On sceptical vs credulous acceptance for abstract argument systems

Sylvie Doutre and **Jérôme Mengin**
IRIT – Université Paul Sabatier
118, route de Narbonne
31062 Toulouse Cedex 4 – France
{doutre,mengin}@irit.fr

Abstract

At a high level of abstraction, many systems of argumentation can be represented by a set of abstract arguments, and a binary relation between these abstract arguments describing how they contradict each other. Acceptable sets of arguments, called *extensions*, can be defined as sets of arguments that do not contradict one another, and attack all their attackers. We are interested in this paper in answering the question: is a given argument in all extensions of an argumentation system? In fact, what is likely to be useful in AI systems is not a simple yes/no answer, but some kind of well-argued answer, called a proof: if an argument is in every extension, why is it so? We describe a close connection between this problem and proofs that some meta-argument is in at least one extension of a *meta*-argumentation system, describing relationships between sets of arguments of the initial system.

Introduction

At a high level of abstraction, many systems of argumentation can be represented by a set of abstract arguments (whose internal structure is not necessarily known), and a binary relation between abstract arguments describing how arguments contradict each other. In particular, several problems related to defeasible reasoning or logic programming can be studied in such an abstract argumentation framework (see e.g. (Bondarenko *et al.* 1997)).

Given that some arguments contradict others, and considering that in general contradiction is not desirable, one of the most important questions concerning abstract argumentation systems is to define which arguments are acceptable. The most widespread definition of acceptability associated with non-monotonic logics or logic programs considers that the acceptable sets are the *stable extensions*, which correspond to kernels of the contradiction graph (Dimopoulos & Torres 1996; Dimopoulos, Magirou, & Papadimitriou 1997; Berge 1973). However, this stable *semantics* has some features that can be undesirable in some contexts: notably, it can happen that no set of arguments is stable. Under another semantics introduced in (Dung 1995), acceptable sets of arguments are called *preferred extensions*. The preferred semantics captures well some of the intuitions behind the stable semantics, and avoids several of its drawbacks.

The next questions to address are then of the form: is a given argument in some/all extensions of an argumentation system? In fact, what is likely to be useful in AI systems is not a simple yes/no answer, but some kind of well-argued answer: if an argument is acceptable, why is it so? In the world of mathematics, a well-argued answer is called a proof. In classical, monotonic logic, a proof is represented by a sequence of formulas, such that the sequence describes stepwise progress towards the conclusion.

In the case of abstract argumentation systems, proofs of acceptability usually have the form of a game between two players: one tries to establish the acceptability of an argument, the other tries to establish the opposite by putting forwards arguments that contradict those of the former; the player that tries to establish the validity of an argument can defeat its opponent by providing arguments that contradict its opponent's ones. Several proof theories of this type have been proposed for acceptability problems under various semantics (Kakas & Toni 1999; Prakken & Sartor 1997; Amgoud & Cayrol 2002; Vreeswijk & Prakken 2000; Cayrol, Doutre, & Mengin 2003). They usually consider the credulous acceptance problem, where one tries to establish that an argument is in at least one extension of the theory; or some particular cases of the more difficult sceptical acceptance problem, where one tries to establish that an argument is in every extension of the theory.

We address here the problem of proving that an argument is in every preferred extension of some argumentation system. We formally establish a close connection between this problem and proofs of credulous acceptance in a *meta*-argumentation system describing relationships between sets of arguments of the initial system.

The paper is built as follows: Dung's abstract argumentation framework is briefly presented in the next section. The meta-argumentation system is introduced in the third Section. The fourth Section describes a general proof theory for sceptical acceptance using this meta-argumentation system. We finish the paper with some concluding remarks.

Preferred extensions of abstract argumentation systems

This section is a short presentation of Dung's abstract argumentation framework and of the preferred semantics. More details, in particular other semantics, can be found in e.g. (Dung 1995; Bondarenko *et al.* 1997; Doutre 2002).

Definition 1 *(Dung 1995) An* argument system *is a pair* (A, R) *where* A *is a set whose elements are called* arguments *and* R *is a binary relation over* A *(*$R \subseteq A \times A$*). Given two arguments* a *and* b*,* $(a, b) \in R$ *or equivalently* aRb *means that* a attacks b *(*a *is said to be an* attacker *of* b*). Moreover, an argument* a R*-attacks a set* S *of arguments if there exists* $b \in S$ *such that* aRb*. A set* S *of arguments* R*-attacks an argument* $a \in A$ *if there exists* $b \in S$ *such that* bRa*. Finally, a set* S *of arguments* R-attacks *a set* S' *of arguments if there exists* $a \in S$ *such that* a R*-attacks* S'*.*

In the following definitions and notations, we assume that an argument system (A, R) is given.

An argument system can be simply represented as a directed graph whose vertices are the arguments and edges correspond to the elements of R.

Notation 1 *For every set* $S \subseteq A$*,* $R^+(S) = \{a \in A \mid S\ R\text{-attacks } a\}$*,* $R^-(S) = \{a \in A \mid a\ R\text{-attacks } S\}$ *and* $R^\pm(S) = R^+(S) \cup R^-(S)$*. Moreover,* $\mathrm{Refl}(A, R) = \{a \in A \mid aRa\}$ *is the set of arguments that attack themselves.*

Definition 2 *A set* $S \subseteq A$ *is* conflict-free *if and only if there are no arguments* a *and* b *in* S *such that* a *attacks* b*. An argument* $a \in A$ *is* defended *by a set* $S \subseteq A$ *(or* S defends a*) iff for each argument* b *in* A *that attacks* a *there exists an argument in* S *that attacks* b*. A set* $S \subseteq A$ *is* admissible *if and only if* S *is conflict-free and* S *defends all its elements. It is a* preferred extension *if and only if* S *is maximal for set-inclusion among the admissible sets.*

Note that the preferred semantics coincide with several important other semantics when the graph of the argumentation system has no odd cycle. And that most existing semantics coincide when the graph has no cycle at all.

Notation 2 *Given* (A, R)*,* $\mathrm{Adm}(A, R)$ *denotes the set of the admissible sets of* (A, R)*.*

Dung exhibits interesting properties of preferred extensions: every admissible set is contained in a preferred extension; every argument system possesses at least one preferred extension.

In the rest of the paper, when no particular semantics is mentionned, it is assumed that extensions refer to the preferred extensions.

Example 1 *Consider the following system:*

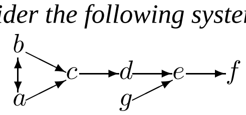

The preferred extensions are $\{a, d, g, f\}$ *and* $\{b, d, g, f\}$*:* d*,* g *and* f *are in all of them; whereas* a *and* b *are in at least one preferred extensions each, but not in every preferred extension.*

The crucial problem, given any definition of acceptability, is to be able to decide which arguments are acceptable. In the case of the preferred semantics, two problems can be formally defined as follows:

Definition 3 *Let* $a \in A$ *and* $S \subseteq A$*.* a *(resp.* S*) is* credulously accepted (w.r.t. (A, R)) *under the preferred semantics iff* a *(resp.* S*) is contained in at least one preferred extension of* (A, R)*.* a *(resp.* S*) is* sceptically accepted (w.r.t. (A, R)) *under the preferred semantics iff* a *(resp.* S*) is contained in every preferred extension of* (A, R)*.*

Note that an argument a is credulously accepted if and only if $\{a\}$ is; whereas a set of arguments S is sceptically accepted if and only if every $x \in S$ is sceptically accepted. In the next section, we relate the sceptical acceptance of an argument to the credulous acceptance of a set of arguments.

Sceptical acceptance as credulous meta-acceptance

The credulous acceptance problem has been well-studied; several proof theories and algorithms exist to answer questions like: is a given argument in at least one extension of a given argumentation system? (See e.g. (Vreeswijk & Prakken 2000; Cayrol, Doutre, & Mengin 2003).) We are interested here in answering questions like: is a given argument in every extension of a given argumentation system?

The last question is easily (if not efficiently) answered if we can enumerate all the extensions of the system: we consider a first extension E_1 and test if $x \in E_1$. If it is, this suggests that x may indeed be in every extension (as opposed to the case where $x \notin E_1$). We then consider a second extension E_2: it may happen that $x \notin E_2$, so the existence of E_2 *a priori* casts a doubt over the fact that x is in every extension. However, if it turns out that $x \in E_2$, this reinforces the possibility that x may be in every extension. Continuing the process, each extension E starts, with its sole existence, by being an argument suggesting that x may not be in every extension, to become, if it turns out $x \in E$, an argument reinforcing the possibility that x is in every extension. Of course, enumerating all the extensions will generally not be efficient. We study in the remainder of this section how we can refine this approach, by enumerating smaller sets that can be interpreted as "meta"-arguments for or against the possibility that x is in every extension.

Since every argumentation system has at least one (preferred) extension, an argument x must be in at least one extension in order to be in all of them, so x must be in at least one admissible set. Now, suppose we have found one admissible set P that contains x; so we know that x is in at least one extension $E \supseteq P$. What could prevent x from being in every extension? If there is an extension E' such that $x \notin E'$, then $P \not\subseteq E'$, so there must be a conflict between P and E' (otherwise, since P and E' defend themselves, $P \cup E'$ would be admissible, which is not possible since E' is maximally admissible and $P \not\subseteq E'$). Thus if x is not in every extension, there must be some admissible set P' that attacks P and such that P' is not in any extension that contains x (take for instance $P' = E'$). In a sense, P can be seen as a "meta"-argument suggesting that x may well be in every extension of the system; whereas P' can be seen as a counter-argument: it suggests that, since there is an admissible set that contradicts P, there may be some maximal admissible set of arguments that does not contain x. This "meta" counter-argument is in turn contradicted if there is some admissible set of arguments P'' that contains both P' and x.

This approach can be formalized by defining a relation R_x on the admissible subsets of some argument system (A, R): let $X, Y \in \mathrm{Adm}(A, R)$, then XR_xY if:

1. $x \in Y \setminus X$ and X R-attacks Y; or
2. $x \in X \setminus Y$ and $X \supseteq Y$.

In case 1., Y suggests that x may be in every extension: it is at least in all the extensions that contain Y; but X suggests that there may in fact be some extensions that do not contain x: those that contain X cannot contain Y because X R-attacks Y.

In case 2., Y suggests that x may not be in every extension, since it is admissible and does not contain x; but X shows that Y can be extended to an admissible set that does contain x.

We will say that X R_x-attacks Y if XR_xY. Note that this is not equivalent to X R-attacks Y, which means that there is $(x, y) \in X \times Y$ such that xRy.

Example 2 *Consider the element d of the argumentation system depicted on the left of Fig. 1, and the relation R_d depicted on the right of the same figure. The "base" argumentation system has four non-empty admissible sets: $\{a\}$, $\{a, d\}$, $\{b\}$ and $\{b, d\}$. Since $d \in \{a, d\}$, d may be in every extensions. This may be contradicted by the fact that $\{b\}$ R_d-attacks $\{a, d\}$: this suggests that there may be an extension that does not contain d. However, $\{b, d\}$ R_d-attacks $\{b\}$: it is a larger admissible set that contains $\{b\}$ and contains d. In fact, the set $\{\{a, d\}, \{b, d\}\}$ is R_d-admissible, and d is in every extension.*

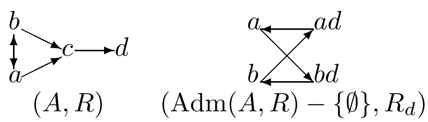

Figure 1: The argumentation system of Example 2.

Example 3 *Consider the following system:*

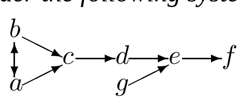

The argument d is in two admissible sets: $\{a, d\}$ and $\{b, d\}$. In fact, like on the previous example, d is in every extension. However, the set $\{\{a, d\}, \{b, d\}\}$ is not R_d-admissible, since it is R_d-attacked by e.g. $\{b, g\}$. A R_d-admissible set is $\{\{a, d, f, g\}, \{b, d, f, g\}\}$. This is somewhat surprising: the status of d only depends on the arguments a, b, c and d. However, when looking for an R_d-admissible set, we have to consider f and g as well. This is because the R_d-admissible set has to defend itself against $\{b, g\}$, although it is b, not g, that R-attacks $\{a, d\}$.

The previous example suggests that it may be sufficient to restrict R_x to some R-admissible sets only. More precisely, we define $\mathcal{A}_x = \mathcal{A}_x^{\mathrm{PRO}} \cup \mathcal{A}_x^{\mathrm{OPP}}$, where:

$\mathcal{A}_x^{\mathrm{PRO}}$ is the set of the R-admissible parts of A that contain x;

$\mathcal{A}_x^{\mathrm{OPP}}$ is the set of the R-admissible parts X of A that do not contain x and are of the form $X = \cup_{Y \in \mathcal{Y}} Y$, where the $Y \in \mathcal{Y}$ are parts of A minimal such that Y is R-admissible and R-attacks some element of $\mathcal{A}_x^{\mathrm{PRO}}$.

We are now able to express problems of sceptical acceptance in terms of admissibility, or credulous acceptance, in the meta-graph:

Proposition 1 *An argument x of an argumentation framework (A, R) is in every preferred extension of (A, R) if and only if there exist $P \in \mathrm{Adm}(A, R)$ and $\mathcal{P} \in \mathrm{Adm}(\mathcal{A}_x, R_x)$ such that $x \in P$ and $P \in \mathcal{P}$.*

Proof 1 *Suppose first that x is in every preferred extension of (A, R): (A, R) has at least one preferred extension, which must contain x; let P be this preferred extension. The next proposition shows that there is $\mathcal{P} \in \mathrm{Adm}(\mathcal{A}_x, R_x)$ such that $P \in \mathcal{P}$.*

For the converse, suppose that P and $\mathcal{P}$ exist, and let E be a preferred extension of (A, R). If $P \subseteq E$, then $x \in E$. If $P \not\subseteq E$, then $E \cup P$ is not admissible, thus, since E and P are both admissible, there is a conflict between them, and E R-attacks P. Let $\mathcal{Y}$ be the set of the minimal parts of E that are admissible and attack some elements of $\mathcal{A}_x^{\mathrm{PRO}}$. Then $X = \cup_{Y \in \mathcal{Y}} Y$ is such that $x \notin X$, $X \in \mathcal{A}_x^{\mathrm{OPP}}$ (X is admissible since $X \subseteq E$ and E is admissible) and X R-attacks P, thus XR_xP. Since $\mathcal{P}$ is admissible in $(\mathcal{A}_x, R_x)$, there is $P' \in \mathcal{A}_x^{\mathrm{PRO}}$ such that $P'R_xX$, that is, $P' \supseteq X$. Suppose that $P' \not\subseteq E$, then E R-attacks P', but then there is some $Y \in \mathcal{Y}$ such that Y R-attacks P', thus $Y \subseteq X \subseteq P'$: this contradicts the fact that P' is R-admissible. Thus $P' \subseteq E$ and, since $x \in P'$, $x \in E$.

This result guarantees that if we can find an R-admissible part $P \subseteq A$ that contains x, and an R_x-admissible part of $\mathcal{A}_x$ that contains P, then x is in every extension. However, this result alone does not guarantee the completeness of the approach: if we find P, but then cannot find $\mathcal{P}$, is it that x is not in every preferred extension of (A, R), or could be that we just picked the wrong P? The following result shows that the method is complete:

Proposition 2 *If an argument x of an argumentation framework (A, R) is in every preferred extension of (A, R), then for every $P \in \mathrm{Adm}(A, R)$ such that $x \in P$ there exists $\mathcal{P} \in \mathrm{Adm}(\mathcal{A}_x, R_x)$ such that $P \in \mathcal{P}$.*

Proof 2 *Let $P \in \mathrm{Adm}(A, R)$ such that $x \in P$, and let $\mathcal{P}$ be the set containing P and the preferred extensions of (A, R): $\mathcal{P}$ is without R_x-conflict since all its elements contain x. Furthermore, suppose that $P' \in \mathcal{A}_x$ is such that P' R_x-attacks $\mathcal{P}$: P' is contained in a preferred extension F of (A, R); and there is $E \in \mathcal{P}$ such that $P'R_xE$. But then, $x \notin P'$, thus FR_xP'. This proves that $\mathcal{P}$ is an admissible part of $(\mathcal{A}_x, R_x)$.*

We close this section with an example that shows why, when defining $\mathcal{A}_x^{\mathrm{OPP}}$, we cannot simply consider minimal admissible sets Y that attack elements of $\mathcal{A}_x^{\mathrm{PRO}}$, but need to consider unions of such Y's.

Example 4 *Consider the following system:*

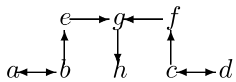

It has four preferred extensions: $\{a, c, e, h\}$, $\{a, d, e, f, h\}$, $\{b, c, g\}$ and $\{b, d, f, h\}$; so h is not in the intersection of the extensions. However, let $\mathcal{P} = \{\{a, c, e, h\}, \{a, d, e, g, h\}, \{b, d, f, h\}\}$, every element of $\mathcal{P}$ is admissible and contains h; furthermore, the minimal admissible sets of arguments that attack at least one element of $\mathcal{P}$ are $\{a\}$, $\{b\}$, $\{c\}$ and $\{d\}$: each of them is contained in an element of $\mathcal{P}$. There is one admissible set of arguments that attacks some element of $\mathcal{P}$ and which is not contained in any element of $\mathcal{P}$: it is the union $\{b, c\}$ of $\{b\}$ and $\{c\}$.

A proof-theory for sceptical acceptance

In this section, we describe a proof theory for the problem of the sceptical acceptance of an argument, using our characterization of this problem in terms of credulous meta-acceptance.

In classical, monotonic logic, a proof is represented by a sequence of formulas, such that the sequence describes stepwise progress towards the conclusion. In the case of abstract argumentation systems, proofs of acceptability usually have the form of a game between two players, one called PRO, the other one called OPP: PRO tries to establish the acceptability of an argument, while OPP tries to establish the opposite by putting forwards arguments that contradict those of PRO; PRO can defeat its opponent by providing arguments that contradict its opponent's ones. Several proof theories of this type have been proposed for acceptance problems under various semantics, (Kakas & Toni 1999; Prakken & Sartor 1997; Amgoud & Cayrol 2002; Vreeswijk & Prakken 2000; Cayrol, Doutre, & Mengin 2003).

Since we have characterized the sceptical acceptance problem as a credulous acceptance problem in a meta-graph, any proof theory designed for the credulous acceptance problem can be used to solve the sceptical acceptance problem. We illustrate this below with a proof theory proposed in (Cayrol, Doutre, & Mengin 2003).

Argument games have been formalized in (Jakobovits & Vermeir 1999) using sequences of *moves*, called *dialogues*. We recall below a definition of a proof theory for the credulous acceptance proposed in (Cayrol, Doutre, & Mengin 2003).

Definition 4 *Let (A, R) be an argumentation system. A* move *in A is a pair $[P, X]$ where $P \in \{\mathrm{PRO}, \mathrm{OPP}\}$ and $X \in A$. A* dialogue *d in (A, R) for a finite set of arguments $S = \{a_1, a_2, \ldots, a_n\} \subseteq A$ is a countable sequence of moves of the form*

$$[\mathrm{PRO}, a_1] \ldots [\mathrm{PRO}, a_n][\mathrm{OPP}, b_1][\mathrm{PRO}, b_2][\mathrm{OPP}, b_3] \ldots \\ \ldots [\mathrm{OPP}, b_{2i-1}][\mathrm{PRO}, b_{2i}][\mathrm{OPP}, b_{2i+1}] \ldots$$

such that:

1. *the first n moves are played by PRO to put forward the elements of S;*
2. *the subsequent moves are played alternatively by OPP and PRO;*
3. *the ith argument put forward by OPP is $b_{2i-1} \in R^-(P_i) \setminus R^+(P_i)$, where $P_i = S \cup \{b_2, b_4, \ldots, b_{2i-2}\}$ is the set of arguments put forward by PRO so far;*
4. *the $n + i$th argument put forward by PRO is $b_{2i} \in R^-(b_{2i-1}) \setminus (P_i \cup R^{\pm}(P_i) \cup \mathrm{Refl}(A, R))$.*

A finite dialogue is won by PRO *if OPP cannot respond to PRO's last move in accordance with rule 3. above.*

Rule 3. means that OPP can attack any argument put forward by PRO, with any argument not attacked by arguments already put forward by PRO.

Rule 4. means that PRO must defend itself against OPP's last attack, with an argument that it has not already put forward, and that is not in conflict with the arguments it has already put forward.

The following proposition ensures the soundness and completeness of the above proof theory for set-credulous acceptance.

Proposition 3 *(Cayrol, Doutre, & Mengin 2003)*[1] *Let (A, R) be an argument system. If d is a dialogue for S won by PRO, then $\mathrm{PRO}(d)$ is an admissible set containing S. If S is included in a preferred extension of (A, R) then there exists a dialogue for S won by PRO.*

Notice that another winning criterion defined by (Jakobovits & Vermeir 1999) (the *winning strategy*) could be used in order to design proofs in which one can see precisely how any argument of a proof (not only an argument of the set S) is defended against its attackers.

Let us now describe how this type of dialogue can be used as a proof theory for the sceptical acceptance problem. Suppose that we want to prove that some argument x of an argumentation system (A, R) is in every extension of (A, R). According to the results of the preceding section, all we need to do is find an admissible set P that contains x, and then find a dialogue for $\{P\}$ won by PRO with respect to the argumentation system $(\mathcal{A}_x, R_x)$.

In order to find the initial admissible set P that contains x, Prop. 3 says that we can look for a dialogue d for x won by PRO w.r.t. (A, R): we can then take $P = \mathrm{PRO}(d)$. In order to establish that x is in every extension of the theory, we then start a dialogue with the move $[\mathcal{PRO}, P]$, where $\mathcal{PRO}$ denotes the player that tries to establish the acceptability of P in the meta-graph. In fact, a more detailed dialogue can start with the move $[\mathcal{PRO}, d]$, showing not only P but the entire dialogue that established the admissibility of P.

In order to continue the meta-dialogue, we need a move of the form $[\mathcal{OPP}, d_1]$, where $\mathcal{OPP}$ denotes the player who tries to establish that P is not credulously accepted in the meta-graph, and where d_1 must be a dialogue in (A, R) for an argument that R-attacks P.

[1]The result proved in (Cayrol, Doutre, & Mengin 2003) concerns the credulous acceptance of a single argument. The result proposed here is a straightforward extension of it to the problem of the credulous acceptance of a set of arguments.

$\mathcal{PRO}$ must then put forward a dialogue for $\mathrm{PRO}(d_1) \cup \{x\}$ won by PRO, thereby showing that the admissible set found by $\mathcal{OPP}$ in the preceding move can be "returned" in favor of $\mathcal{PRO}$.

This type of meta-dialogue is best illustrated on an example.

Example 5 *Consider the following system:*

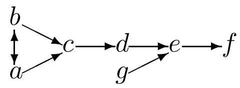

A meta-dialogue proving that d is in every extension of the theory is depicted on Fig. 2 next page (note that the moves of the dialogues in (A, R) are in columns, whereas the moves of the meta-dialogue in $(\mathcal{A}_x, R_x)$ are in line).

We close this section with a remark about the possibilities for $\mathcal{PRO}$ to defend itself. Rule 3. of the definition of a dialogue says that the $1+i$th argument b_{2i} put forward by $\mathcal{PRO}$ must be in $\in R_x^-(b_{2i-1}) \setminus (\mathcal{P}_i \cup R^{\pm}(\mathcal{P}_i) \cup \mathrm{Refl}(\mathcal{A}_x, R_x))$, where $\mathcal{P}_i = S \cup \{b_2, b_4, \ldots, b_{2i-2}\}$ is the set of arguments put forward by $\mathcal{PRO}$ so far. However, since the graph $(\mathcal{A}_x, R_x)$ is bipartite, $\mathrm{Refl}(\mathcal{A}_x, R_x) = R^{\pm}(\mathcal{P}_i) \cap R_x^-(b_{2i-1}) = \emptyset$.

Conclusion

The complexity of the credulous/sceptical acceptance problems has been studied in (Dimopoulos, Nebel, & Toni 2002; Dimopoulos & Torres 1996; Dunne & Bench-Capon 2002). The credulous acceptance problem is shown to be NP-complete, the sceptical acceptance problem is Π_2^p-complete. This is in accordance with our characterization of sceptical acceptance in terms of meta-credulous acceptance.

We have proposed here a definition of proofs that an argument a is in every extension of an argumentation system. Proofs that an argument is *not* in every extension (and algorithms to find such proofs) have been proposed in (Vreeswijk & Prakken 2000; Cayrol, Doutre, & Mengin 2003): such a proof is an attacker x of a together with a proof that x is in at least one extension. Note that this proof theory is not complete: there may be cases where a is not in every extension although none of its attackers are in any extension. Completeness holds in some cases, in particular when the graph has no cycle of odd length.

Proofs of credulous acceptance, like those of (Amgoud & Cayrol 2002; Prakken & Sartor 1997; Vreeswijk & Prakken 2000; Cayrol, Doutre, & Mengin 2003), can also be used as proofs of sceptical acceptance when the argument system only has one extension. This is in particular the case when the graph has no cycle at all.

An important perspective is to design an algorithm that computes sceptical proofs, thereby answering queries of the form: is a given argument in every extension of the system? This can be done using for example the algorithm for answering queries about credulous acceptance of (Cayrol, Doutre, & Mengin 2003): this algorithm returns a proof that an argument is in at least one extension. It can be used at the meta-level, and calls itself at the base-level (however, at the base-level, the algorithm needs to be slightly modified in order to find *all* minimal proofs that an argument is in some extension of the base system). Such an algorithm could be combined with an algorithm looking for proofs that an argument is not in every extension, thus providing proofs in both cases, at least when the graph has no odd cycle. A complete characterization of proofs that an argument is not in every extension when there are odd cycles remains to be done.

References

Amgoud, L., and Cayrol, C. 2002. A reasoning model based on the production of acceptable arguments. *Annals of Mathematics and Artificial Intelligence* 34:197–215.

Berge, C. 1973. *Graphs and Hypergraphs*. North-Holland: North-Holland Mathematical Library.

Bondarenko, A.; Dung, P.; Kowalski, R.; and Toni, F. 1997. An abstract, argumentation-theoretic approach to default reasoning. *Artificial Intelligence* 93:63–101.

Cayrol, C.; Doutre, S.; and Mengin, J. 2003. On Decision Problems related to the preferred semantics for argumentation frameworks. *Journal of Logic and Computation* 13(3):377–403.

Dimopoulos, Y., and Torres, A. 1996. Graph theoretical structures in logic programs and default theories. *Theoretical Computer Science* 170:209–244.

Dimopoulos, Y.; Magirou, V.; and Papadimitriou, C. 1997. On kernels, Defaults and Even Graphs. *Annals of Mathematics and AI* 1–12.

Dimopoulos, Y.; Nebel, B.; and Toni, F. 2002. On the Computational Complexity of Assumption-based Argumentation for Default Reasoning. *Artificial Intelligence* 141(1–2):57–78.

Doutre, S. 2002. Autour de la sémantique préférée des systèmes d'argumentation. PhD Thesis, Paul Sabatier University, Toulouse.

Dung, P. 1995. On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77:321–357.

Dunne, P., and Bench-Capon, T. 2002. Coherence in Finite Argument Systems. *Artificial Intelligence* 141(1–2):187–203.

Jakobovits, H., and Vermeir, D. 1999. Dialectic Semantics for Argumentation Frameworks. In *Proc. ICAIL'99*, 53–62. ACM Press.

Kakas, A., and Toni, F. 1999. Computing Argumentation in Logic Programming. *Journal of Logic and Computation* 9(4):515–562.

Prakken, H., and Sartor, G. 1997. Argument-based logic programming with defeasible priorities. *Journal of Applied Non-Classical Logics* 7:25–75.

Vreeswijk, G., and Prakken, H. 2000. Credulous and Sceptical Argument Games for Preferred Semantics. In *Proc. JELIA'2000*, 239–253. LNAI 1919.

$$\left[\mathcal{PRO},\ \begin{matrix}[\mathrm{PRO}, d]\\ [\mathrm{OPP}, c]\\ [\mathrm{PRO}, a]\end{matrix}\right] \qquad [\mathcal{OPP}, [\mathrm{PRO}, b]] \qquad \left[\mathcal{PRO},\ \begin{matrix}[\mathrm{PRO}, d]\\ [\mathrm{PRO}, b]\end{matrix}\right]$$

$\mathcal{PRO}$'s first move shows that $\{a, d\}$ is admissible

$\mathcal{OPP}$ then plays an admissible set $\{b\}$ that attacks $\{a, d\}$

$\mathcal{PRO}$ concludes by proving that $\{b, d\}$ is admissible in (A, R)

Figure 2: A meta-dialogue proving that d is in every extension on Ex. 5