

Computational Properties of Resolution-based Grounded Semantics

Pietro Baroni

Dip. Elettronica per l'Automazione
University of Brescia
Via Branze 38
25123 Brescia, Italy
baroni@ing.unibs.it

Paul E. Dunne

Dept. of Computer Science
University of Liverpool
Ashton Building
Liverpool L69 7ZF, United Kingdom
P.E.Dunne@liverpool.ac.uk

Massimiliano Giacomin

Dip. Elettronica per l'Automazione
University of Brescia
Via Branze 38
25123 Brescia, Italy
giacomin@ing.unibs.it

Abstract

In the context of Dung's theory of abstract argumentation frameworks, the recently introduced resolution-based grounded semantics features the unique property of fully complying with a set of general requirements, only partially satisfied by previous literature proposals. This paper contributes to the investigation of resolution-based grounded semantics by analyzing its computational properties with reference to a standard set of decision problems for abstract argumentation semantics: (a) checking the property of being an extension for a set of arguments; (b) checking agreement with traditional grounded semantics; (c) checking the existence of a non-empty extension; (d) checking credulous acceptance of an argument; (e) checking skeptical acceptance of an argument. It is shown that problems (a)-(c) admit polynomial time decision processes, while (d) is NP-complete and (e) coNP-complete.

1 Introduction

In the context of Dung's theory of abstract argumentation frameworks [Dung, 1995] a large variety of semantics have been proposed in the literature, their motivation and mutual comparison often being based on informal intuitions and/or specific examples. To overcome this limitation, recently a set of general principles for semantics evaluation and comparison have been identified [Baroni and Giacomin, 2007]. These principles are defined in term of formal properties of the extensions prescribed by a semantics and cover a wide spectrum of notions, ranging from set-theoretical (I-maximality) to topological ones (directionality), and from defense-related concepts (admissibility and reinstatement) to skepticism-related requirements (skepticism adequacy and resolution adequacy). The analysis of a comprehensive set of literature proposals (namely, complete, grounded, stable, preferred, ideal, semi-stable, CF2 and prudent semantics) has shown that none of them is able to respect all the desirable properties altogether. Subsequently, a newly proposed semantics called *resolution-based grounded semantics* (GR^* in the following) has been proved [Baroni and Giacomin, 2008] to satisfy all the desiderata of [Baroni and Giacomin, 2007].

While this might be regarded as a significant advantage of GR^* with respect to other proposals, an analysis of its computational properties is necessary to assess its potential practical relevance and to complete its comparison with other semantics from this important perspective. This paper provides such an analysis: it turns out that GR^* is satisfactory also from this viewpoint, its complexity in some significant computational problems being lesser or the same with respect to other multiple-status semantics. The paper is organized as follows: Sec. 2 introduces basic concepts, notations, and the considered decision problems, whose complexity analysis is provided in Sec. 3, while conclusions are given in Sec. 4.

2 Notation and Definitions

Definition 1 collects the basic notions we use in the paper.

Definition 1 An argumentation framework (AF) is a pair $\mathcal{G} = \langle \mathcal{A}, \mathcal{R} \rangle$, written also $\mathcal{G}(\mathcal{A}, \mathcal{R})$, where \mathcal{A} is a finite set of arguments and $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ is the attack relationship for \mathcal{G} . Given $S \subseteq \mathcal{A}$, the complement of S is denoted as $S^c = \mathcal{A} \setminus S$. The restriction of \mathcal{G} to $S \subseteq \mathcal{A}$ is the AF $\mathcal{G} \downarrow_S = \langle S, \mathcal{R} \cap (S \times S) \rangle$.

A pair $\langle x, y \rangle \in \mathcal{R}$ is referred to as 'y is attacked by x' or 'x attacks y'. For S, T subsets of arguments in $\mathcal{G}(\mathcal{A}, \mathcal{R})$, we say that $t \in T$ is attacked by S – written $\text{attacks}(S, t)$ – if $\exists s \in S : \langle s, t \rangle \in \mathcal{R}$, and analogously that S is attacked by t – written $\text{attacks}(t, S)$ – if $\exists s \in S : \langle t, s \rangle \in \mathcal{R}$. We define $S^+ = \{t \in \mathcal{A} \mid \text{attacks}(S, t)\}$ and $S^- = \{t \in \mathcal{A} \mid \text{attacks}(t, S)\}$. We use $\alpha(S)$ to denote the set $S \cup S^+$. For $x \in \mathcal{A}$, every argument y in $\{x\}^- \cap \{x\}^+$ is involved in a mutual attack with x , i.e. $\{\langle x, y \rangle, \langle y, x \rangle\} \subseteq \mathcal{R}$. A set $S \subseteq \mathcal{A}$ is conflict-free if $\nexists x, y \in S : \langle x, y \rangle \in \mathcal{R}$. Given $S, T \subseteq \mathcal{A}$, S is stable in T with respect to \mathcal{G} , denoted as $\text{st}_{\mathcal{G}}(S, T)$, if $\forall x \in (T \setminus S) \ x \in (S \cap T)^+$.

An argument $x \in \mathcal{A}$ is acceptable with respect to a set $S \subseteq \mathcal{A}$ if $\forall y \in \{x\}^- \ \text{attacks}(S, y)$. The function $\mathcal{F}_{\mathcal{G}} : 2^{\mathcal{A}} \rightarrow 2^{\mathcal{A}}$ which, given a set $S \subseteq \mathcal{A}$, returns the set of the acceptable arguments with respect to S , is called the characteristic function of \mathcal{G} . We also use the notation $\mathcal{F}_{\mathcal{G}}^1(S) \triangleq \mathcal{F}_{\mathcal{G}}(S)$ and for $i > 1$, $\mathcal{F}_{\mathcal{G}}^i(S) = \mathcal{F}_{\mathcal{G}}(\mathcal{F}_{\mathcal{G}}^{i-1}(S))$.

An argumentation semantics \mathcal{S} specifies for any AF $\mathcal{G}(\mathcal{A}, \mathcal{R})$ a set $\mathcal{E}_{\mathcal{S}}(\mathcal{G}) \subseteq 2^{\mathcal{A}}$ of extensions, each extension being a set of arguments which can "survive together" the conflict represented by \mathcal{R} . A minimal requirement shared by all literature semantics is that any extension is conflict-free.




Figure 1: \mathcal{G}_1 : an AF with two resolutions.




Figure 2: \mathcal{G}_2 : an AF with many resolutions.

It is proved in [Dung, 1995] that $\mathcal{F}_{\mathcal{G}}()$ has a least fixed point, i.e., in the case of a finite AF $\mathcal{G}(\mathcal{A}, \mathcal{R})$, there is a finite value i for which $\mathcal{F}_{\mathcal{G}}^i(\emptyset) = \mathcal{F}_{\mathcal{G}}^{i+k}(\emptyset)$ for all $k \geq 0$. The *grounded extension* of \mathcal{G} , denoted as $GE(\mathcal{G})$ is the (unique) least fixed point of $\mathcal{F}_{\mathcal{G}}()$. We use $\mathcal{E}_{GR}(\mathcal{G})$ to denote the set whose (sole) element is $GE(\mathcal{G})$. Letting $S = GE(\mathcal{G})$ we will also denote as $CUT(\mathcal{G})$ the argumentation framework with arguments $\mathcal{A}_S = \mathcal{A} \setminus \alpha(S)$ and attack relation $\mathcal{R}_S = \mathcal{R} \cap (\mathcal{A}_S \times \mathcal{A}_S)$, i.e. $CUT(\mathcal{G}) = \mathcal{G} \downarrow_{(\mathcal{A} \setminus \alpha(S))}$.

The concept of *resolution-based semantics* was introduced in [Baroni and Giacomin, 2008] as a mechanism for treating mutual attacks between arguments. Given $\mathcal{G}(\mathcal{A}, \mathcal{R})$, let $M_{\mathcal{G}} \subseteq \mathcal{R}$ be the set of mutual attacks in \mathcal{G} , i.e. $M_{\mathcal{G}} = \{\langle x, y \rangle \in \mathcal{R} \mid x \neq y \wedge \langle y, x \rangle \in \mathcal{R}\}$ (note that self-attacking arguments are *not* considered to define mutual attacks). A (partial) resolution of \mathcal{G} is defined by any subset $\beta \subset M_{\mathcal{G}}$ for which *at most one element* of each of the pairs $\langle x, y \rangle, \langle y, x \rangle$ is in β . The AF \mathcal{G}_{β} arising from the partial resolution β is $\langle \mathcal{A}, \mathcal{R} \setminus \beta \rangle$. A *full resolution* is any partial resolution in which *exactly one element* of each mutual attack occurs, hence the AF \mathcal{G}_{γ} arising from any full resolution γ of \mathcal{G} contains no mutually attacking arguments.

To exemplify, the AF \mathcal{G}_1 shown in Fig. 1 includes only one mutual attack and has two non-empty resolutions (both full), namely $\{\langle c, d \rangle\}$ and $\{\langle d, c \rangle\}$. The AF \mathcal{G}_2 shown in Fig. 2 includes five mutual attacks. Given that to define any resolution there are three choices for each mutual attack (selecting one of the attacks or neither) and excluding the empty resolution it follows that \mathcal{G}_2 admits 242 resolutions. On the other hand, to define a full resolution there are two choices for each mutual attack entailing that \mathcal{G}_2 has 32 full resolutions, including for instance $\{\langle a, b \rangle, \langle c, a \rangle, \langle c, b \rangle, \langle g, f \rangle, \langle g, h \rangle\}$. It is easy to see that $GE(\mathcal{G}_1) = \{a\}$ and $GE(\mathcal{G}_2) = \{e\}$, hence $CUT(\mathcal{G}_1)$ and $CUT(\mathcal{G}_2)$ are obtained by suppressing arguments and attacks included in the dashed boxes in Figs. 1 and 2 respectively.

The definition of *resolution-based grounded semantics* GR^* is based on the grounded extensions of each different full resolution of an AF and selects those ones which are minimal with respect to set inclusion.

Definition 2 Given an AF $\mathcal{G}(\mathcal{A}, \mathcal{R})$ $\mathcal{E}_{GR^*}(\mathcal{G}) = \min \{S \subseteq \mathcal{A} : GE(\mathcal{G}_{\gamma}) = S \text{ for some full resolution } \gamma \text{ of } \mathcal{G}\}$.

Our main focus concerns computational properties of GR^* with respect to the informal questions and corresponding formal decision problems listed below:

- (a) Given $\mathcal{G}(\mathcal{A}, \mathcal{R})$ and $S \subseteq \mathcal{A}$ decide if S is an extension of GR^* . Formally this relates to the decision problem VER_{GR^*} with instances $\langle \mathcal{G}, S \rangle$ accepted if and only if $S \in \mathcal{E}_{GR^*}(\mathcal{G})$.

- (b) Given $\mathcal{G}(\mathcal{A}, \mathcal{R})$ determine whether its resolution-based grounded extensions are exactly the grounded extension of \mathcal{G} . Formally this relates to the decision problem $COIN_{GR, GR^*}$ with instances \mathcal{G} accepted if and only if $\mathcal{E}_{GR}(\mathcal{G}) = \mathcal{E}_{GR^*}(\mathcal{G})$.
- (c) Given $\mathcal{G}(\mathcal{A}, \mathcal{R})$ determine whether there is at least one non-empty resolution-based grounded extension of \mathcal{G} . Formally this relates to the decision problem NE_{GR^*} with instances \mathcal{G} accepted if and only if $\mathcal{E}_{GR^*}(\mathcal{G}) \neq \{\emptyset\}$.
- (d) Given $\mathcal{G}(\mathcal{A}, \mathcal{R})$ and $x \in \mathcal{A}$ determine whether x is *credulously* accepted with respect to $\mathcal{E}_{GR^*}(\mathcal{G})$. Formally this relates to the decision problem $CAGR^*$ with instances $\langle \mathcal{G}, x \rangle$ accepted if and only if $\exists S \in \mathcal{E}_{GR^*}(\mathcal{G})$ s.t. $x \in S$.
- (e) Given $\mathcal{G}(\mathcal{A}, \mathcal{R})$ and $x \in \mathcal{A}$ determine whether x is *skeptically* accepted with respect to $\mathcal{E}_{GR^*}(\mathcal{G})$. Formally this relates to the decision problem SA_{GR^*} with instances $\langle \mathcal{G}, x \rangle$ accepted if and only if $\forall S \in \mathcal{E}_{GR^*}(\mathcal{G})$ $x \in S$.

We show that while the problems (a), (b) and (c) admit polynomial time decision processes, in contrast (d) is NP-complete and (e) coNP-complete.

3 Decision properties of \mathcal{E}_{GR^*}

3.1 Polynomial time decidable problems

We need several preliminary lemmata concerning properties of GR . First, any argument attacked by the grounded extension receives at least one non-mutual attack from it.

Lemma 1 Given an AF $\mathcal{G}(\mathcal{A}, \mathcal{R})$, $\forall x \in (GE(\mathcal{G}))^+ \exists y \in GE(\mathcal{G}) : \langle y, x \rangle \in \mathcal{R} \wedge \langle x, y \rangle \notin \mathcal{R}$.

Proof: The proof is based on the property $GE(\mathcal{G}) = \bigcup_{i=1}^{\infty} \mathcal{F}_{\mathcal{G}}^i(\emptyset)$ [Dung, 1995] and proceeds by induction on i , showing that the thesis holds for any $x \in (\mathcal{F}_{\mathcal{G}}^i(\emptyset))^+$. As to the basis step, note that the thesis holds trivially for $x \in (\mathcal{F}_{\mathcal{G}}^1(\emptyset))^+$ since $\forall y \in \mathcal{F}_{\mathcal{G}}^1(\emptyset) \{y\}^- = \emptyset$. As to the inductive step, consider $x \in (\mathcal{F}_{\mathcal{G}}^{i+1}(\emptyset))^+$ with $i \geq 1$. Then $\exists y \in \mathcal{F}_{\mathcal{G}}^{i+1}(\emptyset) : \langle y, x \rangle \in \mathcal{R}$. If $\langle x, y \rangle \notin \mathcal{R}$ the thesis obviously holds, otherwise since $y \in \mathcal{F}_{\mathcal{G}}^{i+1}(\emptyset)$ it holds that $x \in (\mathcal{F}_{\mathcal{G}}^i(\emptyset))^+$ and, by the inductive hypothesis, $\exists y' \in \mathcal{F}_{\mathcal{G}}^i(\emptyset) : \langle y', x \rangle \in \mathcal{R} \wedge \langle x, y' \rangle \notin \mathcal{R}$. \square

The following lemma states that if $GE(\mathcal{G})$ is stable in a set T then the part of $GE(\mathcal{G})$ outside T coincides with the grounded extension of the AF obtained from \mathcal{G} by suppressing the arguments in T and those attacked by $GE(\mathcal{G}) \cap T$.

Lemma 2 Given an AF $\mathcal{G}(\mathcal{A}, \mathcal{R})$ and a set $T \subseteq \mathcal{A}$ such that $st_{\mathcal{G}}(GE(\mathcal{G}), T)$, it holds that $GE(\mathcal{G}) \cap T^C = GE(\mathcal{G} \downarrow_{T^C \setminus (GE(\mathcal{G}) \cap T)})$.

Proof: To shorten notation let $\bar{T} \triangleq T^C \setminus (GE(\mathcal{G}) \cap T)^+$. We will exploit again the property $GE(\mathcal{G}) = \bigcup_{i=1}^{\infty} \mathcal{F}_{\mathcal{G}}^i(\emptyset)$. Let us first show by induction on i that for any i ($\mathcal{F}_{\mathcal{G}}^i(\emptyset) \cap T^C$) $\subseteq GE(\mathcal{G} \downarrow_{\bar{T}})$. As to the base case, for any argument x in ($\mathcal{F}_{\mathcal{G}}^1(\emptyset) \cap T^C$) it holds that $\{x\}^- = \emptyset$ in \mathcal{G} and hence in any restriction of \mathcal{G} including x . It follows that $x \in \bar{T}$ and $x \in GE(\mathcal{G} \downarrow_{\bar{T}})$. Let us now assume inductively that ($\mathcal{F}_{\mathcal{G}}^i(\emptyset) \cap T^C$) $\subseteq GE(\mathcal{G} \downarrow_{\bar{T}})$ and show that ($\mathcal{F}_{\mathcal{G}}^{i+1}(\emptyset) \cap T^C$) $\subseteq GE(\mathcal{G} \downarrow_{\bar{T}})$. Letting $x \in (\mathcal{F}_{\mathcal{G}}^{i+1}(\emptyset) \cap T^C)$, we will show that x is acceptable with respect to $GE(\mathcal{G} \downarrow_{\bar{T}})$ in $\mathcal{G} \downarrow_{\bar{T}}$, namely $\forall y \in (\{x\}^- \cap \bar{T}) \quad y \in (GE(\mathcal{G} \downarrow_{\bar{T}}))^+$. In fact, since $x \in \mathcal{F}_{\mathcal{G}}^{i+1}(\emptyset)$, it holds that $y \in (\mathcal{F}_{\mathcal{G}}^i(\emptyset))^+$. Now, $\mathcal{F}_{\mathcal{G}}^i(\emptyset) \subseteq GE(\mathcal{G})$ and since $y \notin (GE(\mathcal{G}) \cap T)^+$ it follows that $y \in (\mathcal{F}_{\mathcal{G}}^i(\emptyset) \cap T^C)^+$, which, by the inductive hypothesis, entails $y \in (GE(\mathcal{G} \downarrow_{\bar{T}}))^+$. It follows that $x \in GE(\mathcal{G} \downarrow_{\bar{T}})$ given the well-known fact that the grounded extension is complete [Dung, 1995], i.e. it contains all arguments acceptable with respect to it.

We have now to prove that $GE(\mathcal{G} \downarrow_{\bar{T}}) \subseteq (GE(\mathcal{G}) \cap T^C)$, by showing (again by induction on i) that for any i $\mathcal{F}_{\mathcal{G} \downarrow_{\bar{T}}}^i(\emptyset) \subseteq (GE(\mathcal{G}) \cap T^C)$, which, given the definition of \bar{T} , is obviously equivalent to $\mathcal{F}_{\mathcal{G} \downarrow_{\bar{T}}}^i(\emptyset) \subseteq GE(\mathcal{G})$. Preliminarily, we show that $\forall x \in \bar{T}, \forall y \in \{x\}^-$ such that $y \notin \bar{T}, y \in (GE(\mathcal{G}))^+$. In fact $y \notin \bar{T}$ implies $y \in (T^C)^C \cup (T^C \cap (GE(\mathcal{G}) \cap T)^+) = T \cup (T^C \cap (GE(\mathcal{G}) \cap T)^+)$. We have two possible cases. If $y \in T$ then, in particular, $y \in (T \setminus GE(\mathcal{G}))$ since $x \notin (GE(\mathcal{G}) \cap T)^+$. Then, since $st_{\mathcal{G}}(GE(\mathcal{G}), T)$ it follows $y \in (GE(\mathcal{G}))^+$. If otherwise $y \in (T^C \cap (GE(\mathcal{G}) \cap T)^+)$ it follows in particular $y \in (GE(\mathcal{G}) \cap T)^+$ hence $y \in (GE(\mathcal{G}))^+$. Turning to the inductive proof, consider for the base case any $x \in \mathcal{F}_{\mathcal{G} \downarrow_{\bar{T}}}^1(\emptyset)$: we have that for any $y \in \{x\}^-$ in \mathcal{G} it must be the case that $y \notin \bar{T}$, which, as shown above, entails $y \in (GE(\mathcal{G}))^+$. It follows that x is acceptable with respect to $GE(\mathcal{G})$, hence $x \in GE(\mathcal{G})$. Now assume inductively that $\mathcal{F}_{\mathcal{G} \downarrow_{\bar{T}}}^i(\emptyset) \subseteq GE(\mathcal{G})$ and consider any $x \in \mathcal{F}_{\mathcal{G} \downarrow_{\bar{T}}}^{i+1}(\emptyset)$. For any $y \in \{x\}^-$ in \mathcal{G} we can consider two cases: if $y \in \bar{T}$ it must be the case that $y \in (\mathcal{F}_{\mathcal{G} \downarrow_{\bar{T}}}^i(\emptyset))^+$, which, by the inductive hypothesis, entails $y \in (GE(\mathcal{G}))^+$. Otherwise, $y \notin \bar{T}$ which again entails $y \in (GE(\mathcal{G}))^+$, as shown above. Summing up, it turns out that x is acceptable with respect to $GE(\mathcal{G})$, hence $x \in GE(\mathcal{G})$, which completes the proof. \square

An equality involving the operations of restriction and resolution of an argumentation framework will also be useful.

Lemma 3 Given an AF $\mathcal{G}(\mathcal{A}, \mathcal{R})$ and a set $T \subseteq \mathcal{A}$, for any resolution γ of \mathcal{G} it holds that $\mathcal{G}_{\gamma} \downarrow_T = (\mathcal{G} \downarrow_T)_{\gamma}$.

Proof: $\mathcal{G}_{\gamma} \downarrow_T = (\mathcal{A}, \mathcal{R} \setminus \gamma) \downarrow_T = (\mathcal{A} \cap T, (\mathcal{R} \setminus \gamma) \cap (T \times T)) = (\mathcal{A} \cap T, (\mathcal{R} \cap (T \times T)) \setminus \gamma) = (\mathcal{G} \downarrow_T)_{\gamma}$. \square

We can now show that the grounded extension of a “resolved” AF \mathcal{G}_{γ} can be “decomposed” into the grounded extension of the original AF \mathcal{G} and the grounded extension of the AF resulting from applying the same resolution to $CUT(\mathcal{G})$.

Lemma 4 Given an AF $\mathcal{G}(\mathcal{A}, \mathcal{R})$, for any resolution γ of \mathcal{G} it holds that $GE(\mathcal{G}_{\gamma}) = GE(\mathcal{G}) \cup GE(CUT(\mathcal{G})_{\gamma})$.

Proof: The fact that for any resolution γ $GE(\mathcal{G}) \subseteq GE(\mathcal{G}_{\gamma})$ is proved in [Baroni and Giacomin, 2007] in relation with the property of *resolution adequacy* of GR . Given that $GE(\mathcal{G}_{\gamma})$ is conflict-free, it follows that $GE(\mathcal{G}_{\gamma}) \cap \alpha(GE(\mathcal{G})) = GE(\mathcal{G})$. Clearly $GE(\mathcal{G})$ is stable in $\alpha(GE(\mathcal{G}))$, therefore taking into account Lemma 1 we have also $st_{\mathcal{G}_{\gamma}}(GE(\mathcal{G}_{\gamma}), \alpha(GE(\mathcal{G})))$. Then, letting $\mathcal{A}_S = \mathcal{A} \setminus \alpha(GE(\mathcal{G}))$, noting that $\mathcal{A}_S \setminus (GE(\mathcal{G}_{\gamma}) \cap \alpha(GE(\mathcal{G})))^+ = \mathcal{A}_S$ and using Lemmata 2 and 3, we have $GE(\mathcal{G}_{\gamma}) \cap \mathcal{A}_S = GE(\mathcal{G}_{\gamma} \downarrow_{\mathcal{A}_S}) = GE((\mathcal{G} \downarrow_{\mathcal{A}_S})_{\gamma}) = GE(CUT(\mathcal{G})_{\gamma})$. Putting together the two (disjoint) pieces we obtain $GE(\mathcal{G}_{\gamma}) = GE(\mathcal{G}) \cup GE(CUT(\mathcal{G})_{\gamma})$. \square

The decomposition identified in Lemma 4 can then be applied to the extensions of GR^* .

Corollary 1 Given an AF $\mathcal{G}(\mathcal{A}, \mathcal{R})$, $\mathcal{E}_{GR^*}(\mathcal{G}) = \{GE(\mathcal{G}) \cup T \mid T \in \mathcal{E}_{GR^*}(CUT(\mathcal{G}))\}$.

Proof: Using definitions and exploiting Lemma 4 at the second equality we have: $\mathcal{E}_{GR^*}(\mathcal{G}) = \min\{S \subseteq \mathcal{A} : GE(\mathcal{G}_{\gamma}) = S \text{ for some full resolution } \gamma \text{ of } \mathcal{G}\} = \min\{GE(\mathcal{G}) \cup GE(CUT(\mathcal{G})_{\gamma}) \text{ for some full resolution } \gamma \text{ of } \mathcal{G}\} = \min\{GE(\mathcal{G}) \cup GE(CUT(\mathcal{G})_{\gamma}) \text{ for some full resolution } \gamma \text{ of } CUT(\mathcal{G})\} = \{GE(\mathcal{G}) \cup T \mid T \in \min\{GE(CUT(\mathcal{G})_{\gamma}) \text{ for some full resolution } \gamma \text{ of } CUT(\mathcal{G})\}\} = \{GE(\mathcal{G}) \cup T \mid T \in \mathcal{E}_{GR^*}(CUT(\mathcal{G}))\}$. \square

Corollary 2 For any AF $\mathcal{G}(\mathcal{A}, \mathcal{R})$, $\mathcal{E}_{GR}(\mathcal{G}) = \mathcal{E}_{GR^*}(\mathcal{G})$ if and only if there is a full resolution γ of $CUT(\mathcal{G})$ such that in $CUT(\mathcal{G})_{\gamma}$ every argument has at least one attacker.

Proof: From Corollary 1 we have $\mathcal{E}_{GR}(\mathcal{G}) = \mathcal{E}_{GR^*}(\mathcal{G}) \Leftrightarrow \mathcal{E}_{GR^*}(CUT(\mathcal{G})) = \{\emptyset\}$, i.e. if and only if there is a full resolution γ of $CUT(\mathcal{G})$ such that $GE(CUT(\mathcal{G})_{\gamma}) = \emptyset$. This entails the conclusion by recalling that $GE(\mathcal{G}) = \emptyset$ if and only if $\mathcal{F}_{\mathcal{G}}^1(\emptyset) = \emptyset$ i.e. if and only if $\forall x \in \mathcal{A} \{x\}^- \neq \emptyset$. \square

Corollary 2 provides a condition for $COIN_{GR, GR^*}$ involving the existence of unattacked arguments in all resolutions of $CUT(\mathcal{G})$. While checking the existence of unattacked arguments is easy, using this condition would impose considering all the full resolutions of $CUT(\mathcal{G})$, whose enumeration would give rise to a combinatorial explosion. Next, we will first derive a simpler to check condition, concerning the case of argumentation frameworks consisting of a single strongly-connected component and then exploit this result in the general case. We recall that the *strongly-connected component* (SCC) decomposition of $\mathcal{G}(\mathcal{A}, \mathcal{R})$ partitions \mathcal{A} according to the equivalence classes induced by the relation $\rho(x, y)$ defined over $\mathcal{A} \times \mathcal{A}$ so that $\rho(x, y)$ holds if and only if $x = y$ or there are directed paths from x to y and from y to x in $\mathcal{G}(\mathcal{A}, \mathcal{R})$. We will denote the set of strongly connected components of \mathcal{G} as $SCCS(\mathcal{G})$. It is well-known that the graph obtained by considering strongly connected components as single nodes is acyclic. As a consequence, a partial order \prec over the SCC decomposition $\{\mathcal{A}_1, \dots, \mathcal{A}_k\}$ is defined as $(\mathcal{A}_i \prec \mathcal{A}_j) \Leftrightarrow (i \neq j)$ and $\exists x \in \mathcal{A}_i, y \in \mathcal{A}_j$ such that there is a directed path from x to y .

Exemplifying SCCs in Figs. 1 and 2, we have $\text{SCCS}(\mathcal{G}_1) = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4\} = \{\{a\}, \{b, c, d\}, \{e\}, \{f\}\}$, with $\mathcal{A}_1 \prec \mathcal{A}_2 \prec \mathcal{A}_3 \prec \mathcal{A}_4$ and $\text{SCCS}(\mathcal{G}_2) = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4\} = \{\{a, b, c\}, \{e\}, \{d\}, \{f, g, h\}\}$ with $\mathcal{A}_1 \prec \mathcal{A}_3 \prec \mathcal{A}_4$ and $\mathcal{A}_2 \prec \mathcal{A}_3 \prec \mathcal{A}_4$. Considering the restricted AFs obtained by suppressing the elements in the dashed boxes we have $\text{SCCS}(\text{CUT}(\mathcal{G}_1)) = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3\} = \{\{c, d\}, \{e\}, \{f\}\}$, with $\mathcal{A}_1 \prec \mathcal{A}_2 \prec \mathcal{A}_3$, and $\text{SCCS}(\text{CUT}(\mathcal{G}_2)) = \{\mathcal{A}_1, \mathcal{A}_2\} = \{\{a, b, c\}, \{f, g, h\}\}$, with no precedence relation between \mathcal{A}_1 and \mathcal{A}_2 .

The following lemma states that given a SCC T and any argument $x \in T$ it is possible to find a full resolution γ of \mathcal{G} which resolves all mutual attacks involving elements of T such that in \mathcal{G}_γ all elements of T , with the only possible exception of x , receive an attack from an element of T itself.

Lemma 5 *Given an AF $\mathcal{G}(\mathcal{A}, \mathcal{R})$ and a SCC $T \in \text{SCCS}(\mathcal{G})$, for any $x \in T$ there is a full resolution γ of $\mathcal{G} \downarrow_T$ such that in $\mathcal{G}_\gamma \forall y \in (T \setminus \{x\})$ attacks(T, y).*

Proof: For a generic $x \in T$, define inductively the following sequence of sets: $L_0 = \{x\}$, $L_{i+1} = L_i^+ \setminus (\bigcup_{j=0}^i L_j)$ for $i \geq 0$. Observe that for any $y \in T \exists i : y \in L_i$. In fact, $x \in L_0$ and for any $y \neq x$ there is a path from x to y , T being a SCC. Letting d be the minimal path length from x to y it is evident that $y \in L_d$. We can now build the full resolution γ as follows: for any mutual attack involving consecutive sets in the sequence insert in γ the attack coming from the set with higher index, namely for any $\{\langle y', y'' \rangle, \langle y'', y' \rangle\} \subseteq \mathcal{R}$ such that $y' \in L_i, y'' \in L_{i+1}$ for some i , let $\langle y'', y' \rangle \in \gamma$. Then, resolve arbitrarily any other mutual attack. It is evident that for any $y \neq x$ the path with minimal length from x to y within T is preserved in \mathcal{G}_γ , hence attacks(T, y). \square

We can now obtain the important result concerning an AF consisting of a single SCC anticipated above.

Lemma 6 *Given an AF $\mathcal{G}(\mathcal{A}, \mathcal{R})$ such that $|\text{SCCS}(\mathcal{G})| = 1$, the condition (i) for any full resolution γ of $\mathcal{G} \exists x$ such that in $\mathcal{G}_\gamma \{x\}^- = \emptyset$ is equivalent to the conjunction (ii) of the following three conditions:*

- $\forall x \in \mathcal{A}, \langle x, x \rangle \notin \mathcal{R}$;
- \mathcal{R} is symmetric, i.e. $\langle x, y \rangle \in \mathcal{R} \Leftrightarrow \langle y, x \rangle \in \mathcal{R}$;
- the undirected graph $\overline{\mathcal{G}}$ formed by replacing each (directed) pair $\{\langle x, y \rangle, \langle y, x \rangle\}$ with a single undirected edge $\{x, y\}$ is acyclic.

Proof: We first prove that (i) implies (ii) by showing that if any of the conditions (a-c) is violated then (i) is violated too. If a) does not hold we can apply Lemma 5 to x such that $\langle x, x \rangle \in \mathcal{R}$ and derive the existence of a full resolution γ where $\forall y \neq x \{y\}^- \neq \emptyset$ while $\{x\}^- \supseteq \{x\}$, thus denying (i). If either b) or c) is violated there is a cycle consisting of at least three distinct elements in \mathcal{G} . In fact, if \mathcal{R} is not symmetric, for any $\langle x, y \rangle \in \mathcal{R}$ such that $\langle y, x \rangle \notin \mathcal{R}$ there must be a path from y to x involving at least another distinct argument z , while if the undirected graph contains a cycle it involves necessarily at least three elements. Now it is easy to build a (possibly empty) resolution γ' resolving only the mutual attacks, if any, involving elements of the cycle and

preserving the existence of such a cycle in $\mathcal{G}_{\gamma'}$. Therefore $\mathcal{G}_{\gamma'}$ still consists of exactly one SCC. Consider now any argument x in the cycle: clearly $\{x\}^- \neq \emptyset$ in $\mathcal{G}_{\gamma'}$ and this condition will still hold in any full resolution of $\mathcal{G}_{\gamma'}$. But we can apply now Lemma 5 to x and derive the existence of a full resolution γ'' of $\mathcal{G}_{\gamma'}$ where $\forall y \neq x \{y\}^- \neq \emptyset$. Summing up we have obtained a full resolution $\gamma = \gamma' \cup \gamma''$ of \mathcal{G} such that $\forall x \{x\}^- \neq \emptyset$ in \mathcal{G}_γ .

Turning to the other direction of the proof, we will now show that the conjunction of a), b) and c) implies (i), by trying to build a full resolution γ such that $\forall x \in \mathcal{A} \{x\}^- \neq \emptyset$ in \mathcal{G}_γ and showing that this is impossible. Since c) holds, the undirected graph $\overline{\mathcal{G}}$ obtained from \mathcal{G} is a tree. Let r be the tree root and for any $y \neq r$ denote as $d(y)$ the length of the unique (simple) path from r to y . Let $m = \max_{y \in \mathcal{A} \setminus \{r\}} d(y)$: for any y such that $d(y) = m$ it is clearly the case that y is directly connected in $\overline{\mathcal{G}}$ with exactly one element z such that $d(z) = m - 1$. This entails that y can only attack or be attacked by z in \mathcal{G} , and, by b), actually both cases hold, i.e. $\{\langle y, z \rangle, \langle z, y \rangle\} \subseteq \mathcal{R}$. Then necessarily $\langle y, z \rangle \in \gamma$, otherwise y , not being self-defeating by c), would be unattacked in \mathcal{G}_γ . This entails that for any z such that $d(z) = m - 1$, z does not receive attacks in \mathcal{G}_γ from any argument y such that $d(y) = m$. But now we can iterate the same reasoning on any argument z such that $d(z) = m - 1$ showing that there is exactly one w such that $d(w) = m - 2$, $\{\langle z, w \rangle, \langle w, z \rangle\} \subseteq \mathcal{R}$ and necessarily $\langle z, w \rangle \in \gamma$. Iterating the same reasoning we reach the arguments x such that $d(x) = 1$ and $\{\langle x, r \rangle, \langle r, x \rangle\} \subseteq \mathcal{R}$. For any such argument x it must be the case that $\langle x, r \rangle \in \gamma$ (otherwise x would be unattacked in \mathcal{G}_γ) but then r is unattacked in \mathcal{G}_γ , showing that the construction of the desired γ is impossible. \square

Lemma 6 has provided three simple topological conditions which, on the basis of Corollary 2, allow to check the condition $\mathcal{E}_{GR}(\mathcal{G}) = \mathcal{E}_{GR^*}(\mathcal{G})$ (while avoiding the enumeration of full resolutions) when $|\text{SCCS}(\mathcal{G})| = 1$. To extend this result to a generic \mathcal{G} we need to focus our attention on the strongly connected components which are minimal with respect to \prec (i.e. do not receive attacks from other strongly connected components) and satisfy conditions (a-c) of Lemma 6.

Definition 3 *Given an AF $\mathcal{G}(\mathcal{A}, \mathcal{R})$, $S \in \text{SCCS}(\mathcal{G})$ is minimal relevant if S is a minimal element of \prec and $\mathcal{G} \downarrow_S$ satisfies conditions (a-c) stated in Lemma 6. The set of the minimal relevant SCCs of \mathcal{G} will be denoted as $\text{MR}(\mathcal{G})$.*

To exemplify, referring to Figs. 1 and 2, it is easy to see that $\text{MR}(\mathcal{G}_1) = \{\{a\}\}$, $\text{MR}(\mathcal{G}_2) = \{\{e\}\}$, $\text{MR}(\text{CUT}(\mathcal{G}_1)) = \{\{c, d\}\}$, and $\text{MR}(\text{CUT}(\mathcal{G}_2)) = \{\{f, g, h\}\}$.

The following theorem achieves the desired generalization by showing that verifying the coincidence between GR and GR^* for an AF \mathcal{G} is equivalent to checking whether $\text{CUT}(\mathcal{G})$ has some minimal relevant component.

Theorem 1 *Given an AF $\mathcal{G}(\mathcal{A}, \mathcal{R})$, $\mathcal{E}_{GR}(\mathcal{G}) = \mathcal{E}_{GR^*}(\mathcal{G}) \Leftrightarrow \text{MR}(\text{CUT}(\mathcal{G})) = \emptyset$.*

Proof: Suppose first $\mathcal{E}_{GR}(\mathcal{G}) = \mathcal{E}_{GR^*}(\mathcal{G})$. By Corollary 2, there is a full resolution γ of $\text{CUT}(\mathcal{G})$ such that in $\text{CUT}(\mathcal{G})_\gamma$ every argument has at least one attacker. Let S be any SCC

of $\text{CUT}(\mathcal{G})$ minimal with respect to \prec . Clearly there is a full resolution γ_S of $\text{CUT}(\mathcal{G})\downarrow_S$ such that every element of S has at least one attacker in $(\text{CUT}(\mathcal{G})\downarrow_S)_{\gamma_S}$ and therefore $\text{CUT}(\mathcal{G})\downarrow_S$ does not satisfy conditions (a-c) of Lemma 6. It follows $\text{MR}(\text{CUT}(\mathcal{G})) = \emptyset$.

Turning to the other direction of the proof, by Corollary 2 it is sufficient to show that there is a full resolution γ of $\text{CUT}(\mathcal{G})$ such that in $\text{CUT}(\mathcal{G})_\gamma$ every argument has at least one attacker. To build such a γ , consider first any SCC of $\text{CUT}(\mathcal{G})$ minimal with respect to \prec . Given the hypothesis $\text{MR}(\text{CUT}(\mathcal{G})) = \emptyset$, by Lemma 6 there is a full resolution γ_S of $\text{CUT}(\mathcal{G})\downarrow_S$ such that every element of S has at least one attacker in $(\text{CUT}(\mathcal{G})\downarrow_S)_{\gamma_S}$. Turning now to the SCCs of $\text{CUT}(\mathcal{G})$ which are not minimal with respect to \prec , we can proceed following the (partial) order induced by \prec . In fact, for any such SCC S' we can assume inductively that there is a full resolution γ such that for every SCC $S \prec S'$ every element of S has at least one attacker in $\text{CUT}(\mathcal{G})_\gamma$ and we need to show that the same holds also for S' . Note first that there must be an element x of S' which receives at least an attack from an element y of a SCC S such that $S \prec S'$ and that any such attack must be non mutual (otherwise S and S' would not be distinct SCCs), hence for any resolution γ $\langle y, x \rangle \notin \gamma$. Now, by Lemma 5 we can define a resolution $\gamma' \subseteq (S' \times S')$ such that any argument $z \neq x$ in S' has at least an attacker, while, as shown above, x has at least an attacker in any resolution. Summing up, we have shown a procedure to incrementally build a full resolution γ of $\text{CUT}(\mathcal{G})$ such that any element of $\text{CUT}(\mathcal{G})$ has at least an attacker in $\text{CUT}(\mathcal{G})_\gamma$, as desired. \square

Polynomial complexity results for COIN_{GR,GR^*} and NE_{GR^*} follow directly from Theorem 1.

Corollary 3 $\text{COIN}_{GR,GR^*} \in \text{P}$.

Proof: By Theorem 1, to check $\mathcal{E}_{GR}(\mathcal{G}) = \mathcal{E}_{GR^*}(\mathcal{G})$ do the following steps: i) compute $GE(\mathcal{G})$; ii) compute $\text{CUT}(\mathcal{G})$; iii) compute the SCC decomposition of $\text{CUT}(\mathcal{G})$; iv) identify those SCCs of $\text{CUT}(\mathcal{G})$ which are minimal with respect to \prec ; v) on each of them check conditions (a-c) of Lemma 6. Each of these steps is known (or easily seen) to belong to P . \square

Corollary 4 $\text{NE}_{GR^*} \in \text{P}$.

Proof: By Corollary 1, $\mathcal{E}_{GR^*}(\mathcal{G}) = \{\emptyset\} \Leftrightarrow GE(\mathcal{G}) = \emptyset \wedge \mathcal{E}_{GR}(\mathcal{G}) = \mathcal{E}_{GR^*}(\mathcal{G})$, thus NE_{GR^*} reduces to checking first NE_{GR} , which is known to belong to P , and then (possibly) COIN_{GR,GR^*} , which belongs to P by Corollary 3. \square

We now turn to the problem VER_{GR^*} . Preliminarily, we have to identify some quite technical but useful properties of GR and GR^* in relation with minimal relevant components.

Lemma 7 *Given an AF $\mathcal{G}(\mathcal{A}, \mathcal{R})$ and $S \in \text{MR}(\mathcal{G})$, it holds that (i) for any full resolution γ of $\mathcal{G}\downarrow_S$ $GE(\mathcal{G}_\gamma\downarrow_S)$ is a stable extension of $\mathcal{G}_\gamma\downarrow_S$ and (ii) $\mathcal{E}_{GR^*}(\mathcal{G}\downarrow_S) = \{T \mid T \text{ is a stable extension of } \mathcal{G}\downarrow_S\}$.*

Proof: Recall first that a stable extension of an AF $\mathcal{G}(\mathcal{A}, \mathcal{R})$ is a conflict-free set $T \subseteq \mathcal{A}$ such that $\forall x \in (\mathcal{A} \setminus T)$ attacks(T, x). As for (ii), we show for any $S \in \text{MR}(\mathcal{G})$ that $\{GE(\mathcal{G}_\gamma\downarrow_S) \mid \gamma \text{ is a full resolution of } \mathcal{G}\downarrow_S\} = \{T \mid$

$T \text{ is a stable extension of } \mathcal{G}\downarrow_S\}$, and, since no stable extension can be a proper subset of another one, this set turns out to be equal to $\mathcal{E}_{GR^*}(\mathcal{G}\downarrow_S)$. Condition (i) will arise as an intermediate result.

To show $\{GE(\mathcal{G}_\gamma\downarrow_S) \mid \gamma \text{ is a full resolution of } \mathcal{G}\downarrow_S\} \subseteq \{T \mid T \text{ is a stable extension of } \mathcal{G}\downarrow_S\}$ we observe first that, by the definition of $\text{MR}(\mathcal{G})$, $\mathcal{G}_\gamma\downarrow_S$ is acyclic. In fact, it does not contain self-defeating arguments, any cycle of length 2 in $\mathcal{G}\downarrow_S$ is resolved by γ and no cycles of length > 2 can be present. It is well-known [Dung, 1995] that in an acyclic argumentation framework the grounded extension is also a stable extension, thus $GE(\mathcal{G}_\gamma\downarrow_S)$ is a stable extension of $\mathcal{G}_\gamma\downarrow_S$ (we have thus proved (i)) and we can observe further that $GE(\mathcal{G}_\gamma\downarrow_S)$ is also a stable extension of $\mathcal{G}\downarrow_S$ since in $\mathcal{G}\downarrow_S$ it clearly preserves both the properties of being conflict-free and of attacking all other arguments.

To show that $\{T \mid T \text{ is a stable extension of } \mathcal{G}\downarrow_S\} \subseteq \{GE(\mathcal{G}_\gamma\downarrow_S) \mid \gamma \text{ is a full resolution of } \mathcal{G}\downarrow_S\}$, for any stable extension T of $\mathcal{G}\downarrow_S$ we have to build a full resolution γ of $\mathcal{G}\downarrow_S$ such that $T = GE(\mathcal{G}_\gamma\downarrow_S)$. To obtain such a γ , note that, by the symmetry of $\mathcal{G}\downarrow_S$, for any x in T either x is unattacked or is involved in mutual attacks with some other elements y of S and we can include in γ all the pairs of the form $\langle y, x \rangle$. It turns out that any element of T is unattacked in $\mathcal{G}_\gamma\downarrow_S$, and, T being a stable extension, that any element $y \notin T$ is attacked by T . It follows that $T = \mathcal{F}_{\mathcal{G}_\gamma\downarrow_S}^1(\emptyset) = \mathcal{F}_{\mathcal{G}_\gamma\downarrow_S}^i(\emptyset)$ for any $i \geq 1$ and hence $T = GE(\mathcal{G}_\gamma\downarrow_S)$. \square

To proceed we have now to recall the property of directionality of argumentation semantics, which is known to be satisfied by both GR [Baroni and Giacomin, 2007] and GR^* [Baroni and Giacomin, 2008]. It corresponds to the intuitive requirement that an argument y may affect another argument x only if there is a directed path from y to x . This can be formalized by referring to sets of arguments not receiving attacks from outside.

Definition 4 *Given an AF $\mathcal{G}(\mathcal{A}, \mathcal{R})$, a non-empty set $S \subseteq \mathcal{A}$ is unattacked if and only if $\nexists x \in (\mathcal{A} \setminus S) : \text{attacks}(x, S)$. The set of unattacked sets of \mathcal{G} is denoted as $\mathcal{US}(\mathcal{G})$.*

Definition 5 *A semantics \mathcal{S} satisfies the directionality criterion if and only if for any AF $\mathcal{G}(\mathcal{A}, \mathcal{R})$, $\forall T \in \mathcal{US}(\mathcal{G})$, $\mathcal{AE}_{\mathcal{S}}(\mathcal{G}, T) = \mathcal{E}_{\mathcal{S}}(\mathcal{G}\downarrow_T)$, where $\mathcal{AE}_{\mathcal{S}}(\mathcal{G}, T) \triangleq \{(E \cap T) \mid E \in \mathcal{E}_{\mathcal{S}}(\mathcal{G})\} \subseteq 2^T$.*

In words, the intersection of any extension prescribed by \mathcal{S} for \mathcal{G} with an unattacked set T is equal to one of the extensions prescribed by \mathcal{S} for the restriction of \mathcal{G} to T , and vice versa. The following theorem provides a characterization of the extensions of GR^* in terms of three (still quite technical) conditions. We omit its lengthy proof (which uses Lemma 7 and directionality of GR and GR^*) for space reasons.

Theorem 2 *Given an argumentation framework $\mathcal{G}(\mathcal{A}, \mathcal{R})$ and letting $\Pi_{\mathcal{G}} = \bigcup_{V \in \text{MR}(\text{CUT}(\mathcal{G}))} V$, $U \in \mathcal{E}_{GR^*}(\mathcal{G})$ if and only if the following conditions hold:*

- $U \cap \alpha(GE(\mathcal{G})) = GE(\mathcal{G})$
- $\forall M \in \text{MR}(\text{CUT}(\mathcal{G})) \text{ st}_{\text{CUT}(\mathcal{G})}(U, M)$
- $(U \cap \Pi_{\mathcal{G}}^C) \in \mathcal{E}_{GR^*}(\text{CUT}(\mathcal{G})\downarrow_{\Pi_{\mathcal{G}}^C \setminus (U \cap \Pi_{\mathcal{G}})^+})$.

Algorithm 1 Verifying that $U \in \mathcal{E}_{GR^*}(\mathcal{G}(\mathcal{A}, \mathcal{R}))$

```
1: procedure  $GR^*\text{-VER}(\mathcal{G}(\mathcal{A}, \mathcal{R}), U)$  returns boolean
2:  $S := GE(\mathcal{G})$ 
3: if  $(U \cap \alpha(S) \neq S)$  then
4:   return false
5: end if
6:  $W := \Pi_{\mathcal{G}}$ 
7:  $T := U \setminus S$ 
8: if  $W = \emptyset$  and  $T = \emptyset$  then
9:   return true
10: end if
11: if  $W = \emptyset$  and  $T \neq \emptyset$  then
12:   return false
13: end if
14: if  $\neg st_{CUT(\varphi)}(T, W)$  then
15:   return false
16: else
17:   return  $GR^*\text{-VER}(CUT(\mathcal{G}) \downarrow_{W^C \setminus (T \cap W)^+}, (T \cap W^C))$ 
18: end if
19: end
```

Theorem 3 $VER_{GR^*} \in P$.

Proof: The proof refers to the recursive Algorithm 1 which is easily seen to correspond to checking whether conditions (a-c) of Theorem 2 hold, while the correctness in the case $W = \emptyset$ comes from Theorem 1. First note that every step of Algorithm 1 is in P. In particular, it is well-known that computing $GE(\mathcal{G})$ is in P, as it clearly is also computing $\alpha(S)$. Computing $\Pi_{\mathcal{G}}$ basically amounts to compute the minimal relevant components of \mathcal{G} , a task we have already commented to be in P in Corollary 3. Verifying whether a set is stable in another one and identifying the arguments attacked by a set is linear in the number of attack relations, while all other operations (e.g. those involved in computing $CUT(\mathcal{G})$, $CUT(\mathcal{G}) \downarrow_{W^C \setminus (T \cap W)^+}$ and $(T \cap W^C)$) only require basic set manipulations. It remains to be seen that the recursion is well-founded and terminates after a polynomial number of calls. To this purpose it is sufficient to observe that at each recursive call an argumentation framework with a strictly lesser number of arguments is considered, since at least the elements of W (which is not empty, otherwise the procedure would terminate) are suppressed. Moreover the procedure clearly terminates without further recursive calls if invoked on $\mathcal{G}(\emptyset, \emptyset)$. It follows that the procedure terminates after a number of calls which is linear in the number of arguments. \square

Let us exemplify an execution of Algorithm 1 with arguments $\mathcal{G} = \mathcal{G}_1$ (Fig. 1) and $U = \{a, d, f\}$. We get $S = \{a\}$, at line 2 and it follows $U \cap \alpha(S) = \{a, d, f\} \cap \{a, b\} = \{a\} = S$. Then the condition of the **if** statement of line 3 is false and we obtain $W = \{\{c, d\}\}$ at line 6 and $T = \{d, f\}$ at line 7. Since $W \neq \emptyset$ we skip the following **if** statements and since T is stable in W the condition at line 14 is false and we enter the else branch at line 17. Here we note that $W^C = \{e, f\}$ and $(T \cap W)^+ = \{c, e\}$ yielding arguments $\mathcal{G}(\{f\}, \emptyset)$ and $\{f\}$ for the recursive invocation of the procedure. Then we have $S = \{f\}$ at line 3, $W = \emptyset$ at line 6, $T = \emptyset$ at line 7, and the procedure returns true at line 9.

In fact, $\{a, d, f\} \in \mathcal{E}_{GR^*}(\mathcal{G}_1)$. Consider instead an execution of Algorithm 1 with arguments $\mathcal{G} = \mathcal{G}_2$ (Fig. 2) and $U = \{a, e, g\}$. We get $S = \{e\}$ at line 2, and, skipping easy observations, we obtain $W = \{\{f, g, h\}\}$ at line 6 and we are led to a recursive invocation with arguments $\mathcal{G}_2 \downarrow_{\{a, b, c\}}$ and $\{a\}$. Now we obtain $S = \emptyset$ at line 3, $W = \emptyset$ at line 6 and $T = \{a\}$ at line 7. Then the condition of the **if** statement of line 11 is satisfied and the procedure returns false. In fact, $\{a, e, g\} \notin \mathcal{E}_{GR^*}(\mathcal{G}_2)$.

By the way, it is not difficult to see that $\mathcal{E}_{GR^*}(\mathcal{G}_1) = \{\{a, c, f\}, \{a, d, f\}\}$ and $\mathcal{E}_{GR^*}(\mathcal{G}_2) = \{\{e, g\}, \{e, f, h\}\}$.

3.2 Intractable decision problems

Turning to the credulous and skeptical acceptance problems, in contrast to the polynomial time methods identified in subsection 3.1 we have:

Theorem 4

- CA_{GR^*} is NP-complete.
- SA_{GR^*} is coNP-complete.

Proof: For part (a), that $CA_{GR^*} \in NP$ follows by observing that any instance $\langle \mathcal{G}(\mathcal{A}, \mathcal{R}), x \rangle$ can be decided by checking $\exists T \subseteq \mathcal{A} : (x \in T) \wedge VER_{GR^*}(\mathcal{G}, T)$. By virtue of Theorem 3 this yields an NP algorithm.

For NP-hardness we use a reduction from 3-SAT. Given an instance $\varphi(Z_n) = C_1 \wedge C_2 \wedge \dots \wedge C_m$ of 3-SAT form the instance $\langle \mathcal{G}(\mathcal{A}_\varphi, \mathcal{R}_\varphi), \varphi \rangle$ of CA_{GR^*} in which,

$$\begin{aligned} \mathcal{A}_\varphi &= \{\varphi\} \cup \{C_j : 1 \leq j \leq m\} \cup \{z_i, \neg z_i : 1 \leq i \leq n\} \\ \mathcal{R}_\varphi &= \{ \langle C_j, \varphi \rangle : 1 \leq j \leq m \} \cup \\ &\quad \{ \langle z_i, C_j \rangle : z_i \text{ occurs in } C_j \} \cup \\ &\quad \{ \langle \neg z_i, C_j \rangle : \neg z_i \text{ occurs in } C_j \} \cup \\ &\quad \{ \langle z_i, \neg z_i \rangle, \langle \neg z_i, z_i \rangle : 1 \leq i \leq n \} \end{aligned}$$

We claim that there is some $T \in \mathcal{E}_{GR^*}(\mathcal{G})$ for which $\varphi \in T$ if and only if there is a satisfying instantiation of $\varphi(Z_n)$.

It is easily seen that $\mathcal{E}_{GR}(\mathcal{G}) = \{\emptyset\}$. Furthermore, noting that $M_{\mathcal{G}}$ contains exactly the set of pairs $\{\langle z_i, \neg z_i \rangle, \langle \neg z_i, z_i \rangle : 1 \leq i \leq n\}$ every full resolution of these yields a distinct set in $\mathcal{E}_{GR^*}(\mathcal{G})$.

Suppose first that $\underline{\alpha} = \langle \alpha_1, \alpha_2, \dots, \alpha_n \rangle$ describes a satisfying assignment for $\varphi(Z_n)$ and consider the full resolution $\gamma(\underline{\alpha})$ given by

$$\langle z_i, \neg z_i \rangle \in \gamma(\underline{\alpha}) \Leftrightarrow \alpha_i = \perp$$

The grounded extension of the AF $\mathcal{G}_{\gamma(\underline{\alpha})} = \langle \mathcal{A}_\varphi, \mathcal{R}_\varphi \setminus \gamma(\underline{\alpha}) \rangle$ contains exactly the arguments $\{z_i : \alpha_i = \top\} \cup \{\neg z_i : \alpha_i = \perp\} \cup \{\varphi\}$: each of the literal arguments (i.e. the $y_i \in \{z_i, \neg z_i\}$ selected) has $\{y_i\}^- = \emptyset$. Furthermore, since $\underline{\alpha}$ satisfies φ , each clause argument C_j attacking φ must be attacked by at least one of these literal arguments. It remains only to note that the resulting subset is minimal among the grounded extensions resulting from full resolutions of \mathcal{G} .

On the other hand suppose that $\gamma \subset M_{\mathcal{G}}$ defines a full resolution for which φ is in the grounded extension, T , of $\mathcal{G}_\gamma = \langle \mathcal{A}_\varphi, \mathcal{R}_\varphi \setminus \gamma \rangle$. From $\varphi \in T$, it follows that $C_j \notin T$ for any $1 \leq j \leq m$, and thus (at least) one literal $y_j \in \{z_j, \neg z_j\}$ among the literals defining C_j must belong to T . It follows

that for each clause, $C_j = y_{j,1} \vee y_{j,2} \vee y_{j,3}$, γ must contain at least one of the attacks $\langle \neg y_{j,k}, y_{j,k} \rangle$ in order for $y_{j,k} \in T$ to hold. Now defining the instantiation $\langle \alpha_1^\gamma, \alpha_2^\gamma, \dots, \alpha_n^\gamma \rangle$ of Z_n via $a_i^\gamma = \top \Leftrightarrow \langle \neg z_j, z_j \rangle \in \gamma$ yields a satisfying assignment of $\varphi(Z_n)$ as required.

For part (b), to decide $\text{SA}_{GR^*}(\mathcal{G}(\mathcal{A}, \mathcal{R}), x)$ by a CONP method, simply involves verifying for every full resolution γ of \mathcal{G} , with T_γ the grounded extension of \mathcal{G}_γ , that

$$T_\gamma \in \mathcal{E}_{GR^*}(\mathcal{G}) \Rightarrow x \in T_\gamma$$

Again, by virtue of Theorem 3 the required test (for $T_\gamma \in \mathcal{E}_{GR^*}(\mathcal{G})$) can be performed in polynomial time.

For CONP-hardness, we use a similar construction applied to deciding *unsatisfiability*: this involves the AF of (a) augmented by a single new argument ψ whose sole attacker is φ . We omit the straightforward proof that ψ is skeptically accepted if and only if φ is unsatisfiable. \square

4 Conclusions and Further Work

We have investigated the computational properties of GR^* with reference to a set of decision problems for abstract argumentation semantics: it turns out that some of them (VER_{GR^*} , COIN_{GR,GR^*} , NE_{GR^*}) are tractable while others (CA_{GR^*} , SA_{GR^*}) are in general not.

Let us discuss this result with respect to complexity properties of other abstract argumentation semantics, focusing first on Dung's traditional grounded (GR), stable (ST) and preferred (PR) semantics. The (unique) grounded extension is known to be computable with a polynomial algorithm, hence all the decision problems considered in this paper are known to be tractable for GR (we have exploited this property in Section 3). On the other hand, the same decision problems are known to be generally intractable for both ST and PR (see in particular [Dimopoulos and Torres, 1996; Dunne and Bench-Capon, 2002; Dunne and Wooldridge, 2009]) with the only exception of $\text{VER}_{ST} \in \text{P}$ (also this property has been exploited in Section 3). In particular CA_{PR} , NE_{PR} , and CA_{ST} are NP-complete, VER_{PR} is CONP-complete, SA_{PR} is Π_2^{P} -complete, and SA_{ST} is D^{P} -complete. We can state therefore that GR^* has better complexity properties than the traditional multiple-status semantics ST and PR . Complexity properties of recently proposed semantics, e.g. ideal, semi-stable or prudent, have not been fully analyzed yet but preliminary non-tractability results exist [Dunne and Wooldridge, 2009]. Actually, as to our knowledge, no other non-trivial multiple-status semantics in the literature has been shown, up to now, to admit polynomial time decision processes (in the general case) for any of the standard decision problems considered here. In this wider perspective one of the main contributions of this paper consists in showing that, differently from what previous literature results could have suggested, argumentation-based reasoning with multiple extensions is not bound to be computationally intractable.

This result has been obtained combining a variety of techniques into an articulated analysis some of whose steps may be interesting on their own. In particular Lemmata 1 and 2 reveal properties of the traditional grounded semantics, while

the role played by SCC decomposition confirms the importance of this topological notion in abstract argumentation [Baroni *et al.*, 2005].

Focusing on the specific scope of the paper, GR^* has been shown to combine its (already known) merits related to principled conceptual requirements with computational advantages which may turn out to be significant for practical applications. This aspect deserves further investigation: defining criteria for matching an abstract argumentation semantics with the requirements of a specific application context is a largely unexplored issue. It is however worth remarking that the results provided in the paper (in particular the procedure summarized in Corollary 3 and Algorithm 1) lend themselves to a straightforward implementation. As the availability of implemented argumentation systems is growing in recent years, integrating GR^* into one of them, applying it in practical test cases and comparing results with other semantics appear to be feasible and very interesting future activities.

On the theoretical side, among further research directions we mention the identification of families of argumentation frameworks where also CA_{GR^*} and SA_{GR^*} are tractable. A preliminary investigation suggests that this is the case for the family of *bipartite* argumentation frameworks.

Acknowledgments

The authors thank the referees for their valuable comments.

References

- [Baroni and Giacomin, 2007] P. Baroni and M. Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence*, 171(10/15):675–700, 2007.
- [Baroni and Giacomin, 2008] P. Baroni and M. Giacomin. Resolution-based argumentation semantics. In P. Besnard, S. Doutre, and A. Hunter, editors, *Proc. 2nd Int. Conf. on Computational Models of Argument (COMMA 2008)*, pages 25–36, Toulouse, F, 2008. IOS Press.
- [Baroni *et al.*, 2005] P. Baroni, M. Giacomin, and G. Guida. SCC-recursiveness: a general schema for argumentation semantics. *Artificial Intelligence*, 168(1-2):165–210, 2005.
- [Dimopoulos and Torres, 1996] Y. Dimopoulos and A. Torres. Graph theoretical structures in logic programs and default theories. *Theoretical Computer Science*, 170(1-2):209–244, 1996.
- [Dung, 1995] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and N -person games. *Artificial Intelligence*, 77(2):321–357, 1995.
- [Dunne and Bench-Capon, 2002] P. E. Dunne and T. J. M. Bench-Capon. Coherence in finite argument systems. *Artificial Intelligence*, 141(1):187–203, 2002.
- [Dunne and Wooldridge, 2009] P. E. Dunne and M. Wooldridge. Complexity of abstract argumentation. In I. Rahwan and G. R. Simari, editors, *Argumentation in Artificial Intelligence*. Springer, Berlin, 2009.