Spectra in Abstract Argumentation: An Analysis of Minimal Change

Ringo Baumann and Gerhard Brewka

University of Leipzig, Informatics Institute, Germany
{baumann,brewka}@informatik.uni-leipzig.de

Abstract. In this paper we present various new results related to the dynamics of abstract argumentation. Baumann [1] studied the effort needed to enforce a set of arguments E, measured in terms of the minimal number of modifications needed to turn an argumentation framework (AF) $\mathcal{A}$ into a framework $\mathcal{A}^*$ such that $\mathcal{A}^*$ has an extension containing E. This value, called the characteristic, depends on the chosen semantics and the type of admitted modifications. Here we study the inverse problem (called the *spectrum problem*): given a collection of semantics and a modification type, what are the corresponding tuples of characteristics one may obtain for an arbitrary argumentation framework $\mathcal{A}$ and set of arguments E? The set of all these tuples is called the spectrum. We define various properties of spectra and show that the investigation of spectra reveals interesting and surprising insights into the relationship among several semantics.

1 Introduction

Argumentation is the interdisciplinary study of how conclusions can be reached through the construction and evaluation of arguments, that is, structures describing a proposition together with the reasons for accepting it. The field has received growing interest within Artificial Intelligence over the last decades. It covers aspects of knowledge representation and multi-agent systems, but also touches on various philosophical questions (for a very good overview see [2]). Dung's abstract argumentation frameworks (AFs) [3] play a dominant role in the field. In AFs arguments and attacks among them are treated as abstract entities. The focus is on conflict resolution and argument acceptability. Various semantics for AFs have been defined, each of them specifying acceptable sets of arguments, so-called *extensions*, in a particular way.

More recently several problems regarding *dynamic* aspects of abstract argumentation have been addressed in the literature [4–8, 1]. One problem which is relevant to the work presented here concerns the acceptability of certain arguments and is called *enforcing problem* [6]. This is, in brief, the question whether it is possible, given a specific set of allowed operations, to modify a given AF such that a desired set of arguments E is contained in an extension of the modified AF. Several necessary and sufficient conditions under which enforcements are possible were identified.

In addition to clarifying the *possibility* of enforcing certain arguments, a natural further question in this context is concerned with the *effort needed* for the enforcements. This more general problem of *minimal change* [1] can be formulated as follows: what is the minimal number of modifications (additions or removals of attacks) needed to reach an enforcement of E? This value, called *characteristic* in [1], depends on the underlying semantics σ and type of allowed modifications Φ. Quite surprisingly, it was shown that, in case of certain semantics and modification types, there are local criteria to determine the minimal number, although infinitely many possibilities to modify a given AF exist.

In this paper we study a further, closely related question in this context which has some similarity with the famous *Spektralproblem*[1] in model theory [9]. Given a certain semantics σ and a modification type Φ, we study whether there is, for a given natural number n, an AF $\mathcal{A}$ and a set of arguments E such that n is the (σ, Φ)-characteristic of E w.r.t. $\mathcal{A}$. In other words, we want to determine the set of all natural numbers which may occur as (σ, Φ)-characteristics, the so-called (σ, Φ)-*spectrum*. This yields interesting insights into particular semantics. To mention one result, we will show that in case of semi-stable semantics and the addition of weak arguments (arguments which do not attack previous arguments) not each natural number may arise as the minimal effort needed to enforce a certain set D. In particular, the characteristic cannot be 1.

What makes our study even more interesting, as we believe, is the fact that it provides useful and at times surprising new insights into the interrelationships among the studied semantics. To this end, we perform our analysis in parallel for a whole group of semantics which we consider as some of the most important semantics for Dung frameworks. Rather than sets of values, spectra thus become sets of tuples of values. Appropriate properties of the spectra - which we will define in Sect. 3 - will help us to identify such relationships.

The rest of the paper is organized as follows. Sect. 2 reviews the necessary background. Sect. 3 introduces the notion of a spectrum and presents our results for the stable/semi-stable/preferred spectra under various types of modifications. In Sect. 4 we discuss related work and conclude.

2 Background

An *argumentation framework* $\mathcal{F}$ 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 the *attack relation*. The set of all AFs is denoted by $\mathscr{A}$. If $(a, b) \in R$ holds we say that a *attacks* b, or b is *defeated* by a in $\mathcal{F}$. An argument $a \in A$ is *defended* by a set $A' \subseteq A$ in $\mathcal{F}$ if for each $b \in A$ with $(b, a) \in R$, b is defeated by some $a' \in A'$ in $\mathcal{F}$. Furthermore, we say that a set $A' \subseteq A$ is *conflict-free* in $\mathcal{F}$ if there are no arguments $a, b \in A'$ such that a attacks b. The set of all conflict-free sets of an AF $\mathcal{F}$ is denoted by $cf(\mathcal{F})$. For an AF $\mathcal{F} = (B, S)$ we use $A(\mathcal{F})$ to refer to B and $R(\mathcal{F})$ to refer to S. Finally, we introduce the union of two AFs as usual, namely $\mathcal{F} \cup \mathcal{G} = (A(\mathcal{F}) \cup A(\mathcal{G}), R(\mathcal{F}) \cup R(\mathcal{G}))$.

Semantics determine acceptable sets of arguments for a given AF $\mathcal{F}$, so-called *extensions*. The set of all extensions of $\mathcal{F}$ under semantics σ is denoted by $\mathcal{E}_\sigma(\mathcal{F})$. For two semantics σ, τ we use $\sigma \subseteq \tau$ to indicate that for any $\mathcal{F} \in \mathscr{A}$, $\mathcal{E}_\sigma(\mathcal{F}) \subseteq \mathcal{E}_\tau(\mathcal{F})$. Due to the limited space we consider stable (st), preferred (pr) and semi-stable (ss) semantics only [3, 10].

[1] Roughly speaking, Scholz investigated the possible sizes finite models of a first-order sentence may have.

Definition 1 (Semantics). *Given an AF $\mathcal{F} = (A, R)$ and $E \subseteq A$. E is a*

1. *stable extension ($E \in \mathcal{E}_{st}(\mathcal{F})$) iff*
   $E \in cf(\mathcal{F})$ and each $a \in A \backslash E$ is defeated by some $e \in E$,
2. *admissible set ($E \in \mathcal{E}_{ad}(\mathcal{F})$) iff*
   $E \in cf(\mathcal{F})$ and each $e \in E$ is defended by E in $\mathcal{F}$,
3. *preferred extension ($E \in \mathcal{E}_{pr}(\mathcal{F})$) iff*
   $E \in \mathcal{E}_{ad}(\mathcal{F})$ and for each $E' \in \mathcal{E}_{ad}(\mathcal{F})$, $E \notin E'$ and
4. *semi-stable extension ($E \in \mathcal{E}_{ss}(\mathcal{F})$) iff*
   $E \in \mathcal{E}_{ad}(\mathcal{F})$ and for each $E' \in \mathcal{E}_{ad}(\mathcal{F})$, $R^+_{\mathcal{F}}(E) \notin R^+_{\mathcal{F}}(E')$ where $R^+_{\mathcal{F}}(E) = E \cup \{b \mid (a, b) \in R, a \in E\}$.

It is well known that $st \subseteq ss \subseteq pr$. Furthermore, there exist sufficient conditions for the agreement of the considered semantics. In particular, $st = ss$ if $st \neq \emptyset$ [10] and $st = pr$ if the considered AFs are SCC-symmetric and self-loop-free (compare [11]).

Expansions were introduced by [6]. They will be our object of investigation since they represent reasonable types of dynamic argumentation scenarios.

Definition 2 (Expansions). *An AF $\mathcal{F}^*$ is an expansion of AF $\mathcal{F} = (A, R)$ (for short, $\mathcal{F} \preceq_E \mathcal{F}^*$) iff $\mathcal{F}^* = (A \cup A^*, R \cup R^*)$ where $A^* \cap A = R^* \cap R = \emptyset$. An expansion is called*

1. *normal ($\mathcal{F} \preceq_N \mathcal{F}^*$) iff $\forall ab\ ((a, b) \in R^* \rightarrow a \in A^* \vee b \in A^*)$,*
2. *strong ($\mathcal{F} \preceq_S \mathcal{F}^*$) iff $\mathcal{A} \preceq_N \mathcal{A}^*$ and $\forall ab\ ((a, b) \in R^* \rightarrow \neg(a \in A \wedge b \in A^*))$,*
3. *weak ($\mathcal{F} \preceq_W \mathcal{F}^*$) iff $\mathcal{A} \preceq_N \mathcal{A}^*$ and $\forall ab\ ((a, b) \in R^* \rightarrow \neg(a \in A^* \wedge b \in A))$.*

For short, normal expansions add new arguments and possibly new attacks which involve at least one of the fresh arguments. Strong (weak) expansions are normal and only add *strong (weak) arguments*, i.e. the added arguments never are attacked by (attack) former arguments. For the purpose of illustration we present the following simple example.

Example 1. The AF $\mathcal{F} = (\{a, b\}, \{(a, b)\})$ is the initial framework. Arbitrary, normal, weak and strong expansions of $\mathcal{F}$ are given by $\mathcal{F}_E$, $\mathcal{F}_N$, $\mathcal{F}_W$ or $\mathcal{F}_S$, respectively.

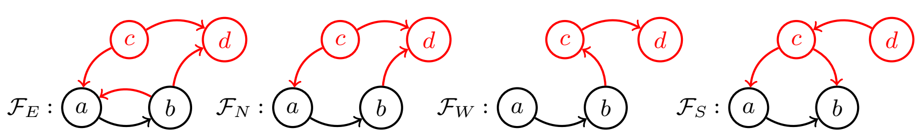

Fig. 1. Notions of Expansions

As usual $\mathcal{F} \prec_X \mathcal{F}^*$ for $X \in \{E, N, S, W\}$ stands for $\mathcal{F} \preceq_X \mathcal{F}^*$ and $\mathcal{F} \neq \mathcal{F}^*$. To simplify notation we will later on often use X to refer to $\preceq_X$. Whenever infix notation is used we stick to $\preceq_X$, though.

The *minimal change problem* [1] is the problem of determining the minimal effort needed to transform a given argumentation framework, using a particular type of modifications, into a framework that possesses an extension containing a specific set of arguments E. The effort is characterized by the (σ, Φ)-*characteristic*:

Definition 3 (Characteristic). *Given a semantics σ, a binary relation $\Phi \subseteq \mathscr{A} \times \mathscr{A}$ and an AF $\mathcal{F}$. The (σ, Φ)-characteristic of a set $C \subseteq A(\mathcal{F})$ is a natural number or infinity defined by the following function*

$$N^{\mathcal{F}}_{\sigma,\Phi} : \wp(A(\mathcal{F})) \to \mathbb{N}_\infty$$

$$C \mapsto \begin{cases} 0, & \exists C' : C \subseteq C' \text{ and } C' \in \mathcal{E}_\sigma(\mathcal{F}) \\ k, & k = \min\{d(\mathcal{F}, \mathcal{G}) \mid (\mathcal{F}, \mathcal{G}) \in \Phi, N^{\mathcal{G}}_{\sigma,\Phi}(C) = 0\} \\ \infty, & \text{otherwise.} \end{cases}$$

Here we define $d(\mathcal{F}, \mathcal{G})$ as the number of added or removed attacks needed to transform $\mathcal{F}$ to $\mathcal{G}$. This means $d(\mathcal{F}, \mathcal{G}) = |R(\mathcal{F}) \, \Delta \, R(\mathcal{G})|$ where Δ is the symmetric difference.

3 The Spectrum Problem

Given a semantics and a type of allowed modifications, the characteristic provides information about the effort needed to enforce a set of arguments C starting from an AF $\mathcal{F}$. Here we study the inverse problem, that is, given a particular characteristic, is there an AF $\mathcal{F}$ and a set of arguments C which possess this characteristic. More generally, we will consider n-tuples of semantics and modification types and ask whether some $\mathcal{F}$ and C possess a given n-tuple of characteristics simultaneously. A tuple of characteristics satisfying this condition is called a *fibre*. A fibre is said to be *finite* if all entries are natural numbers. The set of all fibres provides important insights on how close or far apart the characteristics of a set C may be. That's why this set is called the *spectrum*. Here is the formal definition.

Definition 4. *Given n semantics σ_1,...,σ_n and n binary relations $\Phi_1, ..., \Phi_n \subseteq \mathscr{A} \times \mathscr{A}$. The $(\sigma_1, \Phi_1, ..., \sigma_n, \Phi_n)$-spectrum is a set of n-tuples (so-called fibres) defined as follows:*

$$\mathcal{S}_{(\sigma_i,\Phi_i)_{i=1}^n} = \{(k_1, ..., k_n) \mid \exists \mathcal{F} \in \mathscr{A} \; \exists C \subseteq A(\mathcal{F}) : N^{\mathcal{F}}_{\sigma_i,\Phi_i}(C) = k_i \text{ for all } i \in \{1, ..., n\}\}.$$

For convenience, if $\Phi_1 = ... = \Phi_n$ we simply write $(\sigma_1, ..., \sigma_n, \Phi)$-spectrum or $\mathcal{S}_{(\sigma_1,...,\sigma_n,\Phi)}$. These are exactly the types of spectra which we will consider in this paper. Furthermore, we will restrict ourselves to stable (st), semi-stable (ss) and preferred (pr) semantics, arguably the most important semantics for Dung frameworks. The relations we study will be normal, strong, weak and arbitrary expansions.

We first introduce some basic properties spectra may possess.

Definition 5. *A spectrum $\mathcal{S}_{(\sigma_i,\Phi_i)_{i=1}^n}$ is*

1. *m.d.s. iff any finite fibre $(k_1, ..., k_n) \in \mathcal{S}_{(\sigma_i,\Phi_i)_{i=1}^n}$ is a monotonic decreasing sequence,*

2. *m.d.s.-complete iff* $\mathcal{S}_{(\sigma_i,\Phi_i)_{i=1}^n}$ *is m.d.s. and* $\{(k_1,...,k_n) \in \mathbb{N}^n \mid k_1 \geq ... \geq k_n\} \subseteq \mathcal{S}_{(\sigma_i,\Phi_i)_{i=1}^n}$,
3. *coherent iff there is no fibre* $(k_1,...,k_n) \in \mathcal{S}_{(\sigma_i,\Phi_i)_{i=1}^n}$, *s.t.* $k_i = \infty$ *and* $k_j \neq \infty$ *for some indices* $1 \leq i,j \leq n$ *and*
4. *positive iff any fibre* $(k_1,...,k_n) \in \mathcal{S}_{(\sigma_i,\Phi_i)_{i=1}^n}$ *is finite.*

These properties are interesting for the following reasons: if a spectrum for semantics $\sigma_1,\ldots,\sigma_n$ is m.d.s., then we know that whenever enforcing is possible for all of them it is at least as difficult using σ_i as it is using σ_j given that $i < j$. If it is m.d.s.-complete we know in addition that it can in fact be arbitrarily more difficult. Coherence means that whether some C is enforceable or not does not depend on the choice of the considered semantics. Positive means each set C can actually be enforced.

A few relationships among these properties are clear by definition. First, an m.d.s.-complete spectrum is m.d.s. and second, a positive spectrum is coherent. Further interpretations of the introduced properties are given in the following subsections.

3.1 The (st, ss, pr, Φ)-Spectrum ($\Phi \in \{E, N, S\}$)

In this subsection we will characterize the (st, ss, pr)-spectra w.r.t. strong, normal and arbitrary expansions. In [1, Corrolary 3] it was shown that the stable (semi-stable) characteristic exceeds the semi-stable (preferred) characteristic w.r.t. any binary relation over the set of all finite AFs. Consequently, the considered spectra are m.d.s.

Quite surprisingly, the following proposition shows that the mentioned spectra are even m.d.s.-complete, i.e. the stable (semi-stable) characteristic may take values which exceed the semi-stable (preferred) characteristic by **any** natural number. In a sense this result is negative as it tells us that information about the characteristic of one semantics does not help in determining the characteristic of the other semantics: even if we know that the characteristic w.r.t. preferred semantics for a certain set D is, say, 1 (i.e., only 1 additional attack is needed), there is no possibility to give an upper bound of the characteristic w.r.t. semi-stable or stable semantics. The result underlines the independence of the considered semantics w.r.t. the minimal change problem. It indicates that the choice of the considered semantics may influence the characteristic dramatically, even though the considered semantics possess many similarities.

Proposition 1. *For any* $\Phi \in \{E, N, S\}$, $\mathcal{S}_{(st,ss,pr,\Phi)}$ *is m.d.s.-complete.*

Proof. Let $\Phi \in \{E, N, S\}$ and $k, l, m \in \mathbb{N}$, s.t. $k \geq l \geq m$. Hence, we may assume that $l = m + n$ and $k = m + n + o$ for some $n, o \in \mathbb{N}$. If we may construct AFs $\mathcal{F}$ and corresponding sets $C \subseteq A(\mathcal{F})$, s.t. $N^{\mathcal{F}}_{st,\Phi}(C) = m + n + o$, $N^{\mathcal{F}}_{ss,\Phi}(C) = m + n$ and $N^{\mathcal{F}}_{pr,\Phi}(C) = m$, then $(k, l, m) \in \mathcal{S}_{(st,ss,pr,\Phi)}$ follows. Thus, $\mathcal{S}_{(st,ss,pr,\Phi)}$ is shown to be m.d.s.-complete. We define the AF $\mathcal{F}_{m,n,o} = (A_{m,n,o}, R_{m,n,o})$ where

$$A_{m,n,o} = \{a\} \cup \{b_j \mid 1 \leq j \leq m\} \cup \{c_j, d_j, e_j \mid 1 \leq j \leq n\} \cup \{f_j \mid 1 \leq j \leq o\} \text{ and}$$

$$\begin{aligned} R_{m,n,o} = {} & \{(b_j, a), (b_j, b_j) \mid 1 \leq j \leq m\} \cup \{(c_j, d_j), (d_j, c_j), (e_j, e_j) \mid 1 \leq j \leq n\} \cup \\ & \{(d_j, e_i) \mid 1 \leq j, i \leq n\} \cup \{(d_j, b_i) \mid 1 \leq j \leq n, 1 \leq i \leq m\} \cup \\ & \{(d_j, d_i) \mid j \neq i, 1 \leq j, i \leq n\} \cup \{(f_j, f_j) \mid 1 \leq j \leq o\} \cup \\ & \{(d_j, f_i) \mid 1 \leq j \leq n, 1 \leq i \leq o\}. \end{aligned}$$

Note that if a subindex equals zero, then there are no corresponding arguments and attacks. For the sake of clarity we present here an instantiation of the presented scheme, namely $\mathcal{F}_{3,2,4}$.

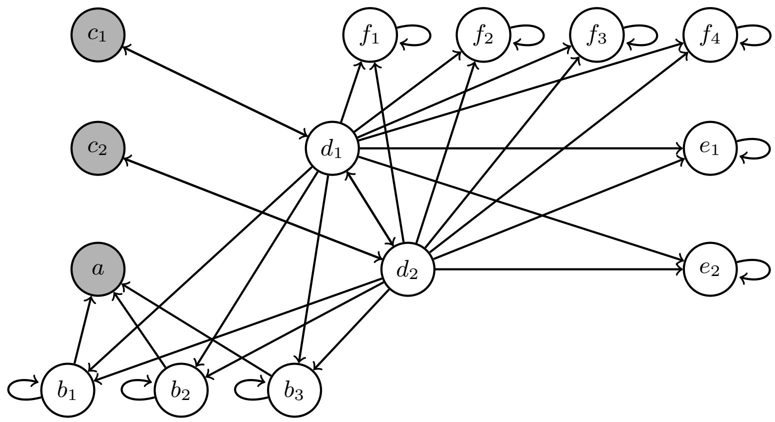

Fig. 2. The AF $\mathcal{F}_{3,2,4}$

The grey highlighted arguments belong to the set $C_2 = \{a, c_1, c_2\}$ which is an instantiation of the scheme $C_n = \{a\} \cup \{c_j \mid 1 \leq j \leq n\}$. We claim that $N^{\mathcal{F}_{m,n,o}}_{st,\Phi}(C_n) = m + n + o$, $N^{\mathcal{F}_{m,n,o}}_{ss,\Phi}(C_n) = m + n$ and $N^{\mathcal{F}_{m,n,o}}_{pr,\Phi}(C_n) = m$. By construction C_n is conflict-free in $\mathcal{F}_{m,n,o}$. Furthermore, C_n does not have proper conflict-free supersets (*). Applying the characterization theorems of [1] (Theorem 9, Def. 8) $N^{\mathcal{F}_{m,n,o}}_{pr,\Phi}(C_n) = V^{\mathcal{F}_{m,n,o}}_{ad,S}(C_n) = \left|R^-_{\mathcal{F}_{m,n,o}}(C_n) \setminus R^+_{\mathcal{F}_{m,n,o}}(C_n)\right| = |\{b_j \mid 1 \leq j \leq m\}| = m$ because these arguments are not counterattacked by C_n. In case of stable semantics $N^{\mathcal{F}_{m,n,o}}_{st,\Phi}(C_n) = V^{\mathcal{F}_{m,n,o}}_{st,S}(C_n) = \left|A(\mathcal{F}_{m,n,o}) \setminus R^+_{\mathcal{F}_{m,n,o}}(C_n)\right| = |\{b_j \mid 1 \leq j \leq m\} \cup \{e_j \mid 1 \leq j \leq n\} \cup \{f_j \mid 1 \leq j \leq o\}| = m + n + o$ since exactly these arguments are not attacked by C_n.

To see that $N^{\mathcal{F}_{m,n,o}}_{ss,\Phi}(C_n) = m + n$ is much more difficult. At first we will show that $N^{\mathcal{F}_{m,n,o}}_{ss,E}(C_n) \geq m+n$ and finally, $N^{\mathcal{F}_{m,n,o}}_{ss,S}(C_n) \leq m+n$. Consequently, $N^{\mathcal{F}_{m,n,o}}_{ss,\Phi}(C_n) = m + n$ for any $\Phi \in \{E, N, S\}$ is proven (Corollary 4 [1]). Consider the n conflict-free sets $S^1_n, ..., S^n_n$ where $S^j_n = \{a\} \cup \{c_i \mid 1 \leq i \leq n\} \setminus \{c_j\} \cup \{d_j\}$. We observe that $C_n \notin S^j_n$ for $n \geq 1$ and furthermore, $R^+_{\mathcal{F}_{m,n,o}}(C_n) \subset R^+_{\mathcal{F}_{m,n,o}}(S^j_n) = A_{m,n,o}$. Assume now $N^{\mathcal{F}_{m,n,o}}_{ss,E}(C_n) = l' < m + n$. Hence, there is an AF $\mathcal{G}$, s.t. $d(\mathcal{F}_{m,n,o}, \mathcal{G}) = l'$, $\mathcal{F}_{m,n,o} \leq \mathcal{G}$ and furthermore, there is a conflict-free superset C'_n of C_n with the property $C'_n \in \mathcal{E}_{ss}(\mathcal{G})$. In consideration of (*) we deduce that $C'_n = C_n \cup G$ where G is a set of fresh arguments. Since any semi-stable extension is admissible we conclude that each b_j has to be attacked by C'_n. This means at least m additional attacks of $\mathcal{G}$ are required for this task.

Let us consider now the remaining $l'' < n = \left|\{S_n^j \mid 1 \leq j \leq n\}\right|$ new attacks. The set S_n^j that we look for satisfies the following conditions: 1. for any $g \in G$, $(d_j, g), (g, d_j) \notin R(\mathcal{G})$, 2. $(d_j, d_j) \notin R(\mathcal{G})$, 3. for any $i \neq j$, $(c_i, d_j), (d_j, c_i) \notin R(\mathcal{G})$ as well as $(a, d_j), (d_j, a) \notin R(\mathcal{G})$ and 4. for any $g \in A(\mathcal{G})\backslash\{A(\mathcal{F}_{m,n,o}) \cup G\}$, $(c_j, g), (g, d_j) \notin R(\mathcal{G})$. Since any new attack may eliminate at most one potential candidate we deduce that there is indeed such a S_n^j satisfying 1. - 4. We will show now that $S_n^j \cup G \in \mathcal{E}_{ad}(G)$ and $R^+_{\mathcal{G}}(C_n \cup G) \subset R^+_{\mathcal{G}}(S_n^j \cup G)$ contradicting $C'_n \in \mathcal{E}_{ss}(\mathcal{G})$. Let us consider the range $R^+_{\mathcal{G}}(C_n \cup G)$. Obviously, there is an index i, s.t. $e_i \notin R^+_{\mathcal{G}}(C_n \cup G)$ since $l'' < n$ was assumed. Note that $e_i \in R^+_{\mathcal{G}}(S_n^j \cup G)$ by construction of $\mathcal{F}_{m,n,o}$ and S_n^j. Furthermore, in consideration of the first part of condition 4. (c_j does not "reach" further arguments) we immediately conclude that $R^+_{\mathcal{G}}(C_n \cup G) \subseteq R^+_{\mathcal{G}}(S_n^j \cup G)$. Altogether, $R^+_{\mathcal{G}}(C_n \cup G) \subset R^+_{\mathcal{G}}(S_n^j \cup G)$ has to hold. Furthermore, $S_n^j \cup G$ is conflict-free in $\mathcal{G}$ for two reasons, first S_n^j satisfies conditions 1. - 3. and second, $C_n \cup G$ is assumed to be admissible and in particular, conflict-free in $\mathcal{G}$. Assume now that $S_n^j \cup G \notin \mathcal{E}_{ad}(\mathcal{G})$. This means, there is argument $g \in A(\mathcal{G})$ which attacks $S_n^j \cup G$ without being counterattacked. Since conflict-freeness is already shown and $A_{\mathcal{F}_{m,n,o}} \subseteq R^+_{\mathcal{G}}(S_n^j \cup G)$ obviously holds, we deduce $g \in A(\mathcal{G})\backslash\{A_{m,n,o} \cup G\}$. In consideration of the second part of condition 4. ($(g, d_j) \notin R(\mathcal{G})$) it follows that g attacks some c_i with $i \neq j$ or an argument $g' \in G$. Since $C_n \cup G$ is assumed to be admissible in $\mathcal{G}$ there is an argument $c' \in C_n \cup G$, s.t. $(c', g) \in R(\mathcal{G})$. If $c' \in G$, then obviously $c' \in S_n^j \cup G$. If $c' \in C_n$, then $c' \in S_n^j \cup G$ because the second part of condition 4. ($(c_j, g) \notin R(\mathcal{G})$) guarantees $c_j \neq c'$. This means, under the assumption $N^{\mathcal{F}_{m,n,o}}_{ss,E}(C_n) = l' < m + n$ we derived a contradiction, namely $C'_n \in \mathcal{E}_{ss}(\mathcal{G}) \wedge C'_n \notin \mathcal{E}_{ss}(\mathcal{G})$. Hence, $N^{\mathcal{F}_{m,n,o}}_{ss,E}(C_n) \geq m + n$ is shown.

Let us prove now that $N^{\mathcal{F}_{m,n,o}}_{ss,S}(C_n) \leq m+n$. Consider therefore a fresh argument c and the AF $\mathcal{G}_{m,n} = (A_{m,n,o} \cup \{c\}, R_{m,n,o} \cup \{(c, b_j) \mid 1 \leq j \leq m\} \cup \{(c, d_j) \mid 1 \leq j \leq n\})$. One can easily verify that $C_n \cup \{c\} \in \mathcal{E}_{ss}(\mathcal{G}_{m,n})$ and furthermore, $\mathcal{F} \preceq_S \mathcal{G}_{m,n}$. Since $d(\mathcal{F}_{m,n,o}, \mathcal{G}_{m,n}) = m+n$ we conclude $N^{\mathcal{F}_{m,n,o}}_{ss,S}(C_n) \leq m+n$. Finally, $N^{\mathcal{F}_{m,n,o}}_{ss,\Phi}(C_n) = m+n$ for any $\Phi \in \{E, N, S\}$ is proven.

The following proposition shows that the spectrum $\mathcal{S}_{(st,ss,pr,\Phi)}$ is coherent, i.e. any fibre either possesses finite values or all values equal infinity. This means, under the considered semantics it is impossible that a set C may be enforced w.r.t. a semantics σ and simultaneously, C is not enforceable w.r.t. another semantics τ. Furthermore, we show that the considered spectra are not positive, i.e. there are unenforcable sets.

Proposition 2. *For any $\Phi \in \{E, N, S\}$, $\mathcal{S}_{(st,ss,pr,\Phi)}$ is coherent but not positive.*

Proof. Given $\Phi \in \{E, N, S\}$. First, we will prove the coherence of $\mathcal{S}_{(st,ss,pr,\Phi)}$. Since $\mathcal{S}_{(st,ss,pr,\Phi)}$ is already shown to be m.d.s.-complete it suffices to prove that for any fibre $(k, l, m) \in \mathcal{S}_{(st,ss,pr,\Phi)}$, if $m < \infty$, then $l < \infty$ and if $l < \infty$, then $k < \infty$. Let $m < \infty$. Hence there is an AF $\mathcal{F}$ and a set $C \subseteq A(\mathcal{F})$, s.t. $N^{\mathcal{F}}_{pr,\Phi}(C) = m$. This means, C has to be conflict-free in $\mathcal{F}$. Applying Corollary 7 in [1] we deduce $l = N^{\mathcal{F}}_{ss,S}(C) \leq |A(\mathcal{F})\backslash C| < \infty$. Since $N^{\mathcal{F}}_{ss,S}(C) \geq N^{\mathcal{F}}_{ss,N}(C) \geq N^{\mathcal{F}}_{ss,E}(C)$ (compare Corollary 4 [1]) holds we are done. In the same way one may show that $l < \infty$ implies $k < \infty$.

To prove that $\mathcal{S}_{(st,ss,pr,\Phi)}$ is not positive it suffices to construct a non-finite fibre. Consider therefore $\mathcal{F} = (\{a\}, \{(a,a)\})$ and $C = \{a\}$. Since C does not possess conflict-free supersets we deduce $N^{\mathcal{F}}_{ad,\Phi}(C) = \infty$ (compare Theorem 9, Def. 6 [1]). Furthermore, by Prop. 5 [1] we get $(\infty, \infty, \infty) \in \mathcal{S}_{(st,ss,pr\Phi)}$ concluding the proof.

The following Theorem summarizes the earlier results. Note that the listed properties fully characterize the considered spectra. This means, it is decidable whether an arbitrary fibre belongs to the considered spectra.

Theorem 1. *For any $\Phi \in \{E, N, S\}$, $\mathcal{S}_{(st,ss,pr,\Phi)}$ is coherent, m.d.s.-complete but not positive.*

3.2 Properties of the (st, ss, pr, W)-Spectrum

The following example taken from [1] shows some first and notable differences between the coinciding spectra w.r.t. normal, strong and arbitrary expansions and the spectrum w.r.t. weak expansions considered in this section.

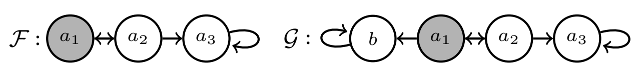

Fig. 3. Non-coherence of the Weak Spectrum

The AFs above exemplify that the (st, ss, pr, W)-spectrum is not coherent since $N^{\mathcal{F}}_{st,W}(\{a_1\}) = \infty$ (unenforceable) and $N^{\mathcal{F}}_{pr,W}(\{a_1\}) = 0$ (already accepted). Furthermore, $1 \leq N^{\mathcal{F}}_{ss,W}(\{a_1\}) \leq 2$ because $\{a_1\}$ and all its proper supersets are not semi-stable in $\mathcal{F}$ but $\{a_1\}$ is semi-stable in $\mathcal{G}$.

Unfortunately, (up to now) there are no characterization theorems for semi-stable semantics. Nevertheless, with the help of the following impossibility result it is shown that $N^{\mathcal{F}}_{ss,W}(\{a_1\}) = 2$ holds. This means, if a desired set of arguments D is not already contained in a semi-stable extension of the initial framework, then the minimal effort needed to enforce D is at least 2 in case of weak expansions.

Proposition 3. $(n, 1, m) \notin \mathcal{S}_{(st,ss,pr,W)}$ *for each* $n, m \in \mathbb{N}_\infty$.

Proof. Since n, m are assumed to be arbitrary natural numbers or ∞ it suffices to prove that $(1) \notin \mathcal{S}_{(ss,W)}$. Assume $(1) \in \mathcal{S}_{(ss,W)}$, i.e. there is an AF $\mathcal{F}$ and a set C with the property $N^{\mathcal{F}}_{ss,W}(C) = 1$. This means there is an AF $\mathcal{G}$, s.t. $\mathcal{F} \preceq_W \mathcal{G}$, $d(\mathcal{F}, \mathcal{G}) = 1$ and a set $C' \supseteq C$ with $C' \in \mathcal{E}_{ss}(\mathcal{G})$. W.l.o.g. $C' = D \cup E$ where $C \subseteq D \subseteq A(\mathcal{F})$ and $E \subseteq A(\mathcal{G}) \backslash A(\mathcal{F})$. Since every semi-stable extension is admissible we deduce $D \in \mathcal{E}_{ad}(\mathcal{F})$. Furthermore, $N^{\mathcal{F}}_{ss,\preceq_W}(C) = 1 \neq 0$ implies there is an admissible set D' in $\mathcal{F}, s.t. R^+_{\mathcal{F}}(D) \subset R^+_{\mathcal{F}}(D')$ (*). We will show now that $D \cup E \notin \mathcal{E}_{ss}(\mathcal{G})$ by proof by cases. Let (d, e) be the new attack. Note that $e \in A(\mathcal{G}) \backslash A(\mathcal{F})$ is implied since $\mathcal{F} \preceq_W \mathcal{G}$ is assumed. Furthermore, $d \in D$ and $e \in E$ is impossible since $D \cup E \in cf(\mathcal{G})$.

1^{st}case: Let $d \in D \backslash D'$ and $e \notin E$. We observe $R^+_{\mathcal{G}}(D \cup E) = R^+_{\mathcal{F}}(D) \cup E \cup \{e\}$. Furthermore, E only contains isolated arguments in $\mathcal{G}$ and hence, $D' \cup E \in \mathcal{E}_{ad}(\mathcal{G})$.

Because of (*) and $d \in D \backslash D'$ we conclude e is defended by D' in $\mathcal{G}$. Thus, $D' \cup E \cup \{e\} \in \mathcal{E}_{ad}(\mathcal{G})$ and obviously, $R^+_{\mathcal{G}}(D' \cup E \cup \{e\}) = R^+_{\mathcal{F}}(D') \cup E \cup \{e\}$. In consideration of (*) it follows that $R^+_{\mathcal{F}}(D) \cup E \cup \{e\} = R^+_{\mathcal{G}}(D \cup E) \subset R^+_{\mathcal{G}}(D' \cup E \cup \{e\})$ and hence, $D \cup E \notin \mathcal{E}_{ss}(\mathcal{G})$ is shown.

2^{nd}case: Let $d \in D \cap D'$ and $e \notin E$. Consequently, $D' \cup E \in \mathcal{E}_{ad}(\mathcal{G})$ and furthermore, $R^+_{\mathcal{G}}(D \cup E) = R^+_{\mathcal{F}}(D) \cup E \cup \{e\} \subset^{(*)} R^+_{\mathcal{F}}(D') \cup E \cup \{e\} = R^+_{\mathcal{G}}(D' \cup E)$ contradicting $D \cup E \in \mathcal{E}_{ss}(\mathcal{G})$.

3^{rd}case: Let $d \in D' \backslash D$ and $e \notin E$. Again, $D' \cup E \in \mathcal{E}_{ad}(\mathcal{G})$ holds and furthermore, $R^+_{\mathcal{G}}(D \cup E) = R^+_{\mathcal{F}}(D) \cup E \subset^{(*)} R^+_{\mathcal{F}}(D') \cup E \cup \{e\} = R^+_{\mathcal{G}}(D' \cup E)$ in contradiction to $D \cup E \in \mathcal{E}_{ss}(\mathcal{G})$.

4^{th}case: Let $d \in D' \backslash D$ and $e \in E$. Hence, $D' \cup (E \backslash \{e\}) \in \mathcal{E}_{ad}(\mathcal{G})$. Furthermore, $R^+_{\mathcal{G}}(D \cup E) = R^+_{\mathcal{F}}(D) \cup E \subset^{(*)} R^+_{\mathcal{F}}(D') \cup (E \backslash \{e\}) \cup \{e\} = R^+_{\mathcal{F}}(D) \cup E = R^+_{\mathcal{G}}(D' \cup E)$. Consequently, $D \cup E \notin \mathcal{E}_{ss}(\mathcal{G})$ is shown.

5^{th}case: Let $d \in A(\mathcal{F}) \backslash (D' \cup D)$ and $e \in E$. Thus, D has to counterattack d in $\mathcal{F}$ since $D \cup E$ is assumed to be admissible in $\mathcal{G}$. Let $d' \in D$ be the counterattacker of d. If $d' \in D'$ we conclude $D' \cup E \in \mathcal{E}_{ad}(\mathcal{G})$. If not, it follows the existence of an argument $d'' \in D'$, s.t. $(d'', d) \in R(\mathcal{F})$ since (*) is assumed. Again, we get $D' \cup E \in \mathcal{E}_{ad}(\mathcal{G})$. In both cases, $R^+_{\mathcal{G}}(D \cup E) = R^+_{\mathcal{F}}(D) \cup E \subset^{(*)} R^+_{\mathcal{F}}(D') \cup E = R^+_{\mathcal{G}}(D' \cup E)$ contradicting $D \cup E \in \mathcal{E}_{ss}(\mathcal{G})$.

6^{th}case: Let $d \in A(\mathcal{F}) \backslash (D' \cup D)$ and $e \notin E$. Consequently, $D' \cup E \in \mathcal{E}_{ad}(\mathcal{G})$ and thus, $R^+_{\mathcal{G}}(D \cup E) = R^+_{\mathcal{F}}(D) \cup E \subset^{(*)} R^+_{\mathcal{F}}(D') \cup E = R^+_{\mathcal{G}}(D' \cup E)$ in contradiction to $D \cup E \in \mathcal{E}_{ss}(\mathcal{G})$.

7^{th}case: Let $d, e \in A(\mathcal{G}) \backslash A(\mathcal{F})$. Since $d(\mathcal{F}, \mathcal{G}) = 1$ it follows that $e \notin E$. Consequently, $D' \cup E \in \mathcal{E}_{ad}(\mathcal{G})$ and furthermore, $R^+_{\mathcal{G}}(D \cup E) = R^+_{\mathcal{F}}(D) \cup R^+_{\mathcal{G}}(E) \subset^{(*)} R^+_{\mathcal{F}}(D') \cup R^+_{\mathcal{G}}(E) = R^+_{\mathcal{G}}(D' \cup E)$ contradicting $D \cup E \in \mathcal{E}_{ss}(\mathcal{G})$.

The proposition above and its usage for the illustrated problem, namely determining the characteristic in a certain argumentation scenario, underline that the investigation of spectra reveals important insights into the minimal change problem. The following impossibility result reveals a further surprising interrelation between the considered semantics, namely that for any $\mathcal{F}$ and any set of arguments E it is impossible that E is already contained in a preferred extension yet unenforceable using semi-stable semantics.

Proposition 4. $(\infty, \infty, 0) \notin \mathcal{S}_{(st,ss,pr,W)}$.

Proof. We will show the stronger result, namely $(\infty, 0) \notin \mathcal{S}_{(ss,pr,W)}$. Assume $(\infty, 0) \in \mathcal{S}_{(ss,pr,W)}$, i.e. there is an AF $\mathcal{F}$ and a set C with the property $N^{\mathcal{F}}_{pr,\leq_W}(C) = 0$ and $N^{\mathcal{F}}_{ss,\leq_W}(C) = \infty$. This means, there exists a set $C' \supseteq C$ with $C' \in \mathcal{E}_{pr}(\mathcal{F})$. Since all considered AFs are assumed to be finite we deduce $C' = \{c'_1, ..., c'_n\}$ for some $n \in \mathbb{N}$. Let $D = \{d_1, ..., d_n\}$ be a set of fresh arguments and consider $\mathcal{G} = (A(\mathcal{F}) \cup D, R(\mathcal{F}) \cup \{(d_i, d_i), (c'_i, d_i) \mid 1 \leq i \leq n\})$. Obviously, $d(F, G) = 2n$ and $F \leq_W G$. Furthermore, the range of C' in $\mathcal{G}$ includes the set D and obviously, no proper subset of C' possess this property too. Consequently, there is no $C'' \in \mathcal{E}_{ad}(\mathcal{G})$, s.t. $R^+_{\mathcal{G}}(C') \subset R^+_{\mathcal{G}}(C'')$ because C' is also preferred in $\mathcal{G}$. Hence, $C' \in \mathcal{E}_{ss}(\mathcal{G})$ contradicting the assumption.

In the light of Prop. 4 the corresponding question about the fibres (∞, ∞, ∞), $(\infty, 0, 0)$ and $(0, 0, 0)$ arises. The following proposition gives the (positive) answer:

Proposition 5. $\{(\infty,\infty,\infty),(\infty,0,0),(0,0,0)\}\subseteq\mathcal{S}_{(st,ss,pr,W)}$.

Proof. Consider the AFs $\mathcal{F}_1=(\{a\},\{(a,a)\})$, $\mathcal{F}_2=(\{a,b\},\{(b,b)\})$ and $\mathcal{F}_3=(\{a\},\varnothing)$. In consideration of Theorem 6 and Definition 7 [1] one may easily verify that the set $\{a\}$ possesses the claimed fibres w.r.t. the AFs $\mathcal{F}_1$, $\mathcal{F}_2$ and $\mathcal{F}_3$.

We have already shown that the minimal effort w.r.t. semi-stable semantics and weak expansions needed to enforce a desired set C cannot be 1. This raises the question about other natural numbers lying between 2 and ∞. The following proposition proves that there are infinitely many numbers n between 2 and ∞, s.t. $(\infty,n,0)$ is a fibre of the (st,ss,pr,W)-spectrum.

Proposition 6. *For any natural number $n\in\mathbb{N}$ there exists $k\in\mathbb{N}$, such that $n\le k\le 2n$ and $(\infty,k,0)\in\mathcal{S}_{(st,ss,pr,W)}$.*

Proof. We define the AF $\mathcal{F}_{\infty,n,0}=(A_{\infty,n,0},R_{\infty,n,0})$ where

$A_{\infty,n,0}=\{c_j,d_j,e_j\mid 1\le j\le n\}$ and

$R_{\infty,n,0}=\{(c_j,d_j),(d_j,c_j),(d_j,e_j)(e_j,e_j)\mid 1\le j\le n\}\cup\{(d_i,e_j)\mid 1\le i,j\le n\}$.

For the sake of clarity we present here an instantiation of the presented scheme, namely $\mathcal{F}_{\infty,3,0}$.

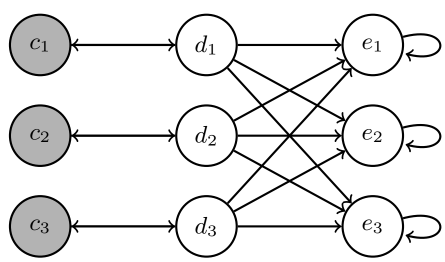

Fig. 4. The AF $\mathcal{F}_{\infty,3,0}$

The grey highlighted arguments belong to the set $C_3=\{c_1,c_2,c_3\}$ which is an instantiation of the scheme $C_n=\{c_j\mid 1\le j\le n\}$ (D_n, E_n are defined analogously). We claim that $N^{\mathcal{F}_{\infty,n,0}}_{st,\preceq_W}(C_n)=\infty$ and $N^{\mathcal{F}_{\infty,n,0}}_{pr,\preceq_W}(C_n)=0$. We observe that no superset of C_n is stable in $\mathcal{F}_{\infty,n,0}$ and furthermore, C_n itself is preferred in $\mathcal{F}_{\infty,n,0}$. Consequently (Theorem 6, Def. 7), the characteristics of C_n in case of stable and preferred semantics hold as claimed.

Consider now the semi-stable semantics. At first we will show that $N^{\mathcal{F}_{\infty,n,0}}_{ss,\preceq_W}(C_n)\ge n$. (proof by contradiction) Assume $N^{\mathcal{F}_{\infty,n,0}}_{ss,\preceq_W}(C_n)=n'<n$. This means, there is an AF $\mathcal{G}$, s.t. $d(\mathcal{F}_{\infty,n,0},\mathcal{G})=n'$, $\mathcal{F}_{\infty,n,0}\preceq_W\mathcal{G}$ and furthermore, there is a superset C'_n of C_n, s.t. $C'_n\in\mathcal{E}_{ss}(\mathcal{G})$. We deduce that $C'_n=C_n\cup G$ where G is a set of fresh arguments since

we consider weak expansions and furthermore, C_n does not possess proper supersets which are conflict-free in $\mathcal{F}_{\infty,n,0}$.

Since $n' < n$ is assumed it follows that there has to be an index j, s.t. $c_j \in C_n$ does not possesses attacks to arguments in $A(\mathcal{G})\backslash A_{\infty,n,0}$ (1) and $d_j \in D_n$ does not possesses attacks to arguments in G (2). Consider $S_n^j = \{c_i \mid 1 \leq i \leq n\}\backslash\{c_j\} \cup \{d_j\}$. Obviously, $R^+_{\mathcal{F}_{\infty,n,0}}(C_n) \subset R^+_{\mathcal{F}_{\infty,n,0}}(S_n^j) = A_{\infty,n,0}$ (3). We will show now that $S_n^j \cup G$ is admissible in $\mathcal{G}$ and it possesses a strictly greater range than C'_n in $\mathcal{G}$. Since we assumed $C'_n \in \mathcal{E}_{ss}(\mathcal{G})$, (2) and we are considering weak expansions the conflict-freeness of $S_n^j \cup G$ in $\mathcal{G}$ is implied. Furthermore, admissibility of $S_n^j \cup G$ in $\mathcal{G}$ holds because S_n^j is admissible in $\mathcal{F}_{\infty,n,0}$ and all potential attackers of arguments in G are counterattacked by at least one argument in $S_n^j \cup G$ (d_j counterattacks any e_i, any d_i where $i \neq j$ is counterattacked by c_i, an attacker $g' \in A(\mathcal{G})\backslash\{A_{\infty,n,0} \cup G\}$ is counterattacked by some $g \in G$ or some $c_i \in S_n^j$ because of the admissibility of C'_n and property (1)). Finally, $R^+_{\mathcal{G}}(C_n \cup G) \subset R^+_{\mathcal{G}}(S_n^j \cup G)$ has to hold because of properties (1) and (3). This contradicts the assumption that $C_n \cup G$ is semi-stable in $\mathcal{G}$.

Let us prove now that $N^{\mathcal{F}_{\infty,n,0}}_{ss,\preceq_S}(C_n) \leq 2n$. Let $C'_n = \{c'_1, ..., c'_n\}$ a set of fresh arguments and consider $\mathcal{G} = (A_{\infty,n,0} \cup C'_n, R_{\infty,n,0} \cup \{(c'_i, c'_i), (c_i, c'_i) \mid 1 \leq i \leq n\})$. Obviously, $d(F_{\infty,n,0}, G) = 2n$ and $F_{\infty,n,0} \preceq_W G$. One can easily verify that $C_n \in \mathcal{E}_{ss}(\mathcal{G})$. Finally, $N^{\mathcal{F}_{\infty,n,0}}_{ss,\preceq_W}(C_n) \leq 2n$ is shown.

It is an open question whether each number greater than 1 can appear as the characteristic of semi-stable semantics in a fibre, i.e. whether $\{(\infty, k, 0) \mid 2 \leq k < \infty\} \subseteq \mathcal{S}_{(st,ss,pr,W)}$. We would like to recall that it is already shown that in case of stable and preferred semantics, either a desired set C is already contained in an extension or C is not enforceable [1]. Consequently, an affirmative answer of the open question would imply a complete characterization of the (st, ss, pr, W)-spectrum.

3.3 A Note on the $(st, ss, pr, \mathcal{U})$-Spectrum

We use $\mathcal{U}$ to denote the universal relation among argumentation frameworks. In other words, we allow for arbitrary modifications including deletions of attacks and arguments. What consequences does this have for the corresponding spectrum? In contrast to the other considered spectra the $(st, ss, pr, \mathcal{U})$-spectrum is the first one proven to be positive. This means there are no cases where the enforcing of a certain set D is impossible. Furthermore, the $(st, ss, pr, \mathcal{U})$-spectrum is m.d.s. in analogy to the spectra w.r.t. arbitrary, normal and strong expansions.

Proposition 7. *The spectrum $\mathcal{S}_{(st,ss,pr,\mathcal{U})}$ is positive and m.d.s.*

Proof. Both properties follow immediately by applying Proposition 11 (positive) and Corollary 3 (m.d.s.) in [1]

A detailed analysis of the $(st, ss, pr, \mathcal{U})$-spectrum is part of future work. Due to the multitude of possibilities to modify a certain argumentation scenario if arbitrary modifications are allowed it is a hard task to show further properties. We want to mention that we conjecture that the considered spectrum is m.d.s.-complete (but were unable to find a proof so far).

4 Related Work and Conclusions

In this paper we presented various new results regarding the minimal change problem for Dung's abstract AFs. We introduced the so-called spectra which describe, for a collection of chosen semantics, the range of possible minimal efforts needed to enforce a set of arguments. We focused on stable, semi-stable and preferred semantics and were able to fully characterize the spectra for strong, normal and arbitrary expansions. This analysis revealed the surprising result that, although the three semantics are closely related, it may be arbitrarily more difficult to enforce arguments using stable rather than semi-stable semantics, and also using semi-stable rather than preferred semantics. The analysis of the spectrum for weak expansions turned out to be more difficult. Nevertheless, we were able to prove several useful (im)possibility results.

The presented work continues existing research on the dynamics of abstract argumentation. The paper [4] defines general principles (postulates) individual approaches may satisfy. The principles are illustrated for the special case of the grounded extension. Principles for the multiple extension case are left to further research. The authors of [5] focus on a particular type of change, namely the *addition of a single new argument* which interacts with previous arguments. They study the impact of such additions on the outcome of the argumentation framework, more particularly on the set of its extensions. The closely related paper [7] contains a theoretical study of the impact the *removal* of a single argument may have on the set of extensions of an argumentation framework. The article [8] develops a general method for handling updates of AFs based on a division. In a nutshell, the updated AF is divided into three parts: an unaffected, an affected, and a conditioning part. The status of arguments in the unaffected sub-framework remains unchanged, while the status of the affected arguments is computed in a special AF composed of an affected part and a conditioning part. It is shown that for specific semantics the extensions of the updated framework can be computed by combining the obtained results.

Booth and colleagues [12] investigated several quantitative distance measures for argumentation. In contrast to our work where the focus is on distances among different argumentation frameworks, the distance in that paper measures how far apart two labellings representing two complete extensions of the same argumentation framework are. This has applications in argument-based belief revision (e.g. if an agent is forced to switch to another extension and tries to identify the one closest to his original extension) and in judgement aggregation. Although the goals of this work are different from ours, it remains to be seen whether results from that work can be reused for our purposes.

Baroni et al. [13] introduce so-called input/output argumentation frameworks, an approach to characterize the behavior of an argumentation framework as sort of a black box with a well-defined external interface. The paper defines the notion of semantics decomposability and analyzes complete, stable, grounded and preferred semantics in this regard. It turns out that, under grounded, complete, stable and credulous preferred semantics, input/output argumentation frameworks with the same behavior can be exchanged without affecting the results of the evaluation of other interacting arguments. Since replaceability is one of the main motivations for studying equivalence notions, we plan to explore connections between equivalence and decomposability in the near future.

To the best of our knowledge the kind of questions analyzed in this paper have not been addressed before in argumentation. The analysis of spectra opens a number of new research directions which we want to pursue in the future. As just mentioned, the full characterization of the weak expansion spectrum is still open. Secondly, it would be useful to include further semantics (like grounded or ideal [3, 14]) in the analysis of spectra. Finally, it would be interesting to consider also a stronger form of enforcement where the enforced set of arguments has to be contained in *all* extensions rather than in *some* extension. We also might want to enforce a set of arguments C and at the same time exclude another, disjoint set D, that is, we might be interested in modifications leading to an AF possessing an extension E such that $C \subseteq E$ and $E \cap D = \emptyset$.

References

1. Baumann, R.: What does it take to enforce an argument? Minimal change in abstract argumentation. In: ECAI. (2012) 127–132
2. Bench-Capon, T.J.M., Dunne, P.E.: Argumentation in artificial intelligence. Artificial Intelligence **171**(10-15) (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. Artificial Intelligence **77**(2) (1995) 321 – 357
4. Boella, G., Kaci, S., van der Torre, L.: Dynamics in argumentation with single extensions: Abstraction principles and the grounded extension. In: Proceedings of the European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU). Volume LNAI 5590. (2009) 107–118
5. Cayrol, C., Dupin de Saint-Cyr, F., Lagasquie-Schiex, M.C.: Change in abstract argumentation frameworks: adding an argument. Journal of Artificial Intelligence Research **38** (2010) 49–84
6. Baumann, R., Brewka, G.: Expanding argumentation frameworks: Enforcing and monotonicity results. In: Proc. COMMA-10, IOS Press (2010) 75–86
7. Bisquert, P., Cayrol, C., Dupin de Saint-Cyr, F., Lagasquie-Schiex, M.C.: Change in argumentation systems: exploring the interest of removing an argument. In Benferhat, S., Grant, J., eds.: Proceedings of the International Conference on Scalable Uncertainty Management (SUM). Volume LNCS 6929., Springer (2011) 275–288
8. Liao, B., Jin, L., Koons, R.C.: Dynamics of argumentation systems: A division-based method. Artificial Intelligence **175**(11) (2011) 1790 – 1814
9. Scholz, H.: Ein ungelöstes Problem in der symbolischen Logik. Journal of Symbolic Logic **17** (1952) 160
10. Caminada, M.W.: Semi-stable semantics. In Dunne, P.E., Bench-Capon, T.J., eds.: Computational Models of Argument. Volume 144 of Frontiers in AI and Applications., IOS Press (2006) 121–130
11. Baroni, P., Giacomin, M.: Characterizing defeat graphs where argumentation semantics agree. In Simari, G., P., T., eds.: 1st International Workshop on Argumentation and Non-Monotonic Reasoning. (2007) 33–48
12. Booth, R., Caminada, M., Podlaszewski, M., Rahwan, I.: Quantifying disagreement in argument-based reasoning. In: Proc. AAMAS-12. (2012) 493–500
13. Baroni, P., Boella, G., Cerutti, F., Giacomin, M., van der Torre, L.W.N., Villata, S.: On input/output argumentation frameworks. In: COMMA. (2012) 358–365
14. Dung, P., Kowalski, R., Toni, F.: Dialectic proof procedures for assumption-based, admissible argumentation. Artificial Intelligence **170**(2) (2006) 114 – 159