need to show, knowing that $v(a) = v(\varphi_a[b/x_b : b \in L^-])$, that $v|_{A_D}(a) = v|_{A_D}(\varphi_a)$. By Lemma 8 it must hold for each $b \in L^-$ that $v(b) = v(x_b)$, hence $v(a) = v(\varphi_a)$ and since $a_D^- \subseteq A_D$, $v|_{A_D} = v|_{A_D}(\varphi_a)$.
2) Consider some $v \in val_2(D)$ and let $v' = v \cup \{x_b \mapsto v(b) \mid b \in L^-\} \cup \{\omega(D^L) \mapsto \mathbf{f}\}$. Obviously $v'(x_b) = v'(\varphi_{x_b})$. Now consider some $a \in A_D$. It holds that $v(a) = v'(a) = v(\varphi_a)$. In order to ensure $v'(\varphi_{\omega(D^L)}) = v'(\omega(D^L)) = \mathbf{f}$, it must hold that $v'(b) = v'(x_b)$ for all $b \in L^-$. Hence $v'(a) = v'(\varphi_a[b/x_b : b \in L^-])$, showing that $v' \in val_2(D^L)$. □

This allows us to apply directional splitting under two-valued models along any desired partition of arguments after a suitable transformation. First the transformation eliminating all links of one direction preserves equality of two-valued models (cf. Proposition 9), and second the computation of two-valued models can be carried out in two stages by directionally splitting the ADF (cf. Theorem 4).

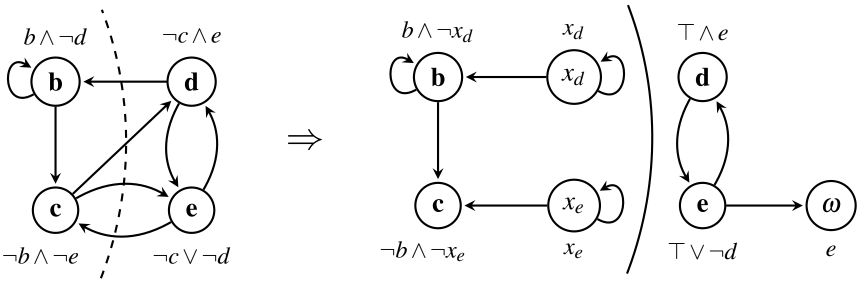

Figure 6. General splitting $S = \{b,c\}$ of the ADF D on the left. The right side depicts the primary slice D^S of the splitting as well as the extended v_1-reduct wrt. S where $v_1 = \{b \mapsto \mathbf{u}, c \mapsto \mathbf{f}, x_d \mapsto \mathbf{u}, x_e \mapsto \mathbf{t}\}$.

Corollary 10. *Given an* ADF *D, let $S \subseteq A_D$ be a general splitting of D. Further let $L = \{(a,b) \in L_D \mid a \in (A_D \setminus S), b \in S\}$ and $X = \{x_b \mid b \in L^-\}$. It holds that*

1. $v_1 \in val_2(D^L|_{(S \cup X)}) \wedge v_2 \in val_2((D^L)^{v_1}) \Rightarrow (v_1 \cup v_2)|_{A_D} \in val_2(D)$,
2. $v \in val_2(D) \Rightarrow \exists v_1, v_2 : (v_1 \cup v_2)|_{A_D} = v \wedge v_1 \in val_2(D^L|_{(S \cup X)}) \wedge v_2 \in val_2((D^L)^{v_1})$.

We now turn to the admissible semantics. A transformation in the fashion of Definition 11 is not possible since we cannot force an interpretation to have equal truth values for two arguments in the three-valued setting. Therefore we have to apply local transformations on each of the sub-frameworks obtained by the splitting.

Definition 12. Given an ADF D, let $S \subseteq A_D$ be a general splitting of D. Further let $B = \{b \in (A_D \setminus S) \mid \exists a \in S : (b,a) \in L_D\}$. The *primary slice* of D wrt. S is defined as

$$D^S = \{\langle a, \varphi_a[b/x_b : b \in B]\rangle \mid a \in S\} \cup \{\langle x_b, x_b\rangle \mid b \in B\}.$$

Moreover, if v is an interpretation of D^S, the *extended v-reduct* of D wrt. S is defined as

$$D^{S,v} = D^v \cup \{\langle \omega(D^{S,v}), \bigwedge_{b \in B, v(x_b)=\mathbf{t}} b \quad \wedge \bigwedge_{b \in B, v(x_b)=\mathbf{f}} (\neg b)\rangle\}.$$

where the new $\omega(D^{S,v})$ is called insurance-argument of $D^{S,v}$.

When S is a general splitting of some ADF D, all arguments not in S which have links to S are simulated in the primary slice of D by new, self-supporting arguments. This has the effect that these arguments can have an arbitrary truth value in an admissible interpretation. Another additional argument in the extended reduct of D ensures that only "valid" interpretations survive the splitting. This construction can be regarded as a kind of guess-and-check-procedure. Figure 6 depicts a general splitting of an exemplary ADF.

Theorem 11. *Given an* ADF *D and a general splitting $S \subseteq A_D$ thereof, let $B = \{b \in (A_D \setminus S) \mid \exists a \in S : (b,a) \in L_D\}$. It holds that*

1. $v_1 \in adm(D^S) \wedge v_2 \in adm(D^{S,v_1}) \wedge v_2(\omega(D^{S,v_1})) = \mathbf{t} \Rightarrow (v_1 \cup v_2)|_{A_D} \in adm(D)$
2. $v \in adm(D) \Rightarrow \exists v_1, v_2 : (v_1 \cup v_2)|_{A_D} = v \wedge v_1 \in adm(D^S) \wedge v_2 \in adm(D^{S,v_1}) \wedge v_2(\omega(D^{S,v_1})) = \mathbf{t}$

Proof. 1) Let $v_1 \in adm(D^S)$ and $v_2 \in adm(D^{S,v_1})$ such that $v_2(\omega(D^{S,v_1})) = \mathbf{t}$. First observe that for any $b \in B$ it holds that if $v_1(x_b) \neq \mathbf{u}$ then $v_1(x_b) = v_2(b)$. Now let $a \in S$. We know that $v_1(a) \leq_i \bigsqcap_{w \in [v_1]_2} w(\varphi_a[b/x_b : b \in B])$. Since by the previous observation $v_1(x_b) \leq_i v_2(b)$ for all $b \in B$, we infer by Lemma 1 and by $a_D^- \subseteq A_D$ that also $(v_1 \cup v_2)|_{A_D}(a) \leq_i \bigsqcap_{w \in [(v_1 \cup v_2)|_{A_D}]_2} w(\varphi_a)$. For all $a \in (A_D \setminus S)$, $(v_1 \cup v_2)|_{A_D}(a) \leq_i \Gamma_D((v_1 \cup v_2)|_{A_D})(a)$ follows by the same reasoning as in the proof of Theorem 6.
2) Consider some $v \in adm(D)$. Let $v_1 = v|_S \cup \{x_b \mapsto v(b) \mid b \in B\}$ and $v_2 = v|_{(A_D \setminus S)} \cup \{x_c \mapsto \mathbf{u} \mid c \in (v|_S)^{\mathbf{u}}\} \cup \{\omega(D^{S,v_1}) \mapsto \mathbf{t}\}$. We argue that $v_1 \in adm(D^S)$ and $v_2 \in adm(D^{S,v_1})$. Since $\varphi_{x_b} = x_b$ it surely holds that $v_1(x_b) \leq_i \Gamma_{D^S}(v_1)(x_b)$. For any $a \in S$, $v_1(a) \leq_i \Gamma_{D^S}(v_1)(a)$ follows from the fact that $v_1(x_b) = v(b)$ by definition. Finally Theorem 6 implies that $v_2(a) \leq_i \Gamma_{D^{S,v_1}}(v_2)(a)$ for all $a \in (A_D \setminus S)$, hence the result follows. □

Example 4. Let D be the ADF on the left side of Figure 6 and consider its general splitting $S = \{b, c\}$. We illustrate the computation of the admissible interpretations $v = \{b \mapsto \mathbf{u}, c \mapsto \mathbf{f}, d \mapsto \mathbf{u}, e \mapsto \mathbf{t}\}$ and $v' = \{b \mapsto \mathbf{u}, c \mapsto \mathbf{f}, d \mapsto \mathbf{t}, e \mapsto \mathbf{t}\}$ of D via the splitting S. First of all we consider the primary slice D^S and observe that $v_1 = \{b \mapsto \mathbf{u}, c \mapsto \mathbf{f}, x_d \mapsto \mathbf{u}, x_e \mapsto \mathbf{t}\}$ is an admissible interpretation thereof. Now the extended v_1-reduct D^{S,v_1} is depicted at the very right part of Figure 6. We observe that the admissible interpretations of D^{S,v_1} having $\omega(D^{S,v_1}) \mapsto \mathbf{t}$ are $v_2 = \{d \mapsto \mathbf{u}, e \mapsto \mathbf{t}, \omega(D^{S,v_1}) \mapsto \mathbf{t}\}$ and $v_2' = \{d \mapsto \mathbf{t}, e \mapsto \mathbf{t}, \omega(D^{S,v_1}) \mapsto \mathbf{t}\}$. Now it indeed holds that $v = (v_1 \cup v_2)|_{A_D}$ and $v' = (v_1 \cup v_2')|_{A_D}$.

On the other hand consider the admissible interpretation $w_1 = \{b \mapsto \mathbf{t}, c \mapsto \mathbf{f}, x_d \mapsto \mathbf{f}, x_e \mapsto \mathbf{u}\}$ of D^S. We get $D^{S,w_1} = \{\langle d, \top \wedge e\rangle, \langle e, \top \vee \neg d\rangle, \langle \omega(D^{S,w_1}), \neg d\rangle\}$ and observe that there is no $w_2 \in adm(D^{S,w_1})$ with $w_2(\omega(D^{S,w_1})) = \mathbf{t}$. This is as expected since there is no $w \in adm(D)$ such that $w(b) = \mathbf{t}$ and $w(c) = \mathbf{f}$.

Note that in general this does not work for the other semantics under consideration. Nevertheless a procedure for gradually computing the preferred interpretations is derivable from Theorem 11. This can be achieved by using the splitting procedure to determine the admissible interpretations and finally selecting the $\leq_i$-maximal elements.

5. Discussion

In this paper, we provided splitting results for ADFs under the semantics proposed in [8], namely two-valued, stable, and three-valued models and admissible, preferred, complete, and grounded interpretations. To summarize, the splitting results show that, given a semantics σ, the results $\sigma(D)$ of the split framework D can be determined in the following way based on the splitting into A_1 and A_2: (1) build the ADF D_1 based on A_1; (2) compute $\sigma(D_1)$; (3) transform D according to $\sigma(D_1)$ to get D_2; (4) compute $\sigma(D_2)$ to get (5) $\sigma(D) = \sigma(D_1) \cup \sigma(D_2)$. Since all operations involved in the splitting can be done efficiently, this lays the basis for an optimization of the evaluation process of ADFs.

Future work includes investigating (general) splitting under the remaining semantics of [8] as well as under the new semantics based on approximation operators [13] and extensions [12]. Moreover, an empirical evaluation is planned in order to verify to which extent an optimization can be achieved by splitting techniques. While directional splitting is very likely to bring considerable savings, the benefits of general splitting are uncertain at this point. Finally, current [20] and new implementations could make use of the results in order to limit search space.

Acknowledgements. The author wishes to thank Stefan Woltran and Wolfgang Dvořák for informative discussions and the anonymous reviewers for helpful comments. This work has been supported by the Austrian Science Fund (FWF) through project I1102.

References

[1] Rahwan, I., Simari, G.R., eds.: Argumentation in Artificial Intelligence. Springer (2009)

[2] Bench-Capon, T.J.M., Dunne, P.E.: Argumentation in artificial intelligence. Artif. Intell. **171** (2007) 619–641

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

[4] Baroni, P., Cerutti, F., Giacomin, M., Guida, G.: AFRA: Argumentation framework with recursive attacks. Int. J. Approx. Reasoning **52** (2011) 19–37

[5] Cayrol, C., Lagasquie-Schiex, M.C.: Bipolarity in argumentation graphs: Towards a better understanding. Int. J. Approx. Reasoning **54** (2013) 876–899

[6] Nielsen, S.H., Parsons, S.: A generalization of Dung's abstract framework for argumentation: Arguing with sets of attacking arguments. In: Argumentation in Multi-Agent Systems. Volume 4766 of Lecture Notes in Computer Science., Springer (2006) 54–73

[7] Brewka, G., Woltran, S.: Abstract dialectical frameworks. In: Proc. KR, AAAI Press (2010) 102–111

[8] Brewka, G., Ellmauthaler, S., Strass, H., Wallner, J.P., Woltran, S.: Abstract dialectical frameworks revisited. In: Proc. IJCAI, IJCAI/AAAI (2013) 803–809

[9] Amgoud, L., Cayrol, C.: On the acceptability of arguments in preference-based argumentation. In: Proc. UAI, Morgan Kaufmann (1998) 1–7

[10] Modgil, S.: Reasoning about preferences in argumentation frameworks. Artif. Intell. **173** (2009) 901–934

[11] Strass, H.: Approximating operators and semantics for abstract dialectical frameworks. Artif. Intell. **205** (2013) 39–70

[12] Polberg, S., Wallner, J.P., Woltran, S.: Admissibility in the abstract dialectical framework. In: Proc. CLIMA XIV. Volume 8143 of Lecture Notes in Computer Science., Springer (2013) 102–118

[13] Strass, H., Wallner, J.P.: Analyzing the computational complexity of abstract dialectical frameworks via approximation fixpoint theory. In: Proceedings KR, To appear (2014)

[14] Baumann, R.: Splitting an argumentation framework. In: Proc. LPNMR 2011. Volume 6645 of Lecture Notes in Computer Science., Springer (2011) 40–53

[15] Baumann, R., Brewka, G., Dvořák, W., Woltran, S.: Parameterized splitting: A simple modification-based approach. In: Correct Reasoning - Essays on Logic-Based AI in Honour of Vladimir Lifschitz. Volume 7265 of Lecture Notes in Computer Science., Springer (2012) 57–71

[16] Lifschitz, V., Turner, H.: Splitting a logic program. In: Proc. ICLP, MIT Press (1994) 23–37

[17] Baumann, R., Brewka, G., Wong, R.: Splitting argumentation frameworks: An empirical evaluation. In: Proc. TAFA. Volume 7132 of Lecture Notes in Computer Science., Springer (2011) 17–31

[18] Baroni, P., Boella, G., Cerutti, F., Giacomin, M., van der Torre, L.W.N., Villata, S.: On input/output argumentation frameworks. In: Proc. COMMA. Volume 245 of Frontiers in Artificial Intelligence and Applications., IOS Press (2012) 358–365

[19] Baroni, P., Giacomin, M., Liao, B.: On topology-related properties of abstract argumentation semantics. a correction and extension to dynamics of argumentation systems: A division-based method. Artif. Intell. (2014)

[20] Ellmauthaler, S., Strass, H.: The DIAMOND system for argumentation: Preliminary report. In: Proc. ASPOCP. (2013) 97–107

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

[22] Baroni, P., Giacomin, M.: On principle-based evaluation of extension-based argumentation semantics. Artif. Intell. **171** (2007) 675–700

[23] Kleene, S.C.: Introduction to Metamathematics. D. Van Nostrand Company (1952)