Dialectical Proof Theories for the Credulous Preferred Semantics of Argumentation Frameworks

Claudette Cayrol, Sylvie Doutre and Jérôme Mengin

Mai 2001

Rapport IRIT/2001-09-R

Dialectical Proof Theories for the Credulous Preferred Semantics of Argumentation Frameworks

Claudette Cayrol, Sylvie Doutre and Jérôme Mengin

Argumentation is a natural form of reasoning, in which two agents cooperate in order to establish the validity of a given argument that could be used to deduce some conclusion of interest. An interesting semantics for logical systems of argumentation is Dung's "preferred semantics", which ameliorates in some ways the better-known stable semantics. In this paper, we present proof theories for the credulous decision problem associated with the preferred semantics: is a given argument in at least one extension of a given argumentation framework? Our proof theories improve on the one by [VP00], in the sense that a proof for a given argument is usually shorter with our system.

Dialectical Proof Theories for the Credulous Preferred Semantics of Argumentation Frameworks

Claudette Cayrol, Sylvie Doutre and Jérôme Mengin

Institut de Recherche en Informatique de Toulouse
Université Paul Sabatier - 118 route de Narbonne - F 31062 Toulouse Cedex 4
{ccayrol,doutre,mengin}@irit.fr

Argumentation is a natural form of reasoning, in which two agents cooperate in order to establish the validity of a given argument that could be used to deduce some conclusion of interest. An interesting semantics for logical systems of argumentation is Dung's "preferred semantics", which ameliorates in some ways the better-known stable semantics. In this paper, we present proof theories for the credulous decision problem associated with the preferred semantics: is a given argument in at least one extension of a given argumentation framework? Our proof theories improve on the one by [VP00], in the sense that a proof for a given argument is usually shorter with our system.

1 Introduction

Argumentation is a natural form of reasoning, in which two agents cooperate in order to establish the validity of a given argument that could be used to deduce some conclusion of interest. During the process of argumentation, each agent forms and asserts arguments that contradict or undermine arguments proposed by the other agent. This dialogue normally goes on until one of the agents cannot reply anything new: the original argument is then considered valid or not, depending on which agent won the dispute, the proponent or the opponent. The connection between dialogue and argument games has been studied by many researchers (see e.g. [CML00] for further references). The formalization of this form of reasoning has recently captured the interest of many researchers in the Artificial Intelligence community. For example, logics of argumentation are used in the construction of systems for legal reasoning, collective decision making or negotiation.

Besides being interesting because they capture this natural form of reasoning, logics of argumentation have also turned out to generalize in some way non-monotonic logics (that had themselves been designed as extensions of classical logic that do not collapse in the presence of inconsistencies). The first formulation of a decision problem related to Reiter's default logic in terms of a dialogue between two agents is probably due to Poole [Poo89]. Similar ideas have been used for building theorem provers for circumscription [Prz89,Gin89]. The close connection between non-monotonic logics and argumentation has been formally established in [BDKT97].

In both non-monotonic logics and logics of argumentation, one has to evaluate a set of pieces of knowledge (defaults or arguments) that can contradict each other. The computation of these contradictions can normally be performed by some classical theorem prover. Then, the evaluation of the defaults or of the arguments can often be based solely on the contradiction graph, where the vertices are the defaults or arguments, and where the directed edges represent the contradictions. One way to formalize this evaluation is to define acceptable sets of pieces of knowledge: usually, one would say that acceptable sets do not contain any contradiction, and are "strong enough" in some sense.

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 [DMP97, Ber73]. However, the stable semantics has some features that can be undesirable in some contexts: notably, it can happen that no set of pieces of knowledge is stable. [Dun95] defines an abstract framework for studying argumentation, and proposes several semantics. In particular, Dung's preferred semantics seems to capture well the intuition of "strong enough", and avoids several drawbacks of the stable semantics. In the preferred semantics, acceptable sets of arguments are called "preferred extensions".

An important problem related to argumentation systems is the credulous decision problem: given an argument and an argumentation framework, is the argument in at least one acceptable set of arguments? Proof theories for that problem have been described by [KT99] for the preferred semantics but for a slightly different form of argumentation systems. Proof theories in the form of argument games have also been described for the grounded semantics by [PS97, AC00], and by [VP00] for the preferred semantics. [VP00] propose argument games to answer credulous queries, and also to answer sceptical queries in a particular case. In these argument games, proofs that an argument belongs to some extension of an argumentation framework have the form of dialogues between a proponent and an opponent. An important aspect of such proofs is that they give an easy way to understand the implications of the underlying notions of acceptability.

On the algorithmic side, [DM01] show how to optimize the enumeration of the subsets of the set of arguments in order to efficiently answer questions related to the preferred semantics: what are the extensions, is an argument in all or some extensions? A careful study of the algorithm designed to answer the credulous decision problem shows that it can also be seen as a dialogue between an opponent and a proponent. Moreover, it appeared to us that the dialogue performed by the algorithm seems to be usually shorter than the proofs as defined by [VP00]. This prompted us to design another proof theory for the credulous decision problem related to the preferred semantics.

We present below two proof theories for that problem, based on the dialectical framework of [JV99]. They both improve on the one by [VP00] in the sense that proofs for a given argument are usually shorter with our system.

The paper is built as follows: the next section presents Dung's basic definitions and the preferred semantics. Section 3 presents a general framework for defining dialectical proofs. Our proof theories for the credulous decision problem associated with the preferred semantics are defined in Sect. 4. We show in Sect. 5 that the algorithm in [DM01] can compute our proofs, and we explain why they are usually shorter than the ones by [VP00]. A part of this report has been published in [CDM01].

2 The Preferred Semantics

Definition 1 [Dun95] An argumentation framework is a pair (A,R) where A is a set of arguments and R is a binary relation over arguments, i.e. $R \subseteq A \times A$. Given two arguments a and b, $(a,b) \in R$ or aRb means a attacks b (a is said to be an attacker of b). Moreover we say that a set $S \subseteq A$ of arguments attacks an argument a if some argument b in S attacks a. An argument $a \in R$ is self-attacking if $(a,a) \in R$. The set of the self-attacking arguments of (A,R) is denoted by Refl(A,R) (or Refl for short).

An argumentation framework can be simply represented as a directed graph where vertices are the arguments and edges correspond to the elements of R. Figure 1 shows an example of argumentation framework. Given an argument $a \in A$, we denote the set of the successors of a by $R^+(a) = \{b \in A \mid (a,b) \in R\}$, the set of its predecessors by $R^-(a) = \{b \in A \mid (b,a) \in R\}$, and the set $R^+(a) \cup R^-(a)$ by $R^{\pm}(a)$. Moreover, given a set $S \subseteq A$ of arguments and $\varepsilon \in \{+, -, \pm\}$, $R^{\varepsilon}(S) = \cup_{a \in S} R^{\varepsilon}(a)$.

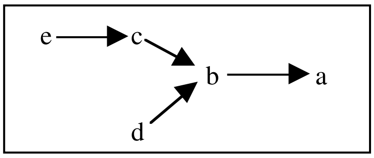

Fig. 1. The graph representation of an argumentation framework called AF1

Definition 2 Let (A,R) be an argumentation framework. An argument $a \in A$ is defended by a set $S \subseteq A$ of arguments (or S defends a) if and only if $\forall b \in A$, if bRa then S attacks b, i.e. $\exists c \in S$ such that cRb. 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. A set $S \subseteq A$ is admissible if S is conflict-free and $\forall x \in S$, S defends x.

Dung defines the preferred semantics of an argumentation framework by the set of preferred extensions. Precisely, a preferred extension is a maximal (with respect to set inclusion) admissible set of arguments. We characterise the preferred extensions on the graph representation of the argumentation framework.

Proposition 1 Given an argumentation framework (A,R), a subset S of A is a preferred extension if and only if the following conditions hold: 1) $R^+(S) \cap S = \varnothing$ (S is conflict-free); 2) $R^-(S) \subseteq R^+(S)$ (S defends every element it contains); 3) for every non-empty $X \subseteq A \setminus S$, $X \cap R^+(S \cup X) \neq \varnothing$ or $R^-(S) \not\subseteq R^+(S \cup X)$ (S is $\subseteq$-maximal such that 1 and 2). The set of the preferred extensions of (A,R) is denoted by Pref(A,R).

Dung exhibits interesting properties of the preferred semantics: in particular, every admissible set is contained in a preferred extension and every argumentation framework possesses at least one preferred extension.

The purpose of this paper is to answer an important question on preferred extensions: given an argument and an argumentation framework (A,R), is the argument in at least one preferred extension of (A,R), or equivalently, is the argument a credulous consequence of (A,R)? We define formally this notion:

Definition 3 Given an argumentation framework (A,R) and an argument a $\in$ A, a is a credulous consequence of (A,R) under the preferred semantics if and only if a is contained in the union of the preferred extensions of (A,R).

Example Given AF1, argument a is defended by S={a,d,e} against b. S is conflict-free and defends all its elements, so it is an admissible set, just like the sets $\emptyset$, {d}, {e}, {a,d}, {d,e} and {b,e}. The preferred extensions of AF1 are {a,d,e} and {b,e}. a, b, d and e are credulous consequences of AF1 under the preferred semantics, c is not.

Since a preferred extension is a maximal admissible set, deciding if an argument a is contained in a preferred extension of (A,R), amounts to deciding if it is contained in an admissible set. A procedure to solve this decision problem can take the form of a game between two players, one trying to build an admissible set containing a (the proponent), the other one trying to show it is not possible (the opponent).

Example We want to decide if a is a credulous consequence of AF1. The proponent starts trying to build an admissible set containing a by advancing a. The opponent says that a is attacked by b. The proponent defends a by advancing c. But the opponent replies that c is attacked by e. The proponent cannot defend c against e, so he advances d another defender of a. The opponent has nothing to say since d is not attacked. The proponent has built an admissible set: {a,d}, so a is a credulous consequence of AF1.

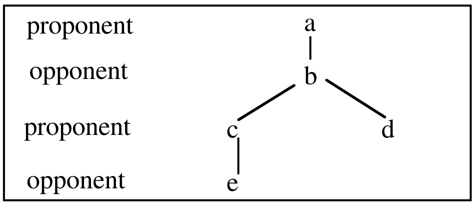

Fig. 2. Argument game to decide if a is a credulous consequence of AF1

3 The Dialectical Framework

Our purpose is to define a dialectical proof theory for the credulous preferred semantics, which takes into account the ideas of the credulous query answering algorithm of [DM01]. A general method for answering such a query takes the form of an argument game between a proponent (PRO) and an opponent (OPP). The proponent starts with the argument to be tested, and attempts to defend that argument against any attack coming from the opponent. The precise rules of the argument game depend on the semantics to be captured.

Argument games have been formalised in [JV99], where a general framework is proposed which enables to define dialectical proofs for winning positions in argumentation frameworks. The formalism developed by [JV99] encompasses the argument games of [PS97] and provides proof theories for two interesting semantics : the "robust" and the "defensible" semantics.

Following the methodological approach of [VP99], but with slightly different definitions, we propose in this section a dialectical framework which will enable us to provide two original proof theories for the credulous preferred semantics, in section 4.

3.1 Dialogue Type

An argument game is formalised by a dialogue between two players PRO and OPP. A dialogue takes place in a given argumentation framework and is governed by rules expressed in a so-called "legal-move" function.

Definition 4 A move in A is a pair [P,X] where P $\in$ {PRO, OPP} and X $\in$ A. For a move m = [P, X], we use pl(m) to denote P and arg(m) to denote X. A dialogue type is a tuple (A,R,Φ) where (A,R) is an argumentation framework and $\Phi : A^* \to 2^A$ is a function[1] (called "legal-move" function). A dialogue d in (A,R,Φ) (or Φ-dialogue for short) is a countable sequence $m_0m_1,\ldots$ of moves in A such that :

(i) $pl(m_0) = PRO$
(ii) $pl(m_i) \neq pl(m_{i+1})$
(iii) $arg(m_{i+1}) \in \Phi(arg(m_0) \ldots arg(m_i))$

d is about the argument $arg(m_0)$.

So, each player plays in turn and PRO plays first. The next move is legal with respect to the preceding moves.

Remark In [JV99], a conflict-free set of arguments may appear in a move. However, an additional requirement is that $arg(m_{i+1})$ must attack $arg(m_i)$.

Notations Let d = $m_0m_1,\ldots\ m_i$ a finite Φ-dialogue.
m_i is denoted by last(d).
$\Phi(arg(m_0) \ldots arg(m_i))$ is denoted by $\Phi(d)$.
PRO(d) will denote the set of arguments advanced by PRO during d.
Let m be a move in A such that $m_0m_1,\ldots\ m_im$ is a Φ-dialogue. The extension of d with m is denoted by the juxtaposition d m.

Any restriction can be included in the "legal-move" function as for instance: an argument advanced in a move attacks the argument advanced in the previous move; PRO cannot repeat himself; no player can introduce a self-attacking argument.

---

[1] A* denotes the set of finite sequences of elements from A.

3.2 Proof Theory

As in any game, we must give winning criteria in order to determine which argument can be successfully defended with Φ-dialogues. We consider two criteria among those proposed by [JV99]. A given dialogue about an argument x can be won, or there can be a winning strategy for x, that is a way for PRO to defend x against all attacks of OPP.

Definition 5 Let d be a Φ-dialogue. d is won by PRO iff d is finite, cannot be continued (that is Φ (d)= $\emptyset$), and last(d) is played by PRO.

The next definition is simpler but equivalent to the one proposed in [JV99].

Definition 6 A Φ−winning strategy for x is a non-empty finite set S of finite Φ-dialogues about x such that : $\forall$ d $\in$ S, $\forall$ d' prefix[2] of d such that last(d') is played by PRO, $\forall$ y $\in$ Φ(d'), $\exists$ d'' $\in$ S such that d'' is won by PRO and d'' is an extension of d'[OPP,y].

The following result provides another characterization of Φ−winning strategies in which only dialogues which cannot be continued are considered.

Proposition 2 There exists a Φ−winning strategy for x iff there exists a finite non-empty set S of finite Φ-dialogues about x won by PRO such that : $\forall$ d $\in$ S, $\forall$ d' prefix of d such that last(d') is played by PRO, $\forall$ y $\in$ Φ(d'), $\exists$ d'' $\in$ S such that d'' is an extension of d'[OPP,y].

4 Proof Theories for Credulous Preferred Semantics

The combination of a dialogue type and a winning criterion determines a proof theory. In this section, we present two specific proof theories dedicated to the credulous decision problem for the preferred semantics. The problem is to decide if an argument x belongs to a preferred extension. The basic idea is to prove that an admissible set of arguments can be built around x, with appropriate strategies for choosing attackers and defenders of the argument x. So, the "legal-move" functions we propose are inspired by the strategies used in the [DM01] algorithm.

Let d be a finite Φ-dialogue. $R^{\pm}$(PRO(d)) contains the arguments which attack or which are attacked by an argument avanced by PRO during d. Since PRO attempts to build an admissible set of arguments, PRO cannot choose any argument in $R^{\pm}$(PRO(d)) for pursuing the dialogue d. Nor any self-attacking argument. Let POSS(d) denote the set of arguments which may be chosen by PRO for extending the admissible set PRO(d). Formally, POSS(d) = A \ (PRO(d) $\cup$ $R^{\pm}$(PRO(d)) $\cup$ Refl).

The role of OPP is to attack one of the previous arguments advanced by PRO in d. But it is useless for OPP to advance an argument which is attacked by PRO(d).

---

[2] The sequence y is prefix of the sequence x, or x is an extension of y iff there exists a sequence z such that x is obtained by the concatenation of y and z, x = yz.

4.1 The $\Phi 1$–Proof Theory

Definition 7 Let $\Phi 1 : A^* \rightarrow 2^A$ defined by :

If $d = m_0 m_1, \ldots m_{2i}$ $\Phi 1(d) = R^-(PRO(d)) \setminus R^+(PRO(d))$

If $d = m_0 m_1, \ldots m_{2i+1}$ $\Phi 1(d) = R^-(arg(m_{2i+1})) \cap POSS(d)$

Combining $\Phi 1$ and the first winning criterion, we obtain $\Phi 1$ -proofs.

Definition 8 A $\Phi 1$ -proof for the argument x is a $\Phi 1$ -dialogue about x won by PRO.

The following results establish the soundness and the completeness of the $\Phi 1$ -proof theory.

Proposition 3 (Soundness) If d is a $\Phi 1$ -proof for the argument x, then PRO(d) is an admissible set containing x.

Proposition 4 (Completeness) If the argument x is in a preferred extension of the argumentation framework (A,R), then there exists a $\Phi 1$ -proof for x.

Example Let AF2 as indicated on figure 3. Let us try to build a $\Phi 1$-proof for a. PRO plays a. OPP can respond with b, c or d, since these arguments are predecessors but not successors of a. Assume OPP responds with b. This argument has two attackers: i and j, which can be advanced by PRO. Assume PRO advances i. Since $R^-(\{a,i\}) \setminus R^+(\{a,i\})=\{c\}$, OPP can only advance c (Note that c attacks a). Then PRO responds with f. OPP cannot play anymore, since $R^-(\{a,i,f\}) \setminus R^+(\{a,i,f\})=\emptyset$. The dialogue cannot be continued, it is won by PRO, thus it constitutes a $\Phi 1$-proof for a.

4.2 The $\Phi 2$–Proof Theory

In order to compare our work with argument games, it is convenient to present proofs in a more traditional way, where at each stage of the proof, the advanced argument attacks the previous one. Such proofs are obtained via the following dialogue type.

Definition 9 Let $\Phi 2 : A^* \rightarrow 2^A$ defined by :

If $d = m_0 m_1, \ldots m_{2i}$ $\Phi 2(d) = R^-(arg(m_{2i})) \setminus R^+(PRO(d))$

If $d = m_0 m_1, \ldots m_{2i+1}$ $\Phi 2(d) = R^-(arg(m_{2i+1})) \cap POSS(d)$

$\Phi 2$ is a restriction of $\Phi 1$ since, according to $\Phi 2$, OPP must advance an argument which attacks the argument advanced in the previous move.

Combining $\Phi 2$ and the second winning criterion, we obtain $\Phi 2$ -proofs.

Definition 10 A $\Phi 2$ -proof for the argument x is a $\Phi 2$ -winning strategy S for x such that $\cup(PRO(d), d \in S)$ is conflict-free.

Both proof theories are equivalent as shown by the following result.

Proposition 5 There exists a $\Phi 1$ -proof for the argument x iff there exists a $\Phi 2$ -proof for x.

Example Given AF2, let us try to build a Φ2-proof of a. PRO plays a. a has three attackers: b, c and d. Thus we have three Φ2–dialogues about a:

- in the first one (d1), OPP responds to a with b. Then PRO can advance i or j. Assume he advances i. Then OPP cannot respond since $R^-(\{i\}) \setminus R^+(\{a,i\})=\emptyset$.
- in the second one (d2), OPP responds to a with c. Then PRO can only advance f. OPP cannot respond since $R^-(\{f\}) \setminus R^+(\{a,f\})=\emptyset$.
- in the third one (d3), OPP responds to a with d. Then PRO can advance i or j. Assume he advances i. Then OPP cannot respond since $R^-(\{i\}) \setminus R^+(\{a,i\})=\emptyset$.

These three dialogues are finite, cannot be continued and are won by PRO. S = {d1, d2, d3} is a Φ2 -winning strategy for a. $\cup$(PRO(d), d$\in$ S) = {a, i, f} is conflict-free. So S is a Φ2-proof for a.

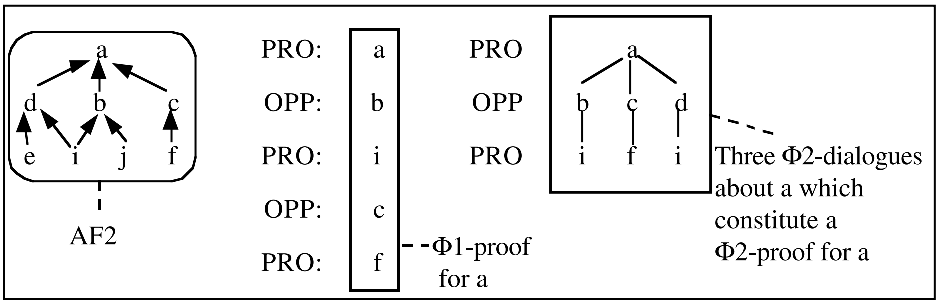

Fig. 3. The argumentation framework AF2, a Φ1-proof and a Φ2-proof for a

5 Related Works

5.1 The credulous query answering algorithm of [DM01]

We have proposed in [DM01] an algorithm to answer credulous queries related to the preferred semantics, of the form: given an argumentation framework (A,R) and an argument a in A, is a in at least one preferred extension of the framework? We have already mentionned that what prompted us to investigate a new proof theory is the fact that the proofs computed by our algorithm are in general shorter than the proofs of [VP00]. In this section, we show that our algorithm computes Φ1-proofs.

Let us first describe our algorithm. It is based on an enumeration of the subsets of A, which is performed by exploring a binary tree. Each node of the tree is labelled with a partition (I, O, U) of A: I is the set of arguments that have been put so far In the extension being built, O is the set of arguments that have been put Out of it, and U is the set of arguments that are still undecided at that stage. Since a preferred extension must be conflict free, the tree explores only nodes such that $R^{\pm}(I)\subseteq O$; in this case we say that (I,O,U) is an R-candidate. Normally, a node labelled with an R-candidate

C=(I,O,U) such that $U\neq\emptyset$ has two children: one is labelled with $C+x=(I\cup\{x\}, O\cup R^{\pm}(x), U\setminus(\{x\}\cup R^{\pm}(x)))$, the other one labelled with $C-x=(I, O\cup\{x\}, U\setminus\{x\})$, for some $x\in U$. Since we look for a preferred extension containing a, the root of the tree is labelled with the partition $(\{a\}, R^{\pm}(a), A \setminus (\{a\} \cup R^{\pm}(a)))$.

The set $Op = R^{-}(I) \setminus R^{+}(I)$ denotes the set of arguments which are predecessors but not successors of I, that is, these arguments attack I and I cannot defend itself against them. When a partition associated to a node n is such that Op is empty, then I is an admissible set: the computation has proved that a is in at least one preferred extension. When the partition is such that there is some $y \in Op$ which has no more undecided predecessor, then no preferred extension can be found on that branch, because an argument in I will never be defended by $(I \cup U)$: the exploration of that branch can be stopped. Thus the algorithm can be sketched in a functional programming style as follows:

function PrefEnum(R,C)
parameters: a binary relation R, and an R-candidate C=(I,O,U)
result: T if I is contained in at least one preferred extension, $\perp$ otherwise
if $Op=\emptyset$ **then** T
elseif there exists $x\in Op$ **such that** $R^{-}(x)\subseteq O$ **then** $\perp$
elseif there exists $y\in R^{-}(Op)$ **such that** yRy **then** PrefEnum(R,C−y)
else for some $x\in Op$ **such that** $R^{-}(x)\not\subset O$ **do**
    select some $y\in R^{-}(x) \setminus O$
    return(PrefEnum(R,C+y) **or** PrefEnum(R,C−y))

It is often possible to choose a "good" y such that only one branch has to be explored, but this is not relevant here. We will now show that when PrefEnum is called on $(\{a\}, R^{\pm}(a), A \setminus (\{a\} \cup R^{\pm}(a)))$, if the computation returns T then it is possible to extract a Φ1-proof for a from the successful branch.

Proposition 6 Let (A,R) be an argumentation framework, and let a be some element of the union of the preferred extensions of (A,R). Let C(0), C(1),…,C(n) be the sequence of R-candidates that label the nodes from the root to the success leaf of the tree explored by PrefEnum when called on $(\{a\}, R^{\pm}(a), A \setminus (\{a\} \cup R^{\pm}(a)))$. For $0\leq j\leq n$, C(j) is of the form C(j) = (I(j),O(j),U(j)); let $Op(j) = R^{-}(I(j)) \setminus R^{+}(I(j))$. Let $(j_i)_{1\leq i\leq m}$ be the longest subsequence of (1,…,n) such that for every $1\leq i\leq m$, $C(j_i)$ is of the form $C(j_i-1) + y(i)$, where $y(i) \in R^{-}(x(i)) \setminus O(j_i-1)$ for some $x(i) \in Op(j_i-1)$ with $R^{-}(x(i)) \not\subset O(j_i-1)$. Let d=(c(0), …, c(2m)) be the dialogue defined by:

- c(0) = [PRO,a]
- c(2i-1) = [OPP,x(i)] and c(2i) = [PRO,y(i)] for $1\leq i\leq m$.

Then d is a Φ1-proof for a.

Example Let us illustrate the above algorithm on the argumentation framework AF2 of figure 3. We look for a preferred extension containing a, so the root of the tree is labelled with the partition C(0)=({a},{b,c,d},{e,f,i,j}). We have Op(0)={b,c,d}. We select $i\in R^{-}(b)$ and we build two branches. One child of the root is labelled with C(0)+i, the other one is labelled with C(0)−i. We look at the one labelled with

$C(0)+i = C(1) = (\{a,i\}, \{b,c,d\}, \{e,f,j\})$. In this case, $Op(1)=\{c\}$. We select $f \in R^-(c)$ and we build two branches. One child of the node associated with the partition C(1) is labelled with $C(1)+f = C(2) = (\{a,i,f\}, \{b,c,d\}, \{e,j\})$. $Op(2)=\emptyset$, this is a leaf of success. The Φ1-proof for a shown on figure 3 can be easily extracted from this successful branch, given that $x(1)=b$, $y(1)=i$, $x(2)=c$, $y(2)=f$.

5.2 – Comparison with [VP00]

[VP00] have proposed an argument game for the credulous decision problem, related to the preferred semantics. The precise rules of that game are the following ones:

The argument advanced in a move (except the first one), attacks one of the previous arguments of the other player. A dispute is a succession of moves satisfying the following conditions: PRO plays first. A line of dispute is a succession of moves such that each player plays in turn and attacks the previous argument proposed by the other player. In a same line of dispute, OPP cannot repeat himself, but OPP can repeat an argument already advanced by PRO. Each player can backtrack; backtracking consists in opening a new line of dispute. OPP can repeat himself in different lines of dispute. PRO can repeat himself but he cannot repeat an argument already proposed by OPP.

OPP wins a dispute if PRO cannot respond to the last move of OPP, or if the last move of OPP is an argument already proposed by PRO and on which there were no backtrack. PRO wins a dispute if OPP does not win.

A given argument is contained in a preferred extension if and only if each dispute beginning with that argument is won by PRO.

Example (from [VP00]) Let AF3 as indicated on figure 4. There are two disputes starting with f. Let us build one of them, as done in [VP00]. PRO plays f. OPP advances n. PRO can respond with i or j. Assume PRO responds with i. Then OPP attacks i with j. To defend i against j, PRO repeats i. According to the [VP00] rules, the dispute cannot be continued, it is won by PRO. Note that it reduces to a single line of dispute. The second dispute starting with f is obtained by exchanging i and j. So, it is also won by PRO.

Similarly, there are two Φ1-proofs for f, which are also Φ2-proofs. Each one is shorter than the corresponding dispute. The three first moves are the same. The Φ1-proof stops at the third move, since at that stage, $R^-(\{f,i\}) \setminus R^+(\{f,i\})=\emptyset$.

Example (from [VP00]) Let AF4 as indicated on figure 5. Let us build a dispute lost by PRO, starting with m, as done in [VP00]. PRO plays m. OPP responds with l. PRO can advance p or k. Assume PRO advances p. Then OPP responds with h. PRO cannot respond to h since h has no predecessor. This line of dispute cannot be continued, but we can open another one where PRO proposes k in response to l. But then OPP attacks k with m and PRO cannot respond, since PRO is not allowed to repeat l, already proposed by OPP. Consequently, m is not a credulous consequence of AF4.

We reach this conclusion faster with our Φ2-proof theory. Actually, to respond to l advanced by OPP, PRO can only play p: k is attacked by m, an argument previously advanced by PRO. OPP attacks p with h and the dialogue cannot be continued. This dialogue is lost by PRO. It corresponds to the first line of dispute.

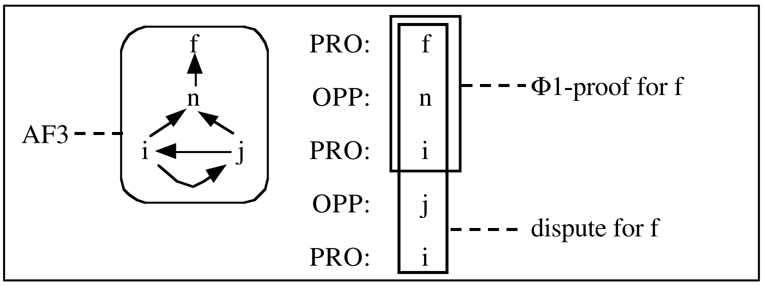

Fig. 4. The argumentation framework AF3, a dispute and a Φ1-proof for f

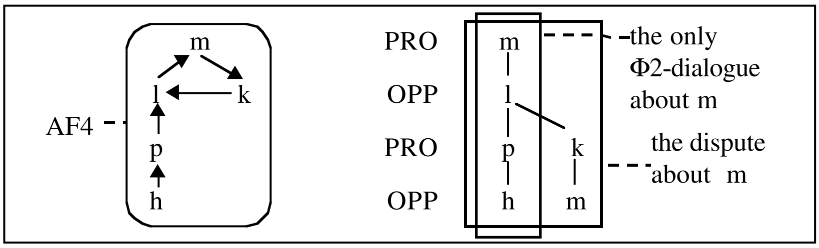

Fig. 5. The argumentation framework AF4, the Φ2-dialogue and the dispute about m

Our proofs are shorter than [VP00]'s proofs for a simple reason: we eliminate the arguments which are successors of PRO(d) from the possible moves of PRO and OPP. It is clear that such arguments are rejected from any admissible set containing PRO(d), so it is useless for OPP to advance these arguments (see figure 4). Moreover, such arguments are forbidden for PRO if we want PRO(d) to be contained in a conflict-free set (see figure 5).

6. Conclusion

The work reported in this paper concerns the credulous decision problem related to the preferred semantics. We have proposed two proof theories for that problem. These proof theories are presented in a dialectical framework, as done in [JV99]. Both proof theories improve on the work by [VP00] in the sense that our proofs for a given argument are usually shorter than the argument games proposed by [VP00].

Moreover, we have shown that our proofs can be computed by the credulous query answering algorithm of [DM01]. These results suggest to follow the same approach for the skeptical decision problem since [DM01] also present an efficient skeptical query answering algorithm. We project to design a dialectical proof theory for the skeptical preferred semantics.

References

[AC00] Amgoud, L., Cayrol, C.: A reasoning model based on the production of acceptable arguments. In Proceedings of the Workshop on Uncertainty Frameworks in Non-monotonic Reasoning, NMR 2000, Breckenbridge, Colorado, April 9-11, (2000).

[BDKT97] Bondarenko, A., Dung, P.M., Kowalski, R.A., Toni, F.: An abstract, argumentation-theoretic approach to default reasoning. Artificial Intelligence 93 (1997), 63-101.

[Ber73] Berge, C.: Graphs and Hypergraphs, vol.6 of North-Holland Mathematical Library, North-Holland (1973).

[CDM01] Cayrol, C., Doutre, S., Mengin, J.: Dialectical Proof Theories for the Credulous Preferred Semantics of Argumentation Frameworks. In Proceedings of the European Conference on Symbolic and Quantative Approaches to Reasoning with Uncertainty, EC-SQARU-2001, Toulouse, Sept. 17-21, (2001).

[CML00] Chesñevar, C.I., Maguitman, A.G., Loui, R.P.: Logical Models of Argument. ACM Computing Surveys 32 (4) (2000), 337-383.

[DMP97] Dimopoulos, Y., Magirou, V., Papadimitriou, C.H.: On kernels, Defaults and Even Graphs. Annals of Mathematics and AI (1997).

[Dun95] Dung, P.M.: On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n-person games. Artificial Intelligence 77 (1995), 321-357.

[DM01] Doutre, S., Mengin, J.: Preferred Extensions of Argumentation Frameworks : Computation and Query Answering. In Proceedings of the Int. Joint Conf. on Automated Reasoning (IJCAR 2001), Siena, Italy, June 18-23 (2001).

[Gin89] Ginsberg, M.: A circumscriptive theorem prover. Artificial Intelligence, 39 (1989) 209-230.

[KT99] Kakas, A.C., Toni, F.: Computing Argumentation in Logic Programming. Journal of Logic and Computation 9 (4) (1999), 515-562.

[JV99] Jakobovits, H., Vermeir, D.: Dialectic Semantics for Argumentation Frameworks. In Proceedings of the 7th International Conference on Artificial Intelligence and Law, (ACM Press) (1999), 53-62.

[Poo89] Poole, D.: Explanation and prediction : an architecture for default and abductive reasoning. Computational Intelligence Journal 5 (1989) 97-110.

[Prz89] Przymusinski, T.: An algorithm to compute circumscription. Artificial Intelligence, 38 (1989) 49-73.

[PS97] Prakken, H., Sartor, G.: Argument-Based Extended logic Programming with Defeasible Priorities. Journal of Applied Non-classical Logics 7 (1997) 25-75.

[Rei80] Reiter, R.: A logic for default reasoning. Artificial Intelligence 13 (1980) 81-132.

[VP00] Vreeswijk, G., Prakken, H.: Credulous and Sceptical Argument Games for Preferred Semantics. In Proceedings of the European Workshop on Logics in Artificial Intelligence (JELIA 2000) LNA1 1919, Springer Verlag (2000).

Appendix

1. Proofs of section 3.2

In [JV99], a __winning strategy for x is defined as follows :

A __winning strategy for x is a prefix-closed non-empty finite set S of finite _-dialogues about x such that :

i) (d (S, (d' extension of d , d' won by PRO

ii) (d[PRO,a] (S, (b (_(d[PRO,a]), (d' (S such that d' is won by PRO and d' is an extension of d[PRO,a] [OPP,b].

First, we prove that condition (i) is not necessary in the above definition, since (ii) implies (i) when the set S is prefix-closed.

Proof : Assume S is a prefix-closed finite set of finite _-dialogues about x such that (ii). Let d (S.

- If last(d) is played by PRO and d cannot be continued, then we take d' = d.
- If last(d) is played by PRO and _(d)≠ (, d is in the form d1[PRO, a]. Due to (ii), (b (_(d), (d' (S such that d' is won by PRO and d' is an extension of d[OPP,b]. Since _(d)≠ (, there exists an extension of d won by PRO.
- If last(d) is played by OPP, d is in the form d1[OPP,b], with d1 (S since S is prefix-closed. last(d1) is played by PRO, and b (_(d1), so by (ii) there exists an extension of d1[OPP,b] (i.e. an extension of d) won by PRO.

Then, we prove that the above definition is equivalent to the following one.

Definition6 A Φ−winning strategy for x is a non-empty finite set S of finite Φ-dialogues about x such that : $\forall$ d $\in$ S, $\forall$ d' prefix of d such that last(d') is played by PRO, $\forall$ y $\in$ Φ(d'), $\exists$ d" $\in$ S such that d" is won by PRO and d" is an extension of d'[OPP,y].

We do not require that S is prefix-closed.

Proof : We prove that the assertions i) and ii) below are equivalent.

i) There exists a non-empty finite set S of finite Φ-dialogues about x such that : $\forall$ d $\in$ S, $\forall$ d' prefix of d such that last(d') is played by PRO, $\forall$ y $\in$ Φ(d'), $\exists$ d" $\in$ S such that d" is won by PRO and d" is an extension of d'[OPP,y].

ii) There exists a prefix-closed non-empty finite set S' of finite Φ-dialogues about x such that :
∀ d[PRO,a] ∈ S', ∀ b ∈ Φ(d[PRO,a]), ∃ d' ∈ S' such that d' is won by PRO and d' is an extension of d[PRO,a] [OPP,b].

(i) => (ii)

We define S' as the set of Φ-dialogues about x which are prefix of an element of S. S' is prefix-closed.

Let d[PRO,a] ∈ S', and b ∈ Φ(d[PRO,a]). There exists d1 ∈ S such that d[PRO,a] is a prefix of d1. Applying (i) there exists d'' ∈ S, d'' won by PRO and d'' extension of d[PRO,a][OPP,b]. By definition of S', d'' belongs to S' too.

(ii) => (i)

Take S = S'.
Let d ∈ S. Let d' prefix of d such that last(d') is played by PRO. Since S is prefix-closed, d' ∈ S. By (ii), ∀ b ∈ Φ(d'), ∃ d'' ∈ S such that d'' is won by PRO and d'' is an extension of d'[OPP,b].

Proposition2 There exists a Φ−winning strategy for x iff there exists a finite non-empty set S of finite Φ-dialogues about x won by PRO such that : ∀ d ∈ S, ∀ d' prefix of d such that last(d') is played by PRO, ∀ y ∈ Φ(d'), ∃ d'' ∈ S such that d'' is an extension of d'[OPP,y].
<u>Proof</u> :
Obvious for the sufficient condition.
Necessary condition : Assume there exists a Φ−winning strategy S for x.
We define S' = {d ∈ S / Φ(d)= ∅}. It is sufficient to prove that for each dialogue d ∈ S', last(d) is played by PRO.

Suppose that there exists d ∈ S' such that last(d) is played by OPP. d = d'[OPP, b]. By definition of S', d ∈ S. d' is a prefix of d, last(d') is played by PRO, and b ∈ Φ(d'). So, ∃ d'' ∈ S such that d'' is won by PRO and d'' is an extension of d'[OPP,b] = d. d'' is different from d since last(d'') is played by PRO. So there is a contradiction with Φ(d)= ∅.

2. Proofs of section 4.1

Proposition3 (Soundness) If d is a Φ1 -proof for the argument x, then PRO(d) is an admissible set containing x.
Lemma1 Given an argumentation framework (A,R) and a Φ1-dialogue d, PRO(d) is conflict-free.
<u>Proof:</u>
We prove the result by induction on the number of elements of PRO(d). If PRO(d) only contains one element, then it is trivially conflict-free. Suppose now that the

property is true for any Φ1-dialogue d such that PRO(d) contains at most n-1 elements, for some n>1. Let d be a Φ1-dialogue such that PRO(d) contains n elements. Suppose first that the last move of d is played by PRO: d has the form d = d'.[OPP,x].[PRO,y], where:

- $x \in \Phi1(d') = R^-(PRO(d')) \setminus R^+(PRO(d'))$
- $y \in \Phi1(d'.[OPP,x]) = R^-(x) \cap POSS(d'.[OPP,x]) = R^-(x) \setminus (PRO(d') \cup Refl \cup R^{\pm}(PRO(d'))$.

Then $PRO(d) = PRO(d') \cup \{y\}$, and $y \notin PRO(d')$, therefore d' contains strictly less than n elements. Thus, by induction hypothesis, PRO(d') is conflict-free. Since $y \notin R^{\pm}(PRO(d'))$, PRO(d) is conflict-free too. Suppose now that the last move of d is played by OPP: then d has the form d = d'.[OPP,x], where d' is a Φ1-dialogue such that PRO(d') = PRO(d), therefore PRO(d') contains n elements. We have just proved that in this case PRO(d') is conflict-free, hence PRO(d) is conflict-free.

Proof of Proposition3:

The first move of d is [PRO,x], thus $x \in PRO(d)$. The lemma proves that PRO(d) is conflict-free. Moreover, $\Phi1(d) = \emptyset = R^-(PRO(d)) \setminus R^+(PRO(d))$, thus $R^-(PRO(d)) \subset R^+(PRO(d))$. Therefore PRO(d) is admissible.

Proposition4 (Completeness) If the argument x is in a preferred extension of the argumentation framework (A,R), then there exists a Φ1 -proof for x.

Lemma2 Let (A,R) be an argumentation framework, and let d be a Φ1-dialogue, the last move of which is played by PRO. Let S be a subset of A minimal such that S is admissible and contains PRO(d). If $S \neq PRO(d)$, then there exists $x,y \in A$ such that the dialogue d' = d.[OPP,x].[PRO,y] is a Φ1-dialogue and S is minimal such that S is admissible and contains PRO(d').

Proof:

Since $PRO(d) \neq S$ and since S contains PRO(d), the minimality of S implies that PRO(d) is not admissible. From lemma1, we know that PRO(d) is conflict-free, thus $R^-(PRO(d)) \not\subset R^+(PRO(d))$. Let $x \in R^-(PRO(d)) \setminus R^+(PRO(d))$, since $PRO(d) \subset S$, and since S is admissible, there exists $y \in S$ such that $x \in R^+(y)$. Since $x \notin R^+(PRO(d))$, $y \notin PRO(d)$. Moreover, S is conflict-free, thus $y \notin R^{\pm}(PRO(d))$ and $y \notin Refl$, therefore $y \in \Phi1(d.[OPP,x])$. Since $\Phi1(d) = R^-(PRO(d)) \setminus R^+(PRO(d))$, $x \in \Phi1(d)$. Thus d'=d.[OPP,x].[PRO,y] is a Φ1-dialogue. We know that $PRO(d') = PRO(d) \cup \{y\} \subset S$ because $y \in S$. There remains to prove that no admissible set S' strictly included in S contains d'. Suppose that S' is a strict subset of S, such that $PRO(d') \subset S'$: then $PRO(d) \subset S$ because $PRO(d) = PRO(d') \setminus \{y\}$, thus, since S is minimal such that S is admissible and contains PRO(d), S' is not admissible.

Proof of proposition4:

Let S be a minimal subset of A such that S is admissible and contains x. Let N be the number of elements of S, let d_1 = [PRO,x] and let $(d_n)_{2\leq n\leq N}$ be a sequence defined by $d_n=d_{n-1}$.[OPP,y].[PRO,z], for $2\leq n\leq N$, where $y \in R^+(PRO(d_{n-1})) \setminus R^+(PRO(d_{n-1}))$ and $z \in S \cap R^-(y)$. The lemma above proves that the sequence is well defined, and that d_N is a $\Phi 1$-dialogue for x won by PRO (since $PRO(d_N) = S$ and $R^-(S) \subset R^+(S)$, $\Phi 1(d_N) = \emptyset$).

3. Proofs of section 4.2

Proposition5 There exists a $\Phi 1$ -proof for the argument x iff there exists a $\Phi 2$ -proof for x.

Proof of the necessary condition :

Let D be a a $\Phi 1$ -proof for the argument x.

We build S a $\Phi 2$–winning strategy for x by considering all the attackers allowed w.r.t. $\Phi 2$ for each argument in PRO(D). Indeed, for each move m_{2k} = [PRO,a_{2k}] played by PRO in D, for each b $\in \Phi 2(m_0 \ldots m_{2K})$:

- either b appears in a move m_{2p+1} of D and the move m_{2p+2} of D will be the child of [OPP,b] in a $\Phi 2$–dialogue,

- or b does not appear in D and the reason is that b is attacked by an argument a advanced by PRO in D. So [PRO,a] will be the child of [OPP,b] in a $\Phi 2$–dialogue.

We need the following notation : Let d be a dialogue and x $\in$ PRO(d). OPP(d, x) denotes the set of arguments which attack the argument x and which are advanced by OPP in d, after the move [PRO,x]. Formally, if m_{2k} =[PRO,x] is a move of d, OPP(d,x) = {y / y $\in R^-(x)$ and m_{2p+1} =[OPP,y] with $k \leq p$ is a move of d}.

Let D be a a $\Phi 1$ -proof for the argument x. D= $m_0 m_1, \ldots m_{2n}$ with $arg(m_0) = x$.

Last(D) is played by PRO and $\Phi 1(D) = \emptyset$.

First, we build by induction on p a finite set Sp of $\Phi 2$–dialogues about x of length 2p, such that for each d $\in$ Sp, PRO(d) $\subseteq$PRO(D).

*Case p = 0 : S0 ={ m_0}

*Assume that Sp has been defined, satisfying the above conditions.

Let d $\in$ Sp. Last(d) is of the form [PRO,a] and [PRO,a] is a move of D (m_{2k} for instance). For each b $\in \Phi 2(d)$, b $\in R^-(a) \setminus R^+(PRO(d))$.

Case1 : b $\in$ OPP(D,a). So there exists $i \geq k$ with b = $arg(m_{2i+1})$. We define s(b) = $arg(m_{2i+2})$. The dialogue d[OPP,b][PRO,s(b)] is an extension of d of length 2p+2, with [PRO,s(b)] $\in$ PRO(D).

Case2 : b $\notin$ OPP(D,a). Since $\Phi 1(D) = \emptyset$, b $\in R^-(a)$ and a $\in$ PRO(D), we have b $\in R^+(PRO(D))$. So there exists m_{2i} in D, i < k, such that $arg(m_{2i}) \in R^-(b)$. We define s(b) = $arg(m_{2i})$. The dialogue d[OPP,b][PRO,s(b)] is an extension of d of length 2p+2, with [PRO,s(b)] $\in$ PRO(D).

Now we prove that the dialogue d [OPP,b][PRO,s(b)] is a $\Phi 2$–dialogue.

Due to the definition of $\Phi 1$–dialogues, $s(b) \in R^-(b)$. Since $b \notin R^+(PRO(d))$, we have $s(b) \notin PRO(d)$. Moreover, $s(b) \notin R^{\pm}(PRO(D))$. $PRO(d) \subseteq PRO(D)$, then $s(b) \notin R^{\pm}(PRO(d))$. Then, $s(b) \in POSS(d)$ and $s(b) \in \Phi 2$ (d [OPP,b]).

We define $Sp+1 = \cup_{d \in Sp} \{ d[OPP,b][PRO,s(b)], b \in \Phi 2(d)\}$.

Now, we prove that the above construction will stop.
Let $d \in Sp$. If d is not won by PRO, d can be extended in $d' \in Sp+1$. $d' = d[OPP,b][PRO,s(b)]$, with $b \in \Phi 2(d)$. Since $s(b) \notin PRO(d)$, PRO(d) is strictly included in PRO(d') and $PRO(d') \subseteq PRO(D)$. D is a $\Phi 1$ -proof, so D is finite and PRO(D) is finite.
So, there exists p such that for each $d \in Sp$, d cannot be continued (i.e. d is won by PRO) . Then $Sp+1 = Sp$, and we define $S = S0 \cup \ldots \cup Sp$.
By construction of S, for each $d \in S$, there exists an extension of d in Sp.

Finally, we prove that S is a $\Phi 2$-winning strategy for x.
Let $d \in S$ and d' prefix of d such that last(d') is played by PRO. By construction of S, $d \in Si$, and $d' \in Sk$, with $k \leq i \leq p$.
We have to prove that $\forall b \in \Phi 2(d')$, $\exists d'' \in S$ such that d" is won by PRO and d" is an extension of d'[OPP,b].
Assume that $b \in \Phi 2(d')$. Then $k < p$ and $d'[OPP,b][PRO,s(b)] \in Sk+1$. So there exists $d'' \in Sp$ such that d" is an extension of d'[OPP,b][PRO,s(b)]. d" is won by PRO and is also an extension of d'[OPP,b].

Moreover, $\cup(PRO(d), d \in S)$ is conflict-free since it is included in PRO(D).

<u>Proof of the sufficient condition :</u>

Assume there exists a $\Phi 2$-winning strategy S for x such that $\cup(PRO(d), d\in S)$ is conflict-free.We may assume that each dialogue in S is won by PRO (replace S with $S' = \{d \in S / \Phi 2(d)= \emptyset\}$).
S is non-empty. Let $d \in S$. We build a $\Phi 1$-proof D for x by providing successive extensions of d.

Case1 : d is a $\Phi 1$-dialogue won by PRO. So, d is a $\Phi 1$-proof for x. D = d.

Case2 : $\Phi 1(d) \neq \emptyset$. Let $b \in R^-(PRO(d)) \setminus R^+(PRO(d))$. We consider d' prefix of d of maximal length such that last(d') = [PRO,a], $b \in R^-(a)$. $b \notin R^+(PRO(d'))$ so $b \in \Phi 2(d')$. Since $\Phi 2(d)= \emptyset$, $b \notin \Phi 2(d)$ and $d' \neq d$.

S is a $\Phi 2$-winning strategy for x, so there exists an extension of d'[OPP,b] in S. In other words , there exists an element of S of the form d'[OPP,b]d".
We extend d by considering d1 = d[OPP,b]d". Since $\cup(PRO(d), d\in S)$ is conflict-free, we are sure that $d'' \in POSS(d)$. So d1 is a $\Phi 1$-dialogue extending d.
We iterate the construction (d1 extended by considering d2, …) until the obtained extension D satisfies $\Phi 1(D) = \emptyset$. Then, D is a $\Phi 1$-proof D for x.

The iteration stops since at each step PRO(di) $\subseteq \cup$(PRO(d), d$\in$ S) and PRO(di) is strictly included in PRO(di+1) when Φ1(di)$\neq \emptyset$. Remember that S is a finite set of finite dialogues.

4. Proofs of section 5.1

Proposition6 Let (A,R) be an argumentation framework, and let a be some element of the union of the preferred extensions of (A,R). Let C(0), C(1),…,C(n) be the sequence of R-candidates that label the nodes from the root to the success leaf of the tree explored by PrefEnum when called on ({a}, $R^{\pm}$(a), A \ ({a} $\cup$ $R^{\pm}$(a))). For $0 \le j \le n$, C(j) is of the form C(j) = (I(j),O(j),U(j)); let Op(j) = R^-(I(j))-R^+(I(j)). Let $(j_i)_{1\le i\le m}$ be the longest subsequence of (1,…,n) such that for every $1 \le i \le m$, C(j_i) is of the form C(j_i-1) + y(i), where y(i) $\in$ R^-(x(i)) \ O(j_i-1) for some x(i) $\in$ Op(j_i-1) such that R^-(x(i)) $\not\subset$ O(j_i-1). Let d=(c(0), …, c(2m)) be the dialogue defined by:

- c(0) = [PRO,a]
- c(2i-1) = [OPP,x(i)] and c(2i) = [PRO,y(i)] for $1 \le i \le m$.

Then d is a Φ1-proof for a.

Proof :

Let us first note that arguments are only added to the set of "In" arguments when there is a call of the form PrefEnum(R,C+y). Therefore, for $1 \le i \le m$, I(j_i)={a, y(1), …, y(i)}. In particular, I(j_m) = PRO(d). Since the algorithm returns true, we must have R^-(I(j_m)) \ R^+(I(j_m)) = $\emptyset$, thus Φ1(d) = $\emptyset$. Let d(0) = [PRO,a], and let, for $1 \le i \le m$, d(i) = [PRO,a].[OPP,x(1)]…[OPP,x(i)].[PRO,y(i)]. With this notation, PRO(d(i)) = I(j_i) for every $1 \le i \le m$. Let $1 \le i \le m$, then Φ1(d(i-1)) = R^-(PRO(d(i-1))) \ R^+(PRO(d(i-1))) = Op(j_{i-1}), thus x(i) $\in$ Φ1(d(i-1)). Moreover, Φ1(d(i-1).[OPP,x(i)]) = R^-(x(i)) $\cap$ (A \ (PRO(d(i-1).[OPP,x(i)]) $\cup$ $R^{\pm}$(PRO(d(i-1).[OPP,x(i)])) $\cup$ Refl)) = R^-(x(i)) \ (I(j_{i-1}) $\cup$ $R^{\pm}$(I(j_{i-1})) $\cup$ Refl) = R^-(x(i)) \ ($R^{\pm}$(I(j_{i-1})) $\cup$ Refl) because x(i) $\notin$ R^+(I(j_{i-1})). An invariant of the algorithm (proved in [DM01]) is that $R^{\pm}$(I(j)) $\subset$ O(j), for $0 \le j \le n$. Moreover, if there is, at some stage, y $\in$ R^+(Op) such that yRy, that is, such that y $\in$ Refl, then PrefEnum is called on C-y, where C is the current candidate. Thus y(i) $\notin$ Refl. Thus y(i) $\in$ Φ1(d(i-1)).