A formal analysis of logic-based argumentation systems

Leila Amgoud Philippe Besnard

IRIT-CNRS, 118, route de Narbonne,
31062 Toulouse Cedex 4 France
{amgoud, besnard}@irit.fr

Abstract. Dung's abstract argumentation model consists of a set of arguments and a binary relation encoding attacks among arguments. Different acceptability semantics have been defined for evaluating the arguments. What is worth noticing is that the model completely abstracts from the applications to which it can be applied. Thus, it is not clear what are the results that can be returned in a given application by each semantics. This paper answers this question. For that purpose, we start by plunging the model in a real application. That is, we assume that we have an inconsistent knowledge base (KB) containing formulas of an *abstract monotonic logic*. From this base, we show how to define arguments. Then, we characterize the different semantics in terms of the subsets of the KB that are returned by each extension. We show a full correspondence between maximal consistent subbases of a KB and maximal conflict-free sets of arguments. We show also that stable and preferred extensions choose *randomly* some consistent subbases of a base. Finally, we investigate the results of three argumentation systems that use well-known attack relations.

1 Introduction

Argumentation has become an Artificial Intelligence keyword for the last fifteen years, especially for handling inconsistency in knowledge bases (e.g. [2, 6, 16]), for decision making under uncertainty (e.g. [4, 7, 13]), and for modeling interactions between agents (e.g. [3, 14, 15]).
One of the most abstract argumentation systems was proposed in [12]. This system consists of a set of arguments and a binary relation encoding attacks among arguments. Different acceptability semantics were also defined for evaluating the arguments. A semantics defines the conditions under which a given set of arguments is declared acceptable. What is worth noticing is that the system completely abstracts from the applications to which it can be applied. Thus:

1. The origin and the structure of both arguments and the attack relation are not specified. This is seen in the literature as an advantage of this model. However, some of its instantiations lead to undesirable results. This is due to a lack of *methodology* for defining these two main components of the system.

2. The different acceptability semantics capture mainly different properties of the graph associated to the system. While it is clear that those properties are nice and meaningful, it is not clear whether they really make sense in concrete applications. It is not even clear what are the results that can be returned, in a given application, by each semantics.

In [1], we proposed an extension that fills in the gap between Dung's system and the applications. The idea was to consider all the ingredients involved in an argumentation problem. We started with an abstract monotonic logic which consists of a set of formulas and a consequence operator. We have shown how to build arguments from a knowledge base using the consequence operator of the logic, and how to choose an appropriate attack relation.

Starting from this class of logic-based argumentation systems, this paper characterizes the different semantics in terms of subsets of the knowledge base that are returned by each extension. The results we got show that there is a full correspondence between maximal consistent subbases of a KB and maximal conflict-free sets of arguments. They also show that stable and preferred extensions choose *randomly* some consistent subbases of a base, which are not necessarily maximal in case of preferred extensions. Finally, the paper studies the properties of three well-known argumentation systems, namely the ones that use resp. rebut, strong rebut and undercut as attack relations. We show that the system that is based on undercut returns sound results at the cost of redundancy: duplicating many arguments with exactly the same support.

The paper is organized as follows: Section 2 recalls our extension of Dung's system. Section 3 presents an analysis of acceptability semantics. Section 4 investigates the properties of particular systems.

2 An extension of Dung's abstract system

In [1], we proposed an extension of Dung's abstract argumentation system. That extension defines all the ingredients involved in argumentation ranging from the logic used to define arguments to the set of conclusions to be inferred from a given knowledge base. A great advantage of this extension is that it is abstract since it is grounded on an *abstract monotonic logic*. According to [17], such logic is defined as a pair $(\mathcal{L}, \mathsf{CN})$ where members of $\mathcal{L}$ are called *well-formed formulas*, and CN is a *consequence operator*. CN is any function from $2^{\mathcal{L}}$ to $2^{\mathcal{L}}$ that satisfies the following axioms:

1. $X \subseteq \mathsf{CN}(X)$ **(Expansion)**
2. $\mathsf{CN}(\mathsf{CN}(X)) = \mathsf{CN}(X)$ **(Idempotence)**
3. $\mathsf{CN}(X) = \bigcup_{Y \subseteq_f X} \mathsf{CN}(Y)$ **(Finiteness)**
4. $\mathsf{CN}(\{x\}) = \mathcal{L}$ for some $x \in \mathcal{L}$ **(Absurdity)**
5. $\mathsf{CN}(\emptyset) \neq \mathcal{L}$ **(Coherence)**

$Y \subseteq_f X$ means that Y is a finite subset of X. Intuitively, $\mathsf{CN}(X)$ returns the set of formulas that are logical consequences of X according to the logic in question. It can easily be shown from the above axioms that CN is a closure operator, and satisfies *monotonicity* (i.e. for $X, X' \subseteq \mathcal{L}$, if $X \subseteq X'$ then $\mathsf{CN}(X) \subseteq \mathsf{CN}(X')$). Note that a wide variety of logics can be viewed as special cases of Tarski's notion of an abstract logic (classical logic, intuitionistic logic, modal logic, temporal logic, ...). Formulas of $\mathcal{L}$ encode both defeasible and undefeasible information while the consequence operator CN is used for defining arguments.

Definition 1 (Consistency) *Let $X \subseteq \mathcal{L}$. X is* consistent *in logic $(\mathcal{L}, \mathsf{CN})$ iff $\mathsf{CN}(X) \neq \mathcal{L}$. It is* inconsistent *otherwise.*

In simple English, this says that X is consistent iff its set of consequences is not the set of all formulas. The coherence requirement (absent from Tarski's original proposal but added here to avoid considering trivial systems) forces $\emptyset$ to always be consistent - this makes sense for any reasonable logic as saying emptiness should intuitively be consistent.

We start with an abstract logic $(\mathcal{L}, \mathsf{CN})$ from which the notions of argument and attacks between arguments are defined. More precisely, arguments are built from a knowledge base, say Σ, containing formulas of the language $\mathcal{L}$.

Definition 2 (Argument) *Let Σ be a knowledge base. An* argument *is a pair (X, x) such that:*

1. $X \subseteq \Sigma$
2. *X is consistent*
3. $x \in \mathsf{CN}(X)$
4. *$\nexists X' \subset X$ s.t. X' satisfies the three above conditions*

Notation: $\mathtt{Supp}$ and $\mathtt{Conc}$ denote respectively the *support* X and the *conclusion* x of an argument (X, x). For $\mathcal{S} \subseteq \Sigma$, $\mathtt{Arg}(\mathcal{S})$ denotes the set of all arguments that can be built from $\mathcal{S}$ by means of Definition 2.

Since CN is monotonic, argument construction is a monotonic process (i.e. $\mathtt{Arg}(\Sigma) \subseteq \mathtt{Arg}(\Sigma')$ whenever $\Sigma \subseteq \Sigma' \subseteq \mathcal{L}$).

We have shown in [1] that in order to satisfy the consistency rationality postulate proposed in [9], the attack relation should be chosen in an appropriate way. Otherwise, unintended results may be returned by the argumentation system. An appropriate relation, called *valid*, should ensure that the set of formulas used in arguments of any non-conflicting (called also *conflict-free*) set of arguments is consistent.

Notation: Let $\mathcal{B} \subseteq \mathtt{Arg}(\Sigma)$. $\mathtt{Base}(\mathcal{B}) = \bigcup_{a \in \mathcal{B}} \mathtt{Supp}(a)$.

Definition 3 (Valid attack relation) *Let $\mathtt{Arg}(\Sigma)$ be a set of arguments built from a knowledge base Σ. An attack relation $\mathcal{R} \subseteq \mathtt{Arg}(\Sigma) \times \mathtt{Arg}(\Sigma)$ is* valid *iff $\forall \mathcal{B} \subseteq \mathtt{Arg}(\Sigma)$, if $\mathcal{B}$ is conflict-free, then $\mathtt{Base}(\mathcal{B})$ is consistent.*

In [1], we have investigated the properties of a valid attack relation. Namely, we have shown that it should depend on the *minimal conflicts* contained in Σ, and also sensitive to them.

Definition 4 *Let* $\mathtt{Arg}(\Sigma)$ *be the set of arguments built from* Σ*, and* $\mathcal{R} \subseteq \mathtt{Arg}(\Sigma) \times \mathtt{Arg}(\Sigma)$.

- $C \subseteq \Sigma$ *is a* minimal conflict *iff i)* C *is inconsistent, and ii)* $\forall x \in C$, $C \backslash \{x\}$ *is consistent.*
- $\mathcal{R}$ *is* conflict-dependent *iff* $\forall a, b \in \mathtt{Arg}(\Sigma)$, $(a, b) \in \mathcal{R}$ *implies that there exists a minimal conflict* $C \in \mathcal{C}_\Sigma$[1] *s.t.* $C \subseteq \mathtt{Supp}(a) \cup \mathtt{Supp}(b)$.
- $\mathcal{R}$ *is* conflict-sensitive *iff* $\forall a, b \in \mathtt{Arg}(\Sigma)$ *s.t. there exists a minimal conflict* $C \in \mathcal{C}_\Sigma$ *with* $C \subseteq \mathtt{Supp}(a) \cup \mathtt{Supp}(b)$*, then either* $(a, b) \in \mathcal{R}$ *or* $(b, a) \in \mathcal{R}$.

In [1], we have shown that when the attack relation is conflict-dependent, from any consistent subset Σ, a conflict-free set of arguments is built.

Dung's abstract system is refined as follows.

Definition 5 (Argumentation system) *Given a knowledge base* Σ*, an* argumentation system *over* Σ *is a pair* $(\mathtt{Arg}(\Sigma), \mathcal{R})$ *s.t.* $\mathcal{R} \subseteq \mathtt{Arg}(\Sigma) \times \mathtt{Arg}(\Sigma)$ *is a valid attack relation.*

Among all the arguments, it is important to know which arguments to rely on for inferring conclusions from a base Σ. In [12], different acceptability semantics have been proposed. The basic idea behind these semantics is the following: for a rational agent, an argument is acceptable if he can defend this argument against all attacks on it. All the arguments acceptable for a rational agent will be gathered in a so-called *extension*. An extension must satisfy a consistency requirement (i.e. conflict-free) and must defend all its elements. Recall that a set $\mathcal{B} \subseteq \mathtt{Arg}(\Sigma)$ *defends* an argument a iff $\forall b \in \mathtt{Arg}(\Sigma)$, if $(b, a) \in \mathcal{R}$, then $\exists c \in \mathcal{B}$ such that $(c, b) \in \mathcal{R}$. The different acceptability semantics defined in [12] are recalled below. Let $\mathcal{B}$ be a conflict-free set of arguments:

- $\mathcal{B}$ is an *admissible* extension iff $\mathcal{B}$ defends all its elements.
- $\mathcal{B}$ is a *preferred* extension iff it is a maximal (for set inclusion) admissible extension.
- $\mathcal{B}$ is a *stable* extension iff it is a preferred extension that attacks any argument in $\mathtt{Arg}(\Sigma) \setminus \mathcal{B}$.

3 Relating acceptability semantics to the knowledge base

The aim of this section is to understand the underpinnings of the different acceptability semantics introduced in [12]. The idea is to analyze the results returned by each semantics in terms of subsets of the knowledge base at hand.

---

[1] Let $\mathcal{C}_\Sigma$ denote the set of all minimal conflicts of Σ.

The first result shows that when the attack relation is conflict-dependent and conflict-sensitive, then from a maximal consistent subbase of Σ, it is possible to build a *unique* maximal (wrt set inclusion) conflict-free set of arguments.

Proposition 1 *Let Σ be a knowledge base, and $(\mathcal{S}_i)_{i \in \mathcal{I}}$ be its maximal consistent subsets. If $\mathcal{R}$ is conflict-dependent and conflict-sensitive, then:*

1. *For all $i \in \mathcal{I}$, $\texttt{Arg}(\mathcal{S}_i)$ is a maximal (wrt set $\subseteq$) conflict-free subset of $\texttt{Arg}(\Sigma)$.*
2. *For all $i, j \in \mathcal{I}$, if $\texttt{Arg}(\mathcal{S}_i) = \texttt{Arg}(\mathcal{S}_j)$ then $\mathcal{S}_i = \mathcal{S}_j$.*
3. *For all $i \in \mathcal{I}$, $\mathcal{S}_i = \texttt{Base}(\texttt{Arg}(\mathcal{S}_i))$.*

Similarly, we show that each maximal conflict-free subset of $\texttt{Arg}(\Sigma)$ is built from a unique maximal consistent subbase of Σ. However, this result is *only* true when the attack relation is chosen in an "appropriate" way, i.e., when it is valid.

Proposition 2 *Let $(\texttt{Arg}(\Sigma), \mathcal{R})$ be an argumentation system over Σ, and $(\mathcal{E}_i)_{i \in \mathcal{I}}$ be the maximal (wrt set $\subseteq$) conflict-free subsets of $\texttt{Arg}(\Sigma)$. If $\mathcal{R}$ is conflict-dependent and valid, then:*

1. *For all $i \in \mathcal{I}$, $\texttt{Base}(\mathcal{E}_i)$ is a maximal (wrt set $\subseteq$) consistent subbase of Σ.*
2. *For all $i, j \in \mathcal{I}$, if $\texttt{Base}(\mathcal{E}_i) = \texttt{Base}(\mathcal{E}_j)$ then $\mathcal{E}_i = \mathcal{E}_j$.*
3. *For all $i \in \mathcal{I}$, $\mathcal{E}_i = \texttt{Arg}(\texttt{Base}(\mathcal{E}_i))$.*

Propositions 1 and 2 provide a full correspondence between maximal consistent subbases of Σ and maximal conflict-free subsets of $\texttt{Arg}(\Sigma)$.

Corollary 1. *Let $(\texttt{Arg}(\Sigma), \mathcal{R})$ be an argumentation system over Σ. If $\mathcal{R}$ is conflict-dependent and valid, then the maximal conflict-free subsets of $\texttt{Arg}(\Sigma)$ are exactly the $\texttt{Arg}(\mathcal{S})$ where $\mathcal{S}$ ranges over the maximal (wrt set inclusion) consistent subbases of Σ.*

This result is very surprising since, apart from stable semantics, all the other acceptability semantics (e.g. preferred and admissible) are not about maximal conflict-free sets of arguments, but rather subsets of them. This means that those semantics do not necessarily yield maximal subsets of the knowledge base Σ.

3.1 Stable semantics

The idea behind stable semantics is that a set of arguments is "acceptable" if it attacks any argument that is outside the set. This condition makes stable semantics very strong and the existence of stable extensions not guaranteed for every argumentation system. From the results obtained in the previous subsection on the link between maximal conflict-free subsets of $\texttt{Arg}(\Sigma)$ and maximal consistent subbases of Σ, it seems that stable semantics is not adequate. Let us illustrate this issue on the following example.

Example 1 *Let Σ be such that $\texttt{Arg}(\Sigma) = \{a, b, c\}$. Also, let the attack relation be as depicted in the figure below. This relation is assumed to be conflict-dependent and valid.*

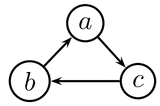

There are three maximal conflict-free sets of arguments. From Corollary 1, there is a full correspondence between the maximal (wrt set $\subseteq$) consistent subsets of Σ and the maximal (wrt set $\subseteq$) conflict-free sets of $\texttt{Arg}(\Sigma)$.

1. $\mathcal{E}_1 = \{a\}$ $\qquad \mathcal{S}_1 = \texttt{Base}(\mathcal{E}_1)$
2. $\mathcal{E}_2 = \{b\}$ $\qquad \mathcal{S}_2 = \texttt{Base}(\mathcal{E}_2)$
3. $\mathcal{E}_3 = \{c\}$ $\qquad \mathcal{S}_3 = \texttt{Base}(\mathcal{E}_3)$

Indeed, the base Σ has "stable sets of formulas" $\mathcal{S}_1, \mathcal{S}_2$ and $\mathcal{S}_3$ while these are not captured by the stable semantics since the argumentation system $(\texttt{Arg}(\Sigma), \mathcal{R})$ has no stable extension. To make the example more illustrative, let us mention that an instance of the example arises from a logic of set-theoretic difference.

The above example shows that stable semantics can miss "stable sets of formulas", i.e. maximal consistent subsets of the knowledge base at hand. However, we show that when an argumentation system $(\texttt{Arg}(\Sigma), \mathcal{R})$ has stable extensions, those extensions capture maximal consistent subsets of Σ.

Proposition 3 *If $\mathcal{R}$ is conflict-dependent and valid, then any stable extension $\mathcal{E}$ of $(\texttt{Arg}(\Sigma), \mathcal{R})$, $\texttt{Base}(\mathcal{E})$ is a maximal (wrt set $\subseteq$) consistent subset of Σ.*

The fact that stable extensions "return" maximal consistent subsets of Σ does not mean that all of them are returned. The following example illustrates this issue, and shows that Dung's system somehow picks some maximal consistent subbases of Σ.

Example 2 *Let Σ be such that $\texttt{Arg}(\Sigma) = \{a, b, c, d, e, f, g\}$ and the attack relation is as depicted in the figure below. Assume also that this relation is both conflict-dependent and valid.*

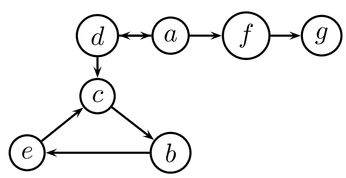

There are 5 maximal conflict-free subsets of $\texttt{Arg}(\Sigma)$:

1. $\mathcal{E}_1 = \{d, e, f\}$ $\qquad \mathcal{S}_1 = \texttt{Base}(\mathcal{E}_1)$
2. $\mathcal{E}_2 = \{b, d, f\}$ $\qquad \mathcal{S}_2 = \texttt{Base}(\mathcal{E}_2)$
3. $\mathcal{E}_3 = \{a, c, g\}$ $\qquad \mathcal{S}_3 = \texttt{Base}(\mathcal{E}_3)$
4. $\mathcal{E}_4 = \{a, e, g\}$ $\qquad \mathcal{S}_4 = \texttt{Base}(\mathcal{E}_4)$
5. $\mathcal{E}_5 = \{a, b, g\}$ $\qquad \mathcal{S}_5 = \texttt{Base}(\mathcal{E}_5)$

Since the attack relation is both conflict-dependent and valid, then each set $\mathcal{E}_i$ returns a maximal consistent subbase $\mathcal{S}_i$ of Σ. From Corollary 1, it is clear that the maximal consistent subbases of Σ are exactly the five $\texttt{Base}(\mathcal{E}_i)$*. The argumentation system* $(\texttt{Arg}(\Sigma), \mathcal{R})$ *has only one stable extension which is $\mathcal{E}_2$. Thus, Dung's approach picks only one maximal subbase, i.e., $\mathcal{S}_2$ among the five. Such a choice is clearly counter-intuitive since there is no additional information that allows us to elicit $\mathcal{S}_2$ against the others. There is a huge literature on handling inconsistency in knowledge bases. None of the approaches can make a choice between the above five subbases $\mathcal{S}_i$ if no additional information is available, like priorities between formulas of Σ.*

The above result shows that from a maximal consistent subbase of Σ, it is not always the case that a stable extension exists as its counterpart in the argumentation framework. For instance, $\mathcal{S}_1$ is a maximal consistent subbase of Σ, but its corresponding set of arguments, $\mathcal{E}_1$, is not a stable extension. This depends broadly on the attack relation that is considered in the argumentation system. The above example shows that when the attack relation is asymmetric then a full correspondence between stable extensions and maximal consistent subbases of Σ is not guaranteed.

Let us now analyze the case of a symmetric attack relation. The following result shows that each maximal consistent subbase of Σ "returns" a stable extension, provided that the attack relation is conflict-dependent and conflict-sensitive.

Proposition 4 *Let $\mathcal{S}$ be a maximal (wrt set inclusion) subset of Σ. If $\mathcal{R}$ is conflict-dependent, conflict-sensitive and symmetric, then* $\texttt{Arg}(\mathcal{S})$ *is a stable extension of* $(\texttt{Arg}(\Sigma), \mathcal{R})$*.*

From the above result, it follows that an argumentation system based on a knowledge base Σ always has stable extensions unless the knowledge base contains only inconsistent formulas.

Corollary 2. *Let* $(\texttt{Arg}(\Sigma), \mathcal{R})$ *be an argumentation system based on a knowledge base Σ such that $\mathcal{R}$ is symmetric, conflict-dependent and conflict-sensitive. If there exists $x \in \Sigma$ s.t. $\{x\}$ is consistent, then* $(\texttt{Arg}(\Sigma), \mathcal{R})$ *has at least one stable extension.*

In order to have a full correspondence between the stable extensions of $(\texttt{Arg}(\Sigma), \mathcal{R})$ and the maximal consistent subsets of Σ, the attack relation should be symmetric, conflict-dependent and valid.

Corollary 3. *Let* $(\texttt{Arg}(\Sigma), \mathcal{R})$ *be an argumentation system based on a knowledge base Σ such that $\mathcal{R}$ is symmetric, conflict-dependent and valid. Each maximal conflict-free subset of* $\texttt{Arg}(\Sigma)$ *is exactly* $\texttt{Arg}(\mathcal{S})$ *where $\mathcal{S}$ ranges over the maximal (wrt set inclusion) consistent subsets of Σ.*

However, in [1], we have shown that when the attack relation is symmetric, it is not valid in general. Namely, when the knowledge base contains n-ary (other then binary) minimal conflicts, then the rationality postulate on consistency [9] is violated due to symmetric relations, that are thereby ruled out. What happens in this case is that the argumentation system not only "returns" stable extensions that capture all the maximal consistent subsets of Σ but also other stable extensions corresponding to inconsistent subsets of Σ.

3.2 Other acceptability semantics

Preferred semantics has been introduced in order to palliate the limits of the stable semantics. Indeed, in [12], it has been shown that each argumentation system has at least one preferred extension. Preferred extensions are maximal sets of arguments that defend themselves against all attacks. The following example shows that preferred semantics may fail to capture consistent subbases of a knowledge base.

Example 3 (Example 1 cont.) *The argumentation system* $(\texttt{Arg}(\Sigma), \mathcal{R})$ *of Example 1 has one preferred extension which is the empty set. The corresponding base is thus the empty set. However, the knowledge base* Σ *has three maximal consistent subbases.*

Let us now analyze the case where the argumentation system $(\texttt{Arg}(\Sigma), \mathcal{R})$ has non-empty preferred extensions. When the attack relation is valid, each preferred extension of $(\texttt{Arg}(\Sigma), \mathcal{R})$ "returns" a consistent subbase of Σ. However, this does not mean that these subbases are maximal wrt set inclusion as shown in the following example.

Example 4 (Example 2 cont.) *The argumentation system of Example 2 has two preferred extensions:* $\mathcal{E}_2 = \{b, d, z\}$ *and* $\mathcal{E}_6 = \{a, y\}$. *The set* $\mathcal{E}_6$ *is a subset of* $\mathcal{E}_3$, $\mathcal{E}_4$ *and* $\mathcal{E}_5$ *whose bases are maximal for set inclusion. Therefore,* $\texttt{Base}(\mathcal{E}_6)$ *is not maximal, whereas* $\texttt{Base}(\mathcal{E}_2)$ *is maximal.*

As for stable extensions, preferred extensions rely on a seemingly random pick. In the above example, it is unclear why $\texttt{Base}(\mathcal{E}_6)$ is picked instead of another consistent subset of Σ. Similarly, why is $\texttt{Base}(\mathcal{E}_2)$ picked instead of $\texttt{Base}(\mathcal{S}_1)$ for instance?

It is worth mentioning that the remaining acceptability semantics (i.e. admissible, grounded and complete) also "return" somehow *randomly* some (but not necessarily all) consistent subbases of Σ *even* when the attack relation is defined carefully. Let us consider the following simple example.

Example 5 *Let* Σ *be such that* $\texttt{Arg}(\Sigma) = \{a, b, c\}$. *Also, let the attack relation be as depicted in the figure below. This relation is assumed to be conflict-dependent and valid.*

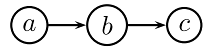

There are two maximal conflict-free sets of arguments: $\mathcal{E}_1$ *and* $\mathcal{E}_2$*.*

1. $\mathcal{E}_1 = \{a, c\}$ $\quad\quad$ $\mathcal{S}_1 = \mathtt{Base}(\mathcal{E}_1)$
2. $\mathcal{E}_2 = \{b\}$ $\quad\quad$ $\mathcal{S}_2 = \mathtt{Base}(\mathcal{E}_2)$

From Corollary 1, there is a full correspondence between the maximal (wrt set inclusion) consistent subbases of Σ *and the maximal (wrt set inclusion) conflict-free sets of* $\mathtt{Arg}(\Sigma)$*. Thus, the knowledge base* Σ *has exactly two maximal consistent subbases* $\mathcal{S}_1$ *and* $\mathcal{S}_2$*. The argumentation system* $(\mathtt{Arg}(\Sigma), \mathcal{R})$ *has one grounded extension which is* $\mathcal{E}_1$*. Thus, conclusions supported by* a *and* c *will be inferred from* Σ*. Here again it is not clear why the base* $\mathcal{S}_2$ *is discarded.*

Recently, several new acceptability semantics have been introduced in the literature. In [8], a weaker version of stable semantics has been defined. The new semantics, called *semi-stable*, is between stable and preferred semantics. It has been shown that each stable extension is a semi-stable one, and each semi-stable extension is a preferred one. When the attack relation has desirable features then, semi-stable extensions return "some" consistent bases which are not necessarily maximal wrt set inclusion. Example 1 has been handled in [8]. It has been shown that the empty set is the only semi-stable extension of the system. However, we have shown that the knowledge base has three maximal consistent subsets.

Corollary 4. *For any semi-stable extension* $\mathcal{E}$ *of* $(\mathtt{Arg}(\Sigma), \mathcal{R})$*, where* $\mathcal{R}$ *is valid,* $\mathtt{Base}(\mathcal{E})$ *is consistent.*

In [5], six new acceptability semantics have been proposed. That work was merely motivated by the fact that Dung's semantics treat differently odd-length cycles and even-length cycles. Thus, they suggested semantics that handle the two cases in a similar way. Those semantics yield extensions that are not necessarily maximal conflict-free subsets of the original set of arguments. This means that their corresponding bases are not maximal for set inclusion.

The very last acceptability semantics that has been proposed in the literature is the so-called *ideal* semantics [11]. This semantics always gives a single extension. The latter is an admissible extension that is contained in every preferred extension. To show that this semantics suffers from the problems discussed here, it is sufficient to consider Example 5. Indeed, in that example, there is one ideal extension which is $\{a, c\}$. We have shown that this result is ad hoc since it allows to infer formulas from a seemingly randomly chosen consistent subset of the knowledge base Σ.

4 Particular argumentation systems

In this section we will study three argumentation systems. These systems use the three attack relations that are commonly used in the literature: rebut, strong rebut and assumption attack (called here "undercut"). Although these relations have originally been defined in a propositional logic context, we will generalize them to any logic in the sense of Tarski.

The rebut relation captures the idea that two arguments with contradictory conclusions are conflicting. This relation is thus defined as follows:

Definition 6 (Rebut) *Let* $a, b \in \texttt{Arg}(\Sigma)$*. The argument* a rebuts b *iff the set* $\{\texttt{C}onc(a), \texttt{C}onc(b)\}$ *is inconsistent.*

It is easy to show that this relation is conflict-dependent.

Property 1. The relation "Rebut" is conflict-dependent.

It can be checked that this relation is *not conflict-sensitive*. Let us, for instance, consider the following two arguments built from a propositional base: $(\{x \wedge y\}, x)$ and $(\{\neg y \wedge z\}, z)$. It is clear that the set $\{x \wedge y, \neg y \wedge z\}$ is a minimal conflict, but none of the two arguments is rebutting the other. Note also that the rebut relation violates most of the properties studied in [1]. In particular, it is not homogeneous. Since the rebut relation is symmetric, it is not valid in the case that the knowledge base contains n-ary (other than binary) minimal conflicts. Worse yet, this relation is not valid even in the binary case. This can be seen from the following conflict-free set of arguments, $\{(\{x, x \to y\}, y), (\{\neg x, \neg x \to z\}, z)\}$, whose corresponding base is inconsistent. This means that the argumentation system $(\texttt{Arg}(\Sigma), \text{Rebut})$ violates consistency. Consequently, this system is not recommended for inferring conclusions from a knowledge base Σ.

Property 2. If Σ is such that $\exists C = \{x_1, \ldots, x_n\} \subseteq \mathcal{C}_\Sigma$ where either $n > 2$ or $n = 2$ and $\exists x'_1 \in \mathsf{CN}(\{x_1\}, \exists x'_2 \in \mathsf{CN}(\{x_2\}$ s.t. $\{x'_1, x'_2\}$ is inconsistent then $(\texttt{Arg}(\Sigma), \text{R}ebut)$ violates extension consistency.

Let us now consider the second symmetric relation, called "Strong Rebut". This relation is defined as follows:

Definition 7 (Strong Rebut) *Let* $a, b \in \texttt{Arg}(\Sigma)$*. The argument* a strongly rebuts b *iff* $\exists x \in \mathsf{CN}(\texttt{S}upp(a))$ *and* $\exists y \in \mathsf{CN}(\texttt{S}upp(b))$ *such that the set* $\{x, y\}$ *is inconsistent.*

The following result shows that "Strong Rebut" satisfies interesting properties.

Property 3. "Strong Rebut" is conflict-dependent and conflict-sensitive.

Unfortunately, since this relation is symmetric, then it is not valid in the general case, i.e. when n-ary (other than binary) minimal conflicts occur in the knowledge base Σ.

Property 4. Let Σ be such that $\exists C \in \mathcal{C}_\Sigma$ such that $|C| > 2$. An argumentation system $(\texttt{Arg}(\Sigma), \text{Strong Rebut})$ violates consistency.

In the particular case where all the minimal conflicts of the knowledge base base Σ are binary, we have shown in [1] that conflict-sensitive attack relations are valid. Since "Strong Rebut" is conflict-sensitive, then it is a valid relation.

Property 5. Let Σ be such that all the minimal conflicts in $\mathcal{C}_\Sigma$ are binary. Any argumentation system $(\texttt{Arg}(\Sigma), \text{Strong Rebut})$ satisfies consistency.

Finally, we show that when all the minimal conflicts of the knowledge base Σ are binary, then there is a full correspondence between the maximal consistent subsets of Σ and the stable extensions of $(\texttt{Arg}(\Sigma), \text{Strong Rebut})$. Remember that, since "Strong Rebut" is symmetric, then all the semantics coincide as they all give the maximal conflict-free subsets of $\texttt{Arg}(\Sigma)$.

Property 6. Let $(\texttt{Arg}(\Sigma), \text{Strong Rebut})$ be an argumentation system over a knowledge base Σ whose minimal conflicts are all binary. For any stable extension $\mathcal{E}$ of $(\texttt{Arg}(\Sigma), \text{Strong Rebut})$, $\texttt{Base}(\mathcal{E})$ is a maximal consistent subset of Σ. For any maximal consistent subset $\mathcal{S}$ of Σ, $\texttt{Arg}(\mathcal{S})$ is a stable extension of $(\texttt{Arg}(\Sigma), \text{Strong Rebut})$.

Let us now investigate the results of an argumentation system that is built upon the attack relation "Undercut". Recall that the idea behind this relation is that an argument may undermine a premise of another argument. Formally:

Definition 8 (Undercut) *Let $a, b \in \texttt{Arg}(\Sigma)$. The argument a* undercuts *b iff $\exists x \in \texttt{S}upp(b)$ such that the set $\{\texttt{C}onc(a), x\}$ is inconsistent.*

The next property shows the main feature of this attack relation.

Property 7. The relation "Undercut" is conflict-dependent.

It can be checked that this relation is unfortunately not conflict-sensitive. For instance, neither the argument $(\{x \wedge y\}, x)$ undercuts the argument $(\{\neg y \wedge z\}, z)$ nor $(\{\neg y \wedge z\}, z)$ undercuts $(\{x \wedge y\}, x)$. However, the set $\{x \wedge y, \neg y \wedge z\}$ is clearly a minimal conflict. Due to this limitation, undercut is *not valid*.

Property 8. The relation "Undercut" is not valid.

Somewhat surprisingly, the argumentation system $(\texttt{Arg}(\Sigma)$, Undercut) satisfies extension consistency. Indeed, the base of any admissible extension of this argumentation system is consistent. The reason is that the arguments that do not behave correctly w.r.t. Undercut, are defended by other arguments.

Property 9. If $\mathcal{E}$ is an admissible extension of an argumentation system $(\texttt{Arg}(\Sigma)$, Undercut), then $\texttt{Base}(\mathcal{E})$ is consistent.

Unfortunately, the fact that the extensions (under a given acceptability semantics) of $(\texttt{Arg}(\Sigma)$, Undercut) satisfy consistency does not mean that "Undercut" is a good attack relation. This relation works only when "all" the arguments that can be built from a knowledge base are extracted and taken into account. In many applications, like dialogue between two or more agents, this is inadequate. Indeed, in a dialogue, agents seldom utter all the possible arguments.

In [10], it has been shown that when the knowledge base Σ is encoded in *propositional logic*, that is, $(\mathcal{L}, \mathsf{CN})$ is propositional logic, there is a correspondence between the stable extensions of the argumentation system $(\mathtt{Arg}(\Sigma), \text{Undercut})$ based on Σ and the maximal consistent subsets of Σ. The following result shows that this property holds for any logic in the sense of Tarski.

Proposition 5 *Let* $(\mathtt{Arg}(\Sigma), \textit{Undercut})$ *be an argumentation system over a knowledge base* Σ*.*

1. *For all stable extension* $\mathcal{E}$ *of* $(\mathtt{Arg}(\Sigma), \textit{Undercut})$*,* $\mathtt{Base}(\mathcal{E})$ *is a maximal (wrt set inclusion) consistent subset of* Σ*.*
2. *For all stable extension* $\mathcal{E}$ *of* $(\mathtt{Arg}(\Sigma), \textit{Undercut})$*,* $\mathcal{E} = \mathtt{Arg}(\mathtt{Base}(\mathcal{E}))$*.*

The following result shows that each maximal (wrt set inclusion) consistent subset of Σ "returns" exactly one stable extension of $(\mathtt{Arg}(\Sigma), \text{Undercut})$.

Proposition 6 *Let* $(\mathtt{Arg}(\Sigma), \textit{Undercut})$ *be an argumentation system over a knowledge base* Σ*.*

1. *For all maximal (wrt set inclusion) consistent subset* $\mathcal{S}$ *of* Σ*,* $\mathtt{Arg}(\mathcal{S})$ *is a stable extension of* $(\mathtt{Arg}(\Sigma), \textit{Undercut})$*.*
2. *For all maximal (wrt set inclusion) consistent subset* $\mathcal{S}$ *of* Σ*,* $\mathcal{S} = \mathtt{Base}(\mathtt{Arg}(\mathcal{S}))$*.*

From the above results, a full correspondence between the stable extensions of $(\mathtt{Arg}(\Sigma), \text{Undercut})$ and the maximal consistent subsets of Σ is established.

Corollary 5. *Let* $(\mathtt{Arg}(\Sigma), \textit{Undercut})$ *be an argumentation system over a knowledge base* Σ*. The stable extensions of* $(\mathtt{Arg}(\Sigma), \textit{Undercut})$ *are exactly* $\mathtt{Arg}(\mathcal{S})$ *where* $\mathcal{S}$ *ranges over the maximal (wrt set inclusion) consistent subsets of* Σ*.*

Since each maximal consistent subset of Σ "returns" a stable extension of the argumentation system $(\mathtt{Arg}(\Sigma), \text{Undercut})$, then the latter always has stable extensions unless the knowledge base Σ contains only inconsistent formulas.

Corollary 6. *Let* $(\mathtt{Arg}(\Sigma), \textit{Undercut})$ *be an argumentation system over a knowledge base* Σ*. If there exists* $x \in \Sigma$ *s.t.* $\{x\}$ *is consistent, then* $(\mathtt{Arg}(\Sigma), \textit{Undercut})$ *has at least one stable extension.*

The results presented in this section are disappointing. They show that the three classes of argumentation systems that are that have been given extensive attention in the literature suffer from collapse. The argumentation system $(\mathtt{Arg}(\Sigma), \text{Rebut})$ should not be used since it violates the postulate on consistency, while the system $(\mathtt{Arg}(\Sigma), \text{Strong Rebut})$ should only be used in the particular case that all minimal conflicts of Σ are binary. The third argumentation system $(\mathtt{Arg}(\Sigma), \text{Undercut})$ is misleading inasmuch as it satisfies extension consistency while the "Undercut" relation is not valid. However, its stable extensions "return" all and only maximal consistent subbases of the knowledge base at hand. Thus, it can be used specifically under the stable semantics. Anyway, we should keep in mind that the "Undercut" relation satisfies extension consistency at the cost of needing to take into account arguments that have different conclusions but the same support: so much redundancy! The results show that the system based on the "Undercut" relation guarantees extension consistency (at a prohibitive cost) but "Undercut" does not enjoy nice properties.

5 Conclusion

The paper has analyzed the different acceptability semantics introduced in [12] in terms of the subbases that are returned by each extension. The results of the analysis show that there is a complete correspondence between maximal consistent subbases of a KB and maximal conflict-free sets of arguments. Regarding stable extensions, we have shown that, when the attack relation is chosen in an appropriate way, then they return maximal consistent subbases of KB. However, they don't return all of them. A choice is made in an unclear way. The case of preferred extensions is worse since they compute consistent subbases but not necessarily maximal ones. When the attack relation is considered symmetric, inconsistent subbases are returned by conflict-free sets of arguments when ternary or more minimal conflicts are available in a KB. This is due to the binary character of the attack relation.

The paper studies three argumentation systems that rely on well-known attack relations. The results show that the system based on the "Undercut" relation guarantees extension consistency (at a prohibitive cost) but "Undercut" does not enjoy nice properties.

References

1. L. Amgoud and P. Besnard. Bridging the gap between abstract argumentation systems and logic. In *Proceedings of the 3rd International Conference on Scalable Uncertainty*, pages 12–27, 2009.
2. L. Amgoud and C. Cayrol. Inferring from inconsistency in preference-based argumentation frameworks. *International Journal of Automated Reasoning*, 29 (2):125–169, 2002.
3. L. Amgoud, S. Parsons, and N. Maudet. Arguments, dialogue, and negotiation. In *Proceedings of the 14th European Conference on Artificial Intelligence (ECAI'00)*, pages 338–342. IOS Press, 2000.
4. L. Amgoud and H. Prade. Using arguments for making and explaining decisions. *Artificial Intelligence Journal*, 173:413–436, 2009.
5. P. Baroni, M. Giacomin, and G. Guida. Scc-recursiveness: a general schema for argumentation semantics. *Artificial Intelligence Journal*, 168:162–210, 2005.
6. P. Besnard and A. Hunter. *Elements of Argumentation.* MIT Press, 2008.
7. B. Bonet and H. Geffner. Arguing for decisions: A qualitative model of decision making. In *Proceedings of international Conference on Uncertainty in Artificial Intelligence, (UAI'96)*, pages 98–105, 1996.

8. M. Caminada. Semi-stable semantics. In *Proceedings of the 1st International Conference on Computational Models of Argument*, pages 121–130, 2006.
9. M. Caminada and L. Amgoud. On the evaluation of argumentation formalisms. *Artificial Intelligence Journal*, 171 (5-6):286–310, 2007.
10. C. Cayrol. On the relation between argumentation and non-monotonic coherence-based entailment. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pages 1443–1448, 1995.
11. P. Dung, P. Mancarella, and F. Toni. Computing ideal skeptical argumentation. *Artificial Intelligence Journal*, 171:642–674, 2007.
12. P. M. Dung. On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n-person games. *Artificial Intelligence Journal*, 77:321–357, 1995.
13. A. Kakas and P. Moraitis. Argumentation based decision making for autonomous agents. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multi-Agents systems (AAMAS'03)*, pages 883–890, 2003.
14. S. Kraus, K. Sycara, and A. Evenchik. Reaching agreements through argumentation: a logical model and implementation. *Journal of Artificial Intelligence*, 104:1–69, 1998.
15. H. Prakken. Coherence and flexibility in dialogue games for argumentation. *Journal of Logic and Computation*, 15:1009–1040, 2005.
16. G. Simari and R. Loui. A mathematical treatment of defeasible reasoning and its implementation. *AIJ*, 53:125–157, 1992.
17. A. Tarski. *On Some Fundamental Concepts of Metamathematics*. Logic, Semantics, Metamathematic. Edited and translated by J. H. Woodger, Oxford Uni. Press, 1956.