On the Revision of Argumentation Systems: Minimal Change of Arguments Status

Sylvie Coste-Marquis, Sébastien Konieczny,
Jean-Guy Mailly, and Pierre Marquis

CRIL
Université d'Artois – CNRS
Lens, France
{coste,konieczny,mailly,marquis}@cril.fr

Abstract. In this paper, we investigate the revision issue for argumentation systems à la Dung. We focus on revision as minimal change of the arguments status. Contrarily to most of the previous works on the topic, the addition of new arguments is not allowed in the revision process, so that the revised system has to be obtained by modifying the attack relation, only. We introduce a language of revision formulae which is expressive enough for enabling the representation of complex conditions on the acceptability of arguments in the revised system. We show how AGM belief revision postulates can be translated to the case of argumentation systems. We provide a corresponding representation theorem in terms of minimal change of the arguments status. Several distance-based revision operators satisfying the postulates are also pointed out.

Keywords: argumentation, belief revision

1 Introduction

In this paper, we investigate the revision issue for abstract argumentation systems à la Dung [1]. Argumentation systems are directed graphs, where nodes correspond to arguments and arcs to attacks between arguments. In such systems, the status (acceptance) of each argument depends on the chosen acceptability semantics (grounded, preferred, stable – among others).

In his book [2] Gärdenfors introduced abstractly belief change as the operation allowing to change the epistemic status of a piece of information with respect to the epistemic state of an agent. There are three possible status: accepted, refused or undefined. And revision, contraction and expansion are defined as the possible transitions between these status, as illustrated by Figure 1.

Then Gärdenfors instantiates this general definition within a logical framework. But it is interesting to note that before defining belief change operators, one has to make precise the logical setting into which the status "accepted", "refused" and "undefined" are defined.

Similarly, if one wants to instantiate Gärdenfors general definition of belief change within Dung's argumentation theory, it is first necessary to define what

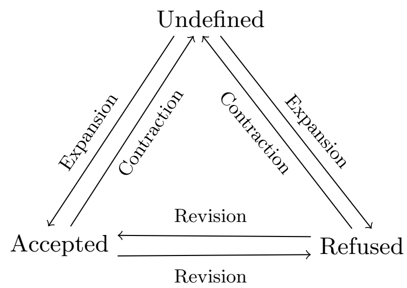

Fig. 1. Gärdenfors' epistemic transitions

are the available pieces of information and what these status mean. In Dung's argumentation theory, the basic pieces of information are the arguments of the system, and their status depends on the acceptability semantics under consideration.

Thus, it does not make sense to study the revision of argumentation systems directly on the attack graph, independently of any semantics. Stated otherwise, the revision of a given argumentation system under two different semantics may easily lead to two different results. For instance, in the case of the argumentation system given on Figure 2, we note that under the stable and preferred semantics, d belongs to every extension, whereas d does not belong to any extension in the case of the grounded semantics. So, revise to accept d does not need any change for stable and preferred semantics, but a change is required for the grounded semantics.

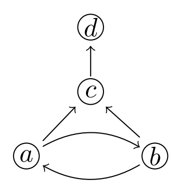

Fig. 2. The change needed to revise this system is not the same according to the semantics.

In this paper, we focus on revision as minimal change of the arguments status. To be more precise, under a chosen semantics, and given an argumentation system and a revision formula expressing how the status of some arguments has to be changed, we want to derive one or several argumentation systems which satisfy the revision formula, and are such that the corresponding extensions are as close as possible to the extensions of the input system. Contrarily to most of the previous works on the topic, the addition of new arguments is not allowed in the revision process, so that the revised system has to be obtained by modifying the attack relation, only. Especially, the revision formula does not indicate *why* the status of arguments have changed.

Minimal change of the attack graph can be considered as a criterion for defining the revised systems, but not as the main one, since, as explained above, the acceptance status of arguments is a more fundamental information. Accordingly, ensuring a minimal change of these status is more important than (and different from) ensuring a minimal change of the attack graph.

Such revision operators, where minimal change bears on the arguments status, can be very useful for applications of argumentation on social network debates [3]. When an agent A initiates a debate about an argument α, if another agent B does not agree with A about α but considers that A is trustworthy, B has to revise her beliefs to include α in at least one of her extensions. This kind of reasoning uses credulous inference, but it can be replaced by a skeptical inference and so the revised system can be computed with one of the operators defined in this paper.

The rest of the paper is organized as follows. After a short introduction to the Dung's theory of abstract argumentation in Section 2, we introduce a language of revision formulae which is expressive enough for enabling the representation of complex conditions on the acceptability of arguments in the revised system in Section 3. In Section 4 we show how AGM belief revision postulates can be translated to the case of argumentation systems. And we provide a corresponding representation theorem in terms of minimal change of the arguments status. In Section 5 several distance-based revision operators satisfying the postulates are pointed out. Then Section 6 discusses how to associate argumentation systems to the obtained sets of extensions, and discuss the issue of minimal change of the attack graph. Section 7 discusses some related work. Section 8 concludes the paper. Proofs are omitted for space reasons.

2 Preliminaries

We start with a very short introduction to Dung's theory of argumentation (see [1] for more details). A (finite) argumentation system (also referred to as an argumentation framework) is a pair $AF = \langle A, R \rangle$ where A is a (finite) set of so-called arguments and R is a binary relation over A (a subset of $A \times A$). In the following, A is supposed to contain at least two elements and the attack relation R is supposed to be irreflexive, i.e., self-contradicting arguments are rejected. AFs_A denotes the set of all such systems on the set of arguments A.

An argument $a \in A$ is acceptable with respect to a set of arguments $S \subseteq A$ whenever it is defended by the set, i.e., for every $b \in A$ s.t. $(b, a) \in R$, there exists $c \in S$ such that $(c, b) \in R$. We say that a subset S of A is conflict-free if and only if for every $a, b \in S$, we have $(a, b) \notin R$. A subset S of A is admissible for AF if and only if S is conflict-free and acceptable with respect to S. "Solutions" of an argumentation systems are sets S of arguments that can be, so to say, accepted together. Several semantics σ (especially, the complete semantics, the preferred semantics, the stable semantics, the grounded semantics) can be considered for capturing formally this notion, and each of them gives rise to a specific notion of extensions. For instance:

– S is a complete extension of AF if and only if it is an admissible set and every argument which is acceptable w.r.t. S belongs to S,
– S is a preferred extension of AF if and only if it is maximal (with respect to set inclusion) in the set of admissible sets for AF,
– S is a stable extension of AF if and only if S is conflict-free and $\forall a \in A \setminus S$, $\exists b \in S$ such that $(b, a) \in R$,
– S is the (unique) grounded extension of AF if and only if it is the smallest element (with respect to set inclusion) among the complete extensions.

$Ext_\sigma(AF)$ denotes the set of extensions of AF for the semantics σ. The epistemic status of any argument $a \in A$ with respect to the epistemic state represented by AF is then given by: a is accepted if a belongs to every $\varepsilon \in Ext_\sigma(AF)$, a is refused if a does not belong to any $\varepsilon \in Ext_\sigma(AF)$, and a is undefined in the remaining case.

Moreover, one introduce a notation about minimal elements of a set. First, $<$ denotes the strict part of $\leq$ and $\simeq$ denotes the indifference relation associated. Given a set E and a pre-order $\leq$, the minimal elements of E w.r.t. this pre-order are $\min(E, \leq) = \{e \in E | \nexists e' \in E, e' < e\}$.

3 On Revision Formulae

We want to define a revision setting for Dung's argumentation systems in which sophisticated revision formulae can be taken into account, and not only the fact that a given argument should be accepted or refused. To this end, we consider a logical language $\mathcal{L}_A$, where negation is used to denote the fact that a given argument should be refused, and formulae can be connected using conjunction and disjunction.

Definition 1 *Given $A = \{\alpha_1, \ldots, \alpha_k\}$ a set of arguments, $\mathcal{L}_A$ is the language generated by the following context-free grammar in BNF:*

$$arg ::= \alpha_1 | \ldots | \alpha_k$$
$$\Phi \quad ::= arg | \neg\Phi | (\Phi \wedge \Phi) | (\Phi \vee \Phi)$$

For instance, $\varphi_1 = (a \wedge ((\neg b \vee c) \wedge (b \vee \neg c)))$ expresses that in the revised epistemic state one wants a to be accepted and b and c to be both accepted or both refused. The epistemic status of such formulae φ_1 from $\mathcal{L}_A$ in an argumentation system $AF \in \mathsf{AFs}_A$ for a given semantics σ is given by:

Definition 2 *Let $\varepsilon \subseteq A$ and $\varphi \in \mathcal{L}_A$. The concept of* satisfaction *of φ by ε, noted $\varepsilon \mid\!\sim \varphi$, is defined inductively as follows:*

– *If $\varphi = a \in A$, then $\varepsilon \mid\!\sim \varphi$ iff $a \in \varepsilon$,*
– *If $\varphi = (\varphi_1 \wedge \varphi_2)$, $\varepsilon \mid\!\sim \varphi$ iff $\varepsilon \mid\!\sim \varphi_1$ and $\varepsilon \mid\!\sim \varphi_2$,*
– *If $\varphi = (\varphi_1 \vee \varphi_2)$, $\varepsilon \mid\!\sim \varphi$ iff $\varepsilon \mid\!\sim \varphi_1$ or $\varepsilon \mid\!\sim \varphi_2$,*
– *If $\varphi = \neg\psi$, $\varepsilon \mid\!\sim \varphi$ iff $\varepsilon \not\mid\!\sim \psi$.*

Then for any AF in AFs_A*, and any semantics* σ*, we say that:*

- φ *is* accepted *w.r.t. AF, noted* $AF \mid\!\sim_\sigma \varphi$*, if* $\varepsilon \mid\!\sim \varphi$ *for every* $\varepsilon \in Ext_\sigma(AF)$*,*
- φ *is* refused *w.r.t. AF, noted* $AF \mid\!\sim_\sigma \neg\varphi$*, if* $\varepsilon \mid\!\sim \varphi$ *for no* $\varepsilon \in Ext_\sigma(AF)$*,*
- φ *is* undefined *w.r.t. AF in the remaining case.*

From now on we call *candidate*[1] any subset ε of A. Candidates can be interpreted as models or counter-models of revision formulae. Continuing the previous example, if $A = \{a, b, c\}$, φ_1 is satisfied by the candidates from $\{\{a\}, \{a, b, c\}\}$. Thus, for the grounded semantics, φ_1 is accepted w.r.t. AF_1 with $R_1 = \{(b, c), (c, b)\}$ but is refused w.r.t. AF_2 with $R_2 = \{(a, b), (b, a)\}$.

Obviously enough, for any set $M = \{\varepsilon_1, \ldots, \varepsilon_k\}$ of candidates, there exists a formula $\varphi \in \mathcal{L}_A$ so that $M = \mathcal{A}_\varphi$, where $\mathcal{A}_\varphi = \{\varepsilon \subseteq A \mid \varepsilon \mid\!\sim \varphi\}$ is the set of candidates which satisfy φ. However, in the general case, $\mathcal{A}_\varphi$ is not the set of all σ extensions of an AF in AFs_A. Consider for instance, $A = \{a, b, c\}$ and $\varphi_1 = (a \wedge b \wedge c) \vee (a \wedge \neg b \wedge \neg c)$[2]. $\{a\}$ and $\{a, b, c\}$ are the two candidates satisfying φ_1, and there is no AF in AFs_A such that $Ext_\sigma(AF) = \{\{a\}, \{a, b, c\}\}$ for $\sigma =$ grounded, $\sigma =$ preferred or $\sigma =$ stable.

A concept of σ-consistency can then be defined as follows:

Definition 3 *A formula* $\varphi \in \mathcal{L}_A$ *is* σ-consistent[3] *iff there exists a set* $\mathcal{S} \neq \emptyset$ *of argumentation systems AF in* AFs_A *such that* $\mathcal{A}_\varphi = \bigcup_{AF \in \mathcal{S}} Ext_\sigma(AF)$*.*

When $A = \{a, b, c\}$, $\varphi_1 = (a \wedge ((\neg b \vee c) \wedge (b \vee \neg c)))$ is σ consistent for $\sigma =$ grounded, $\sigma =$ preferred or $\sigma =$ stable since $\{\{a\}, \{a, b, c\}\} = Ext_\sigma(AF_3) \cup Ext_\sigma(AF_4)$ where $R_3 = \{(a, b), (a, c)\}$ and $R_4 = \emptyset$. Contrastingly, $\varphi_2 = \neg a \wedge \neg b \wedge \neg c$ is grounded-consistent and preferred-consistent but not stable-consistent. $\varphi_3 = a \wedge \neg a$ neither is grounded-consistent nor is preferred-consistent, but is stable-consistent (consider AF_5 such that $R_5 = \{(a, b), (b, c), (c, a)\}$.)

4 On the Revision of Argumentation Systems

A revision operator on argumentation systems is a mapping associating a set of argumentation systems to the input argumentation system and the input revision formula:

Definition 4 *Given any set of arguments A, a* revision operator *on argumentation systems* $\star$ *is a mapping from* $\mathsf{AFs}_A \times \mathcal{L}_A$ *to* 2^{AFs_A}*.*

Clearly, the result of the revision of an argumentation system is not a unique argumentation system in the general case, but a set of argumentation systems. The reason is quite simple: there can be several possible results which have exactly the same maximum plausibility. So in this case there is no reason to

[1] That is to say "candidate to be an extension"
[2] Equivalent to $(a \wedge ((\neg b \vee c) \wedge (b \vee \neg c)))$.
[3] or simply *consistent* if the semantics is fixed.

choose a priori one of them (we will return to this point later on.) If this is problematic for a particular application, a selection function can be used as a tie-break rule for ensuring the unicity of the result (just like, for instance, the maxichoice selection function is considered in AGM belief revision [2].)

In order to define revision operators, our approach follows a two-step process. Intuitively, one first selects from the candidates satisfying the revision formula φ those which are as close as possible to the σ-extensions of AF. Then, one generates some argumentation systems such that the union of their σ extensions precisely coincides with the selected candidates. Of course, it is not the case that each mapping from $\mathsf{AFs}_A \times \mathcal{L}_A$ to 2^{AFs_A} is a reasonable revision operator. For instance, the constant, yet trivial operator defined by $AF \star \varphi = \emptyset$ should be discarded.

In order to identify interesting revision operators, one has to identify the logical properties that guarantee a rational behaviour. Such an axiomatic approach is standard in logic and the AGM postulates [4,5] have been pointed out for characterizing valuable revision operators in a logical setting. As in [6], we can revisit these postulates in a set-theoretic framework, here suited to the argumentation case.

Let $\mathcal{S}$ be a set of argumentation systems AF in AFs_A. For each semantics σ, we define the set $Ext_\sigma(\mathcal{S})$ of σ-extensions of $\mathcal{S}$ as $\bigcup_{AF \in \mathcal{S}} Ext_\sigma(AF)$. The counterpart of AGM postulates in the argumentation case is given by:

(AE1) $Ext_\sigma(AF \star \varphi) \subseteq \mathcal{A}_\varphi$
(AE2) If $Ext_\sigma(AF) \cap \mathcal{A}_\varphi \neq \emptyset$, then $Ext_\sigma(AF \star \varphi) = Ext_\sigma(AF) \cap \mathcal{A}_\varphi$
(AE3) If φ is σ-consistent, then $Ext_\sigma(AF \star \varphi) \neq \emptyset$
(AE4) If $\mathcal{A}_\varphi = \mathcal{A}_\psi$, then $Ext_\sigma(AF \star \varphi) = Ext_\sigma(AF \star \psi)$
(AE5) $Ext_\sigma(AF \star \varphi) \cap \mathcal{A}_\psi \subseteq Ext_\sigma(AF \star (\varphi \wedge \psi))$
(AE6) If $Ext_\sigma(AF \star \varphi) \cap \mathcal{A}_\psi \neq \emptyset$, then $Ext_\sigma(AF \star (\varphi \wedge \psi)) \subseteq Ext_\sigma(AF \star \varphi) \cap \mathcal{A}_\psi$

(AE1) states that the σ-extensions of the resulting set of argumentation systems must be among the candidates satisfying φ. **(AE2)** demands that if there are σ-extensions of the input system satisfying φ, then the resulting σ-extensions must coincide with them. **(AE3)** requires the resulting set of σ-extensions to be non-empty as soon φ is σ-consistent. **(AE4)** says that the revision by equivalent formulae must lead to the same results. The last two postulates **(AE5)** and **(AE6)** express a minimal change principle with respect to the arguments status: one expects to change as few as possible the status of arguments in the input system.

Interestingly, as in the logical case, we can derive a representation theorem which characterizes exactly the revision operators satisfying the postulates in a more constructive way. To this end, we first need to extend the notion of faithful assignment [5]:

Definition 5 *A* faithful assignment *is a mapping associating any argumentation system $AF = \langle A, R \rangle$ (under a semantics σ) with a total pre-order $\leq^\sigma_{AF}$ on the set of candidates such that:*
1. *if $\varepsilon_1 \in Ext_\sigma(AF)$ and $\varepsilon_2 \in Ext_\sigma(AF)$, then $\varepsilon_1 \simeq^\sigma_{AF} \varepsilon_2$,*
2. *if $\varepsilon_1 \in Ext_\sigma(AF)$ and $\varepsilon_2 \notin Ext_\sigma(AF)$, then $\varepsilon_1 <^\sigma_{AF} \varepsilon_2$.*

The representation theorem can then be stated as follows:

Proposition 1 *Given a semantics σ, a revision operator $\star$ satisfies the rationality postulates **(AE1) - (AE6)** iff there exists a faithful assignment which matches every system $AF = \langle A, R\rangle$ to a total pre-order $\leq^{\sigma}_{AF}$ so that*

$$Ext_\sigma(AF \star \varphi) = \min(\mathcal{A}_\varphi, \leq^{\sigma}_{AF})$$

This theorem is useful for defining operators satisfying the rationality postulates, as those presented in the next section.

5 Distance-Based Revision

Let us now present some distance-based revision operators satisfying the rationality postulates **(AE1) - (AE6)**. Let d be any distance on 2^A, for instance, the Hamming distance given by $d_H(\varepsilon_1, \varepsilon_2) = |(\varepsilon_1 \setminus \varepsilon_2) \cup (\varepsilon_2 \setminus \varepsilon_1)|$. Given $\varepsilon \in 2^A$ and $\mathcal{E} \subseteq 2^A$, d can be extended to a "distance" between ε and $\mathcal{E}$, by stating that $d(\varepsilon, \mathcal{E}) = \min_{\varepsilon' \in \mathcal{E}} d(\varepsilon, \varepsilon')$. For any argumentation system $AF \in \mathsf{AFs}_A$, this distance induces a total pre-order between candidates $\varepsilon_1, \varepsilon_2 \in 2^A$ given by

$$\varepsilon_1 \leq^{\sigma,d}_{AF} \varepsilon_2 \text{ iff } d(\varepsilon_1, Ext_\sigma(AF)) \leq d(\varepsilon_2, Ext_\sigma(AF)).$$

On this ground, revision operators can be defined by:

Definition 6 *Let σ be any given semantics. A* distance-based revision operator *$\star^d$ is any revision operator for which there exists a distance d on 2^A such that for every AF and every φ, we have $Ext_\sigma(AF \star^d \varphi) = \min(\mathcal{A}_\varphi, \leq^{\sigma,d}_{AF})$.*

Proposition 2 *Let σ be any semantics. Any distance-based revision operator $\star^d$ satisfies the rationality postulates **(AE1) - (AE6)**.*

Let us now define another family of distance-based operators, which take advantage of labellings. Let us first recall that, instead of using extensions, the solutions of an argumentation system can be expressed using the concept of labelling [7]. Formally, a labelling is a mapping L associating a label *in*, *undec* or *out* with every argument of the set A. The stronger notion of reinstatement labelling depends on the attack relation R: an argument a is labelled *in* iff every argument attacking a is *out*; an argument a is *out* iff there exists an argument *in* attacking a; an argument is *undec* iff it is neither *in* nor *out*. These reinstatement labellings correspond to Dung's complete extensions in a bijective way. Thus, for any complete extension ε, the associated reinstatement labelling is such that every argument $a \in \varepsilon$ is *in*, every argument attacked by an argument *in* is *out*, every other argument is *undec*. Conversely, for every reinstatement labelling L, the corresponding extension $E(L)$ is the set of arguments labelled *in* by the labelling. Other semantics can also be encoded with labellings. We introduce some notations: L_φ is the set of labellings L such that $E(L) \in \mathcal{A}_\varphi$. Given a set of labellings Lab, $E(Lab) = \{E(L) | L \in Lab\}$. Finally, given a system AF

and a semantics σ, $Labs_\sigma(AF)$ denotes the set of labellings corresponding to the σ-extensions of AF.

Labellings, which bring richer information than extensions, can be used to define interesting distance-based revision operators. Indeed, consider the following notion of edition distance:

Definition 7 *Let m, n, o be three integers, let L_1 and L_2 be two labellings, and let X and Y be any of in, out and undec.*
An edition distance $d_{(m,n,o)}$ *between labellings is defined as:*

$$d_{(m,n,o)}(L_1, L_2) = \sum_{a \in A} ad(L_1(a), L_2(a)),$$

where

- $ad(in, out) = m$
- $ad(in, undec) = n$
- $ad(out, undec) = o$
- $ad(X, Y) = ad(Y, X)$
- $ad(X, X) = 0$

Proposition 3 *Let m, n, o be three integers. $d_{(m,n,o)}$ is a distance.*

An interesting point to note is that these edition distances are not necessarily neutral or symmetrical. We call neutral a distance such that $ad(in, undec) + ad(undec, out) = ad(in, out)$ and symmetrical a distance such that $ad(in, undec) = ad(undec, out)$. Defining non symmetrical distances is a way for instance to favor acceptance of arguments over reject, by choosing $ad(in, undec) = 1$, $ad(out, undec) = 9$, and $ad(in, out) = 10$.

For any distance $d_\mathcal{L}$ between labellings, we can define a pre-order $\leq_{AF}^{\sigma, d_\mathcal{L}}$ between labellings as we did it for candidates:

$$L_1 \leq_{AF}^{\sigma, d_\mathcal{L}} L_2 \text{ iff } d_\mathcal{L}(L_1, L_\sigma(AF)) \leq d_\mathcal{L}(L_2, L_\sigma(AF)).$$

Definition 8 *Let σ be any given semantics. A* labelling-distance-based revision operator $\star_{d^\mathcal{L}}$ *is any revision operator for which there exists a distance $d^\mathcal{L} = d_{(m,n,o)}$ on 2^A such that for every AF and every φ, we have $Labs_\sigma(AF \star_{d^\mathcal{L}} \varphi) = \min(L_\varphi, \leq_{AF}^{\sigma, d_\mathcal{L}})$.*

The following example illustrates the impact of the chosen distance on the revised system:

Example 1 *Let σ be the stable semantics. We revise the system AF_6 below by the formula $\varphi = \neg d \wedge \neg e$.*

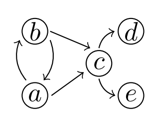

$Ext_\sigma(AF_6) = \{\{a, d, e\}, \{b, d, e\}\}$, the associated stable labellings are $\{(a, in), (b, out), (c, out), (d, in), (e, in)\}$ and $\{(a, out), (b, in), (c, out), (d, in),$

$(e, in)\}$. When we revise AF_6 by φ using the distance-based operator induced by the distance $d_{(1,9,10)}$ on labellings, one gets as result a system of whom the labellings are $\{(a, in), (b, out), (c, out), (d, out), (e, out)\}$ and $\{(a, out), (b, in), (c, out), (d, out), (e, out)\}$. When the distance $d_{(9,1,10)}$ is used, one obtains $\{(a, in), (b, out), (c, out), (d, undec), (e, undec)\}$ and $\{(a, out), (b, in), (c, out), (d, undec), (e, undec)\}$ as labellings of the result systems.
If the generation of the systems take account of the labellings, the structure of the result graph will be different : when the refused arguments are out, it means that there exists an attack from an accepted argument to a refused argument. When the arguments are undec, those attacks do not exist.
With the first distance $d_{(1,9,10)}$, it is less costly to change an argument from in to out than to undec. Such a distance allow to choose candidates which refuses arguments by . Contrarily, the distance $d_{(9,1,10)}$ allow to choose candidates which accept more arguments.
The choice of a particular distance allow to influence result of the revision.

Like operators based on extensions, distance-based operators using labeling exhibit good logical properties:

Proposition 4 *Let σ be any semantics. Any labelling-distance-based revision operator $\star_{d^{\mathcal{L}}}$ satisfies the rationality postulates **(AE1)** - **(AE6)**.*

6 Revision at the System Level

The operators defined in the previous section focus on the candidates that are as close as possible to the extensions of the input system. However, they do not indicate how to generate the corresponding argumentation systems, i.e., the argumentation systems such that the union of their extensions coincides with the selected candidates[4].

In order to do this job, we consider a mapping $\mathcal{AF}_\sigma$ from $2^{2^{\mathcal{A}}}$ to 2^{AFs_A}, called generation operator, that associates with any set $\mathcal{C}$ of candidates a set of argumentation systems such that $Ext_\sigma(\mathcal{AF}_\sigma(\mathcal{C})) = \mathcal{C}$.

Observe that, whatever the semantics σ, such a mapping $\mathcal{AF}_\sigma$ exists. This comes easily from the fact that:

Proposition 5 *Whatever the semantics σ, for every non empty set $\mathcal{C}$ of candidates from 2^A, such that $\emptyset \notin \mathcal{C}$, there exists a finite set $\mathcal{S} \subseteq \mathsf{AFs}_A$ such that $\mathcal{C} = \bigcup_{AF \in \mathcal{S}} Ext_\sigma(AF)$.*

The point is that every candidate C can be associated with an argumentation system AF such that C is the unique σ-extension of it. For instance, if $A = \{a, b, c\}$ and $C = \{a, b\}$, AF given by $R = \{(a, c), (b, c)\}$ does the job whatever the semantics σ. This method allows to prove that a set of systems exists,

[4] The construction of systems from labellings is still an open question, this section only deals with generation of systems from candidates.

but it does not study the generation of a minimal set of systems corresponding to the candidates. We give a first approach to this problem in the rest.

On this ground, revision operators can then be defined as follows:

Definition 9 *Given a semantics σ, a faithful assignment that matches every argumentation system to a total pre-order $\leq^{\sigma}_{AF}$, and a generation operator $\mathcal{AF}_{\sigma}$, the corresponding* basic revision operator $\star$ *is defined by:*

$$AF \star \varphi = \mathcal{AF}_{\sigma}(\min(\mathcal{A}_{\varphi}, \leq^{\sigma}_{AF}))$$

One of the key results of the paper is that:

Proposition 6 *Every basic revision operator $\star$ satisfies the postulates **(AE1)**-**(AE6)**.*

Basic operators deal only with minimality of change of arguments status. Indeed, the rationality postulates asks for preserving as much as possible the status of arguments in the input system: doing so while ensuring that the revision formula is satisfied does not usually imply a minimal change of the attack relation, and vice-versa. As a matter of illustration, consider the argumentation systems AF_7, AF_8, AF_9, and AF_{10} represented respectively in Figure 3.a, 3.b, 3.c, 3.d.

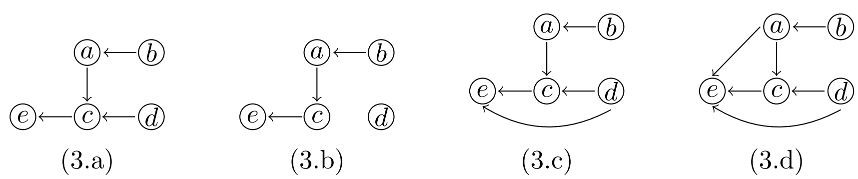

Fig. 3. Minimal Change

Suppose that our goal is to reject e, that is to get a system so that e does not appear in any extension. We consider the revision formula $\varphi = \neg e$. A minimal change on the attack relation of AF_7 leads to AF_8 or AF_9: each of them differs with AF_7 on a single attack. This contrasts with AF_{10} since the change on the attack relation required to go from AF_7 to AF_{10} is strictly greater than the change on the attack relation required to go from AF_7 to AF_9. Each of these four systems has a unique extension whatever the semantics: $\{b, d, e\}$ for AF_7, $\{b, c, d\}$ for AF_8, and $\{b, d\}$ for AF_9 and AF_{10}. Hence, the change on the status of arguments achieved when going from AF_7 to AF_9 or AF_{10} is strictly smaller than the change on the status of arguments achieved when going from AF_7 to AF_8.

A simple generation operator $\mathcal{AF}_{\sigma}$ consists in generating sets $\mathcal{S}$ of increasing cardinality n until the condition $\mathcal{C} = \bigcup_{AF \in \mathcal{S}} Ext_{\sigma}(AF)$ is satisfied and then returns the union of all the minimal-size $\mathcal{S}$ which satisfies it. However, using this simple generation operator does not ensure to obtain an optimal set of argumentation systems, when optimality is understood either as the minimal number of argumentation systems required to generate the candidates, or as minimal change of the attack relation.

Minimizing the change of the attack relation can nevertheless be taken into account in the definition of the generation operator (hence as a second criterion, with a smaller priority than the criterion capturing the minimal change of the arguments status.) This can be easily done by considering a distance dg on the argumentation graphs. One simple example is the Hamming distance, given by $dg_H(AF_1, AF_2) = |(\mathcal{R}_1 \setminus \mathcal{R}_2) \cup (\mathcal{R}_2 \setminus \mathcal{R}_1)|$. But one can also choose more elaborate edition distances such as those given in [8]. The dg_H distance between two argumentation systems corresponds to the number of attacks that must be added or removed to make them identical. Each distance dg induces a pre-order between argumentation systems, defined by $AF_1 \leq_{AF}^{dg} AF_2$ iff $dg(AF_1, AF) \leq dg(AF_2, AF)$. This pre-order can be used to filter out unnecessary argumentation systems produced by the generation operator. However, a simple minimization w.r.t. $\leq_{AF}^{dg}$ of the output of the generation operator is not sufficient in general, as illustrated by the following example:

Example 2 *Suppose that the candidates are $\mathcal{C} = \{C_1, C_2\}$, and the generation operator $\mathcal{AF}_\sigma$ applied on $\mathcal{C}$ gives $\mathcal{AF}_\sigma(\mathcal{C}) = \{AF_{11}, AF_{12}\}$ so that, for the considered semantics σ, $Ext_\sigma(AF_{11}) = \{C_1\}$ and $Ext_\sigma(AF_{12}) = \{C_2\}$. If AF_{11} and AF_{12} are not at equal distance from the input system AF (that is $AF_{11} <_{AF} AF_{12}$ or $AF_{12} <_{AF} AF_{11}$), one of the systems is not minimal and must be removed. So an extension is lost.*

To tackle this problem, we consider selection functions removing argumentation systems when this does not change the set of extensions:

Definition 10 *Given a semantics σ, a set of candidates $\mathcal{C}$, a pre-order $\leq_{AF}$ and a generation operator $\mathcal{AF}_\sigma$, let us define the* $\mathcal{C}$-minimum-cover selection function *$\gamma_\mathcal{C}$ by:*

- $\gamma_\mathcal{C}(\mathcal{AF}_\sigma(\mathcal{C})) \subseteq \mathcal{AF}_\sigma(\mathcal{C})$
- $Ext_\sigma(\gamma_\mathcal{C}(\mathcal{AF}_\sigma(\mathcal{C}))) = Ext_\sigma(\mathcal{AF}_\sigma(\mathcal{C}))$
- *$AF \in \mathcal{AF}_\sigma(\mathcal{C})$ and $AF \notin \gamma_\mathcal{C}(\mathcal{AF}_\sigma(\mathcal{C}))$ only if $\exists AF' \in \mathcal{AF}_\sigma(\mathcal{C})$ so that $AF' <_{AF} AF$*

Based on this notion, full revision operators can be defined as follows:

Definition 11 *Given a semantics σ, a faithful assignment $\leq_{AF}^{\sigma,d}$, a generation operator $\mathcal{AF}_\sigma$, a pre-order $\leq_{AF}^{dg}$ over AFs_A and a $\mathcal{C}$-minimum-cover selection function γ, the corresponding* full revision operator $\circ^{\langle \leq_{AF}^{\sigma,d}, \leq_{AF}^{dg}, \gamma \rangle}$ *is defined by:*

$$AF \circ^{\langle \leq_{AF}^{\sigma,d}, \leq_{AF}^{dg}, \gamma \rangle} \varphi = \gamma_{\min(\mathcal{A}_\varphi, \leq_{AF}^{\sigma,d})}(\mathcal{AF}_\sigma(\min(\mathcal{A}_\varphi, \leq_{AF}^{\sigma,d})))$$

Clearly full revision operators are basic revision operators.

Since the minimization of the attack relation is done without modifying the selected candidates, we get that:

Proposition 7 *Every full revision operator* $\circ^{\langle \leq^{\sigma,d}_{AF}, \leq^{dg}_{AF}, \gamma\rangle}$ *satisfies the postulates* ***(AE1)-(AE6)****.*

For illustration purposes, we now provide some examples of revision operators. For these examples we will use the following $\mathcal{C}$-minimum-cover selection function $\gamma_1^{\leq^{dg}_{AF}}$:

- V is an ordered list defined from $\mathcal{AF}_\sigma(\mathcal{C})$ by sorting it in descending order[5] w.r.t. $\leq^{dg}_{AF}$.
- For i from 1 to n do
  if $\exists V_j \in V$ s.t. $V_j <^{dg}_{AF} V_i$ and $Ext_\sigma(V \backslash V_i) = Ext_\sigma(V)$, then $V := V \backslash V_i$.
- $\gamma_1^{\leq^{dg}_{AF}} := V$.

With this method, the systems removed from V are those which are the farthest from the input AF.

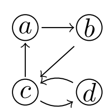

Fig. 4. The system AF_{13}

Let us now illustrate those concepts by considering the revision of the argumentation system AF_{13} represented at Figure 4, using the full revision operator defined from the Hamming distance between candidates, the Hamming distance between graphs and the previous selection function $\gamma_1^{\leq^{dg}_{AF}}$.

The set of extensions of AF_{13} is $\{\{a, d\}\}$ for the preferred and stable semantics, and its grounded extension is $\emptyset$. Hence, whatever the chosen semantics, we have $AF_{13} \mid\!\sim \neg b$. Suppose now that we want b to be accepted, so we make a revision with $\varphi = b$.

For the preferred semantics and the stable semantics, we first build the set of candidates which satisfy b so that the Hamming distance with $\{\{a, d\}\}$ is minimal. This set is $\{\{a, b, d\}\}$. Then we build the argumentation systems which cover this candidate, focusing on those which are at minimal distance to AF_{13}. The result is AF_{14}, given at Figure 5.a.

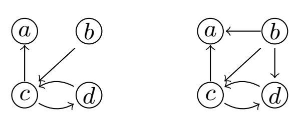

(5.a) (5.b)

Fig. 5. The result of the revision of AF_{13} by b

[5] There exists several ordered lists, depending on the way to sort elements which are equals w.r.t. the pre-order. We choose arbitrarily one of these lists.

We proceed similarly for the grounded semantics: the set of extensions $\{\{b\}\}$ is first computed, and then the unique system, AF_{15}, represented on Figure 5.b is generated.

A last point we would like to discuss is the fact that a basic revision operator $\star$ outputs a set of argumentation systems, and not a single argumentation system in the general case. Actually, this is a consequence of the expressiveness of the language of revision formulae we want to consider. In order to illustrate it, consider $A = \{a, b, c, d\}$, and AF_{16} as represented in Figure 6.

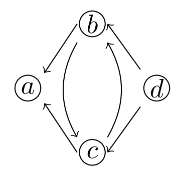

Fig. 6. The system AF_{16}

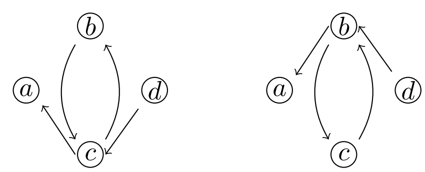

Fig. 7. Revision of AF_{16}

The extensions of AF_{16} are the same for the grounded, stable and preferred semantics, $Ext(AF_{16}) = \{\{a, d\}\}$. Let $\varphi = b \vee c$. Observe that b and c play symmetric roles, both in AF_{16} and in φ. When computing the result of the revision with the full revision operator based on Hamming distance between candidates, Hamming distance on the graph and the selection function $\gamma_1^{\leq_{AF}^{dg}}$, we obtain two candidates $\{a, b, d\}$ and $\{a, c, d\}$, and two corresponding argumentation systems AF_{17} and AF_{18} (given respectively at Figure 7.a and Figure 7.b). Choosing one of these systems would require to accept some arbitrariness given the symmetric roles of b and c.

Finally, a few words about the iteration issue for revision. Given that the output of a revision operator is not a single system in the general case but a set of such systems, revision cannot be directly iterated. Nevertheless, it is easy to extend revision operators so that sets of argumentation systems can be accepted as input. Basically, $\mathcal{S} = \{AF_1, \ldots, AF_n\} \star \varphi$ can be defined as $\bigcup_{AF_i \in \mathcal{S}} \bigcup_{AF \in AF_i \star \varphi} AF$.

7 Related Work

Some previous works have already considered the change issue for argumentation systems à la Dung.

Thus, Boella, Kaci and Van der Torre, [9,10] have studied abstraction and refinement principles. An abstraction is a reduction of the attack relation or the set of arguments, whereas a refinement is the addition of attacks or arguments to the system. The authors focused on the study of semantics which ensure the existence of a unique extension (for instance, the grounded extension), and they formulated some principles of the form "if we do this particular change, then the extension of the result is like this". They identified some principles satisfied by the grounded semantics.

Cayrol, Dupin de Saint-Cyr and Lagasquie-Schiex [11] studied the addition of an argument to an argumentation system. They stated some properties that can be satisfied when a change occurs in an argumentation system, and pointed out those which are satisfied (and under which conditions) when an argument (and the attacks concerning it) is added to the graph. With Bisquert, they did a similar study about the deletion of an argument [12].

Baumann [13] also studied the minimal change problem in abstract argumentation. He reported some bounds on the number of modifications to make on an argumentation system so as to enforce a given set of arguments. These bounds depend on the semantics and the type of change allowed.

All these works are basically concerned by the modification of the attack relation. As such, they differ in a significant way from the change problem as studied in this paper, where change primarily lies on the status of arguments (and secondary only on the attack relation.)

8 Conclusion

In this paper, we investigated the revision problem for abstract argumentation systems à la Dung. We focused on revision as minimal change of the arguments status. We introduced a language of revision formulae which is expressive enough for enabling the representation of complex conditions on the acceptability of arguments in the revised system. We showed how AGM belief revision postulates can be translated to the case of argumentation systems. We provided a corresponding representation theorem in terms of minimal change of the arguments status, and pointed out several distance-based revision operators satisfying the postulates.

As a future work we first plan to encode our revision operators by representing argumentation systems with logical constraints (in a similar way to Besnard and Doutre [14]), so as to be able to benefit from the power of constraint solvers to compute revised systems. The study of other change operations on argumentation systems is another perspective for further research.

Moreover, the association of a minimal set of argumentation systems to a set of candidates is very interesting, including for other applications than revision. So we can deepen this question. The same question for the case of labellings (see Example 1) is still open and will be studied.

References

1. 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
2. Gärdenfors, P.: Knowledge In Flux. Cambridge University Press, Cambridge, UK (1988)
3. Gabbriellini, S., Torroni, P.: MS Dialogues: Persuading and getting persuaded, A model of social network debates that reconciles arguments and trust. (2013)
4. Alchourrón, C.E., Gärdenfors, P., Makinson, D.: On the logic of theory change : Partial meet contraction and revision functions. Journal of Symbolic Logic **50** (1985) 510–530
5. Katsuno, H., Mendelzon, A.O.: Propositional knowledge base revision and minimal change. Artificial Intelligence **52** (1991) 263–294
6. Qi, G., Liu, W., Bell, D.A.: Knowledge base revision in description logics. In: Proceedings of the European Conference on Logics in Artificial Intelligence (JELIA'06). Volume 4160 of Lecture Notes in Computer Science., Springer (2006) 386–398
7. Caminada, M.: On the issue of reinstatement in argumentation. In: Proceedings of the European Conference on Logics in Artificial Intelligence (JELIA'06), Springer (2006) 111–123
8. Coste-Marquis, S., Devred, C., Konieczny, S., Lagasquie-Schiex, M.C., Marquis, P.: On the merging of Dung's argumentation systems. Artificial Intelligence Journal **171** (2007) 730–753
9. Boella, G., Kaci, S., van der Torre, L.: Dynamics in argumentation with single extensions: attack refinement and the grounded extension. In: Proceedings of the International Conference on Autonomous Agents and Multiagents Systems (AAMAS'09). (2009) 1213–1214
10. 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 Conferences on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'09). Volume 5590 of Lecture Notes in Computer Science., Springer (2009) 107–118
11. Cayrol, C., de Saint-Cyr, F.D., Lagasquie-Schiex, M.C.: Change in abstract argumentation frameworks: Adding an argument. Journal of Artificial Intelligence Research **38** (2010) 49–84
12. Bisquert, P., Cayrol, C., de Saint-Cyr, F.D., Lagasquie-Schiex, M.C.: Change in argumentation systems: Exploring the interest of removing an argument. In: Proceedings of the International Conference on Scalable Uncertainty Management (SUM'11). Volume 6929 of Lecture Notes in Computer Science., Springer (2011) 275–288
13. Baumann, R.: What does it take to enforce an argument? minimal change in abstract argumentation. In: Proceedings of the European Conference on Artificial Intelligence (ECAI'12). (2012) 127–132
14. Besnard, P., Doutre, S.: Checking the acceptability of a set of arguments. In: Proceedings of the 10th International Workshop on Non-Monotonic Reasoning (NMR'04). (2004) 59–64