Solving Semantic Problems with Odd-Length Cycles in Argumentation

Pietro Baroni and Massimiliano Giacomin

Università di Brescia
Dipartimento di Elettronica per l'Automazione
Via Branze 38, I-25123 Brescia, Italy
{baroni, giacomin}@ing.unibs.it

Abstract. In the context of Dung's abstract framework for argumentation, two main semantics have been considered to assign a defeat status to arguments: the grounded semantics and the preferred semantics. While the two semantics agree in most situations, there are cases where the preferred semantics appears to be more powerful. However, we notice that the preferred semantics gives rise to counterintuitive results in some other cases, related to the presence of odd-length cycles in the attack relation between arguments. To solve these problems, we propose a new semantics which preserves the desirable properties of the preferred semantics, while correctly dealing with odd-length cycles. We check the behavior of the proposed semantics in a number of examples and discuss its relationships with both grounded and preferred semantics.

1 Introduction

Argumentation theory is a framework for practical and uncertain reasoning which has received a great deal of attention in several application areas, such as the realization of intelligent autonomous agents [1], automated negotiation in multi-agent systems [2] and defeasible reasoning [3]. In a nutshell, commonsense reasoning dealing with incomplete and uncertain information is modeled as the process of constructing and comparing arguments for propositions. The construction of arguments proceeds, from a given set of premises, by chaining rules of inference which may represent just provisional reasons for their conclusions. Due to the uncertainty affecting both premises and rules of inference, it may well be the case that different arguments support contradictory conclusions, therefore the core problem consists in computing the *defeat status* of the arguments, namely in determining which ones of them emerge undefeated from conflict: their conclusions are the most credible ones and are considered as justified, while other arguments, being defeated in the conflict, are rejected.

In order to analyze and compare different kinds of *semantics* underlying the defeat status computation, Dung [4] has proposed an abstract framework where arguments are simply conceived as the elements of a set, whose origin is not specified, and the interaction between them is represented by a binary relation of *attack*: this way, the current set of arguments can be represented by means of a directed graph, called *defeat graph* in [1]. Thus, an argumentation semantics can be introduced in a declarative way by defining what arguments are justified within a generic defeat graph. As pointed out in [5], this definition can follow two alternative approaches, namely a unique-status approach or a multiple-status approach. In the first approach, the defeat status of the arguments is defined in such a way that there is always exactly one possible way to assign them a status. This approach is adopted e.g. in the argumentation system introduced in [1], and is represented by the *grounded semantics* in Dung's framework. On the other hand, in a multiple-status approach several *extensions* are identified. Roughly, an extension is a set of arguments which do not conflict among them and which attack their attackers. An argument is considered as justified if it belongs to all extensions. This is the approach adopted e.g. in [6, 7, 3], and in Dung's framework is captured by the *preferred semantics*.

It has been proved in [4] that the preferred semantics "agrees" with the grounded semantics in those arguments that the latter considers as definitely justified or rejected. On the other hand, the preferred semantics appears to be more powerful with respect to the grounded semantics, in that it is sometimes able to discriminate some of the arguments that are left undecided by the grounded semantics [8].

After recalling concepts and definitions about argumentation semantics in Sect. 2, we point out in Sect. 3 that the preferred semantics improperly deals with odd-length cycles in the defeat graph and we identify some examples where this limitation gives rise to counter-intuitive defeat status assignments. To solve these problems, we propose in Sect. 4 a new semantics which preserves the desirable properties of the preferred semantics, while correctly dealing with odd-length cycles. In Sect. 5, the relationships with grounded and preferred semantics are investigated. Finally, Sect. 6 concludes the paper.

2 The Grounded and Preferred Semantics

In the abstract framework proposed by Dung [4], the primitive notion is that of *argumentation framework*:

Definition 1. *An argumentation framework is a pair* $\mathrm{AF} = < \mathcal{A}, \rightarrow >$*, where* $\mathcal{A}$ *is a set of arguments and* $\rightarrow$ *is a binary relation of 'attack' between them.*

It should be noticed that this definition is generic with respect to the interpretation of $\mathcal{A}$, which is not specified. In any case, we assume $\mathcal{A}$ to be finite, as it is necessarily the case when considering a *real* reasoner.

In the following, nodes that attack a given $\alpha \in \mathcal{A}$ are called *defeaters* of α, and form a set denoted as $\mathrm{parents}(\alpha)$. If $\mathrm{parents}(\alpha) = \emptyset$, then α is called an *initial* node. Following Pollock [1] we define the grounded semantics inductively (an alternative fixed-point definition is given and shown to be equivalent in [4]):

Definition 2. *Given an argumentation framework* AF $=<\mathcal{A}, \rightarrow>$, *we define for all* $i \geq 0$ *the sets* $\mathcal{A}_i \subseteq \mathcal{A}$ *as follows:*

$$\mathcal{A}_i = \begin{cases} \mathcal{A} & \textit{if } i = 0 \\ \{\alpha \in \mathcal{A} \mid \nexists \beta \in \mathcal{A}_{i-1} : \beta \rightarrow \alpha\} & \textit{if } i > 0 \end{cases}$$

Definition 3. *Given an argumentation framework* AF $=<\mathcal{A}, \rightarrow>$, *the set of* undefeated, defeated *and* provisionally defeated *arguments are respectively defined as follows:*

- $\mathrm{U}_{\mathcal{G}}(\mathrm{AF}) = \{\alpha \in \mathcal{A} \mid \exists m : \forall i \geq m \ \alpha \in \mathcal{A}_i\}$
- $\mathrm{D}_{\mathcal{G}}(\mathrm{AF}) = \{\alpha \in \mathcal{A} \mid \exists m : \forall i \geq m \ \alpha \notin \mathcal{A}_i\}$
- $\mathrm{P}_{\mathcal{G}}(\mathrm{AF}) = \{\alpha \in \mathcal{A} \mid \forall m \ \exists i \geq m : \alpha \in \mathcal{A}_i \wedge \exists j \geq m : \alpha \notin \mathcal{A}_j\}$

The idea is that an undefeated argument should be believed given the current set of arguments $\mathcal{A}$, a defeated argument should not be believed, while a provisionally defeated argument is controversial, thus it should not be believed but it should retain the potential to prevent other arguments to be justified. This is shown in the following examples.

Example 1. With reference to the argumentation framework AF_1 shown in Fig. 1, it is easy to see that α belongs to $\mathcal{A}_i$ for all $i \geq 0$, since it has no defeaters, therefore α is undefeated. As a consequence, $\forall i \geq 1 \ \beta \notin \mathcal{A}_i$, therefore β is defeated. This entails in turn that $\forall i \geq 2, \ \gamma \in \mathcal{A}_i$, therefore γ is undefeated.

Example 2. With reference to the argumentation framework AF_2 of Fig. 1, it is easy to see that, for all $k \geq 0$, both α and β belong to $\mathcal{A}_{2k}$ but don't belong to $\mathcal{A}_{2k+1}$, therefore they are provisionally defeated. This alternation of levels is inherited by γ, which turns out to be provisionally defeated as well.

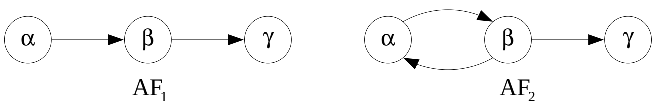

Fig. 1. Two different chains

In the context of the preferred semantics, the notion of 'defence' is introduced by the following definitions:

Definition 4. *Given an argumentation framework* AF $=<\mathcal{A}, \rightarrow>$, *a set* $S \subseteq \mathcal{A}$ *is* conflict-free *if and only if* $\nexists \alpha, \beta \in S$ *such that* $\alpha \rightarrow \beta$.

Definition 5. *Given an argumentation framework* AF $=<\mathcal{A}, \rightarrow>$, *a set* $S \subseteq \mathcal{A}$ *is* admissible *if and only if* S *is conflict-free and* $\forall \alpha \in S$, *if* $\exists \beta \in \mathcal{A}$ *such that* $\beta \rightarrow \alpha$ *then* $\exists \gamma \in S$ *such that* $\gamma \rightarrow \beta$.

Accordingly, a preferred extension is defined as a maximal set which is able to defend all its elements:

Definition 6. *Given an argumentation framework* AF $=<\mathcal{A}, \rightarrow>$, *a* preferred extension *of* AF *is a maximal (with respect to set inclusion) admissible set* $S \subseteq \mathcal{A}$. *The set of preferred extensions will be denoted as* $\mathcal{FP}(\mathrm{AF})$.

As shown in [4], $\mathcal{FP}(\mathrm{AF})$ is never empty, though there are cases where $\mathcal{FP}(\mathrm{AF}) = \{\emptyset\}$, e.g. $\mathcal{FP}(\mathrm{AF}_5)$ in Example 4 below. Also in this semantics three sets of arguments are identified: undefeated arguments belong to all extensions, defeated arguments to none, while provisionally defeated only to some of them.

Definition 7. *Given an argumentation framework* AF $=<\mathcal{A}, \rightarrow>$, *we define the following three sets, forming a partition of* $\mathcal{A}$:

- $\mathrm{U}_{\mathcal{P}}(\mathrm{AF}) = \{\alpha \in \mathcal{A} \mid \forall P \in \mathcal{FP}(\mathrm{AF})\ \alpha \in P\}$
- $\mathrm{D}_{\mathcal{P}}(\mathrm{AF}) = \{\alpha \in \mathcal{A} \mid \forall P \in \mathcal{FP}(\mathrm{AF})\ \alpha \notin P\}$
- $\mathrm{P}_{\mathcal{P}}(\mathrm{AF}) = \{\alpha \in \mathcal{A} \mid \exists P_1, P_2 \in \mathcal{FP}(\mathrm{AF}) : \alpha \in P_1 \wedge \alpha \notin P_2\}$

Turning to Example 1 and Example 2, it is easy to see that the preferred semantics gives the same outcome as the grounded semantics, since we have that $\mathcal{FP}(\mathrm{AF}_1) = \{\{\alpha, \gamma\}\}$, while $\mathcal{FP}(\mathrm{AF}_2) = \{\{\alpha, \gamma\}, \{\beta\}\}$. The relation between grounded and preferred semantics has been analyzed in [4]: in a nutshell, the grounded semantics is more cautious than preferred semantics, since for all argumentation frameworks it turns out that all the arguments undefeated and defeated according to the grounded semantics have the same status in the preferred semantics. On the other hand, there may be arguments provisionally defeated in the grounded semantics which are defeated or undefeated in the preferred semantics, as in the case of *floating arguments* [5] exemplified below.

Example 3. With reference to the argumentation framework AF_3 shown in Fig. 2, it is easy to see that, according to the grounded semantics, all arguments are provisionally defeated. On the other hand, it turns out that $\mathcal{FP}(\mathrm{AF}_3) = \{\{\alpha, \delta\}, \{\beta, \delta\}\}$, therefore we have that $\mathrm{P}_{\mathcal{P}}(\mathrm{AF}_3) = \{\alpha, \beta\}$, $\mathrm{D}_{\mathcal{P}}(\mathrm{AF}_3) = \{\gamma\}$ and $\mathrm{U}_{\mathcal{P}}(\mathrm{AF}_3) = \{\delta\}$.

In the example above, every preferred extension includes an argument which attacks γ, while no argument attacking γ belongs to all extensions: this is a case of 'floating defeat', as it has been called in [8], which determines in turn the 'floating acceptance' of δ. The inability to discriminate floating arguments is not a specific disadvantage of grounded semantics, since Schlechta has proved in [9] that it affects any possible single-status approach.

3 Odd-length Cycles: a Problem in Preferred Semantics

According to the definitions presented in the previous section, if the nodes of a defeat graph are arranged in a cycle of attack relationships, then they are not justified (i.e. they are provisionally defeated) according to both the grounded

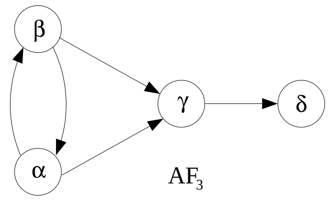

Fig. 2. Argumentation framework with a floating argument

and preferred semantics. This seems to be the intuitively right result, since all arguments in a cycle should be treated equally for obvious symmetry reasons and considering them all justified would yield a contradiction. However this result is obtained in rather different ways in the two semantics. In the context of the grounded semantics, all arguments forming a cycle simply turn out to belong to $\mathcal{A}_i$ if i is even and not to belong to $\mathcal{A}_i$ if i is odd (see Definitions 2 and 3).

On the other hand, the preferred semantics features a sort of asymmetry, since it treats odd-length cycles differently from the even-length ones.

Example 4. Considering the argumentation framework AF_4 of Fig. 3, we have that $\mathcal{FP}(\mathrm{AF}_4) = \{\{\alpha\}, \{\beta\}\}$, therefore both arguments belong to $\mathrm{P}_{\mathcal{P}}(\mathrm{AF}_4)$. With reference to the argumentation framework AF_5, Definition 6 identifies the empty set as the unique preferred extension, therefore all the arguments belong to $\mathrm{D}_{\mathcal{P}}(\mathrm{AF}_5)$. More generally, with odd-length cycles there is a unique empty extension, while with even-length cycles non-empty extensions exist but their intersection is empty.

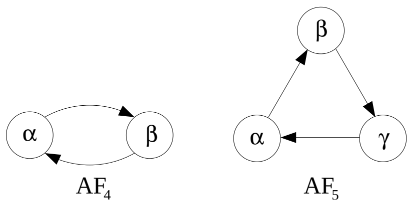

Fig. 3. Even-length and odd-length cycles

The peculiar way to assign a defeat status to odd-length cycles has recently been indicated as "puzzling" by Pollock [10]. As to our knowledge, however, this difference has been considered a mere question of symmetry and elegance in previous literature. We show in the following example that the different treatment of odd-length cycles is a real problem since it gives rise to counter-intuitive results.

Example 5. Considering the argumentation framework AF_6 shown in Fig. 4, it turns out that $\mathcal{FP}(\mathrm{AF}_6) = \{\{\alpha, \delta\}\}$, therefore α and δ are justified according to the preferred semantics. By replacing the cycle (α, β, γ) with a two-length cycle, we obtain the argumentation framework AF_7 whose arguments all belong to $\mathrm{P}_{\mathcal{P}}(\mathrm{AF}_7)$ (and a similar result is obtained with any other even-length cycle).

In the example above α and δ emerge (unreasonably) undefeated, while all nodes would be provisionally defeated with a similar graph encompassing an even-length cycle. It does not seem acceptable that different results in conceptually similar situations depend on the cycle length: symmetry reasons suggest that all cycles should be treated equally. The difference arises because an odd-length cycle has no extensions besides the empty one: as a consequence, there is no extension where δ is out and γ is in (such an extension would instead exist with an even-length cycle). Since δ defeats γ, in this context also α survives. Notice that a similar situation arises by replacing the three-length cycles with any odd-length cycle: in a sense, odd-length cycles are in this case 'weaker' than even-length cycles, since they are not able to prevent δ from being justified. The opposite happens in the following example:

Example 6. With reference to the argumentation framework AF_8 shown in Fig. 4, it turns out that $\mathcal{FP}(\mathrm{AF}_8) = \{\{\delta_2\}\}$, therefore δ_2 is undefeated while all the other arguments are not justified. On the other hand, by replacing the three-length cycle with an even-length cycle, we obtain an argumentation framework whose arguments are all provisionally defeated.

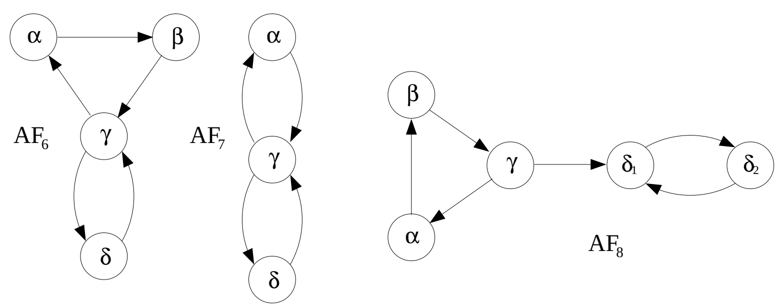

Fig. 4. Two problematic argumentation frameworks for preferred semantics

In the example above, the absence of non-empty extensions for the three-length cycle prevents the existence of extensions including δ_1, leaving δ_2 as the only accepted argument, while this would not happen with an even-length cycle. Notice that in this case odd-length cycles are 'stronger' than even-length cycles, since the status of δ_1 is the same as if it would be attacked by an initial node. In summary, we notice that besides being treated differently with respect to even-length cycles, odd-length cycles exhibit anomalous behaviors: they change their capability of defeating other arguments depending on the topology of the defeat graph.

4 Exploiting the Notion of Maximal Conflict-free Set

The analysis carried out in previous sections shows that neither the grounded nor the preferred semantics are completely satisfactory and suggests the following requirements for the study of an improved semantics, able to preserve the advantages of both of them:

– it should discriminate floating arguments as the preferred semantics;
– it should handle odd-length cycles in the same way as even-length cycles, as the grounded semantics;
– it should correctly handle the problematic examples shown above;
– it should not be more skeptical than the grounded semantics, in particular, it should agree with it upon the status of undefeated and defeated arguments.

On the basis of the results provided by Schlechta [9], mentioned in Sect. 2, the first requirement can only be satisfied by a multiple-status approach. Thus, to satisfy the second requirement we look for a new notion of extension, able to remove the anomalous treatment of odd-length cycles. After identifying our candidate definition, we will check its properties concerning third and fourth requirement. To figure out a proper notion of extension, let us consider again Example 4, in which we have recognized as anomalous the fact that AF_5 admits the empty set as its unique extension. In order to reconcile the treatments of AF_4 and AF_5, we can look for the set $\mathcal{E}$ of non-empty extensions that can be admitted for AF_5. First of all, we cannot tolerate contradictions in any extension, therefore each extension has to include one node exactly. Moreover, all nodes should be treated equally, therefore the only possibility for $\mathcal{E}$ is the set $\{\{\alpha\},\{\beta\},\{\gamma\}\}$. We notice that $\mathcal{E}$ is made up of all conflict-free sets of AF_5 that are maximal, and this suggests to exploit this notion as a basis for a new definition of extension.

Definition 8. *Given an argumentation framework* $\mathrm{AF} =< \mathcal{A}, \rightarrow >$*, we denote as* $\mathcal{FI}(\mathrm{AF})$ *the set made up of the maximal (with respect to set inclusion) conflict-free subsets of* $\mathcal{A}$*.*

The above intuition is confirmed by the fact that, by defining the set of justified arguments as the intersection of all maximal conflict-free sets, the problematic examples of the above section are handled correctly. In particular, with reference to the argumentation framework AF_6 of Example 5, we have that $\mathcal{FI}(\mathrm{AF}_6) = \{\{\alpha,\delta\},\{\gamma\},\{\beta,\delta\}\}$, therefore all the arguments are provisionally defeated, as prescribed by the grounded semantics.

Notice that Definition 8 is strictly weaker than Definition 6, since the absence of conflicts is one of the conditions for admissibility. Actually, while this brings about a correct handling of Example 4 and Example 5, it does not represent a satisfactory solution, since due to the increased number of extensions, it would tend to assign the status of provisionally defeated to a large number of arguments (often all of them): this happens, for instance, even for the argumentation framework AF_1 of Example 1, where we have that $\mathcal{FI}(\mathrm{AF}_1) = \{\{\alpha,\gamma\},\{\beta\}\}$. Notice that, in this case, the requirement of admissibility would have excluded, among the elements of $\mathcal{FI}(\mathrm{AF}_1)$, the set $\{\beta\}$, yielding the intuitively correct result. Thus, we are lead to add some further condition to the Definition 8, in order to capture only a subset of the maximal conflict-free sets. In order to do this, we draw inspiration from the way the defeat status can be computed according to the grounded semantics. Considering again Example 1 as a simple reference, basically, computation proceeds from the frontier of the defeat graph towards the inside: the initial node α is assigned the status of undefeated, causing β, which is attacked by α, to be assigned the status of defeated, and this in turn causes γ, whose unique defeater β is defeated, to be assigned the status of undefeated. Thus, the set $\{\beta\}$ is rejected in this computation schema, which is therefore a promising candidate as a way to identify the extensions among maximal conflict-free sets.

In order to refine this intuition, let us consider again Example 3: according to the first requirement stated above, our approach should capture exactly the preferred extensions $P_1 = \{\alpha, \delta\}$ and $P_2 = \{\beta, \delta\}$. Starting from the frontier of the graph, the construction of these extensions might proceed according to the following steps:

1. Consider the subgraph involving $\{\alpha, \beta\}$, and identify the relevant maximal conflict-free sets $\overline{P_1} = \{\alpha\}$ and $\overline{P_2} = \{\beta\}$;
2. Consider then node γ for possible additions to the sets identified in the previous step: notice that $\overline{P_1}$ includes the defeater α of γ, therefore γ cannot be added to $\overline{P_1}$. For the same reason, γ cannot be added to $\overline{P_2}$ as well;
3. Consider node δ: it can be added to $\overline{P_1}$ obtaining the extension P_1, since its unique defeater γ has not be added to $\overline{P_1}$. In the same way, we obtain P_2 as $\overline{P_2} \cup \{\delta\}$.

Notice that, in steps 1–3, we have considered the *strongly connected components* of the defeat graph, i.e. $\{\alpha, \beta\}$, $\{\gamma\}$ and $\{\delta\}$, respectively. In a sense, we have generalized the defeat status computation prescribed by the grounded semantics, by considering strongly connected components instead of single nodes. In particular, the extensions have been constructed by completing maximal conflict-free sets in an incremental way, starting from the frontier of the graph and proceeding towards the interior. In order to proceed with this analysis in more formal terms, let us introduce the following definitions.

Definition 9. *Given an argumentation framework* $\mathrm{AF} =< \mathcal{A}, \rightarrow>$*, two nodes* $\alpha, \beta \in \mathcal{A}$ *are* path-equivalent *iff either* $\alpha = \beta$ *or there is a path from* α *to* β *and a path from* β *to* α*. The strongly connected components of* AF *are the equivalence classes of vertices under the relation of path-equivalence. The set of the strongly connected components of* AF *is denoted as* $\mathrm{SCC}(\mathrm{AF})$*.*

Definition 10. *Given an argumentation framework* $\mathrm{AF} =< \mathcal{A}, \rightarrow>$ *and a strongly connected component* $S \in \mathrm{SCC}(\mathrm{AF})$*,* $\mathrm{parents}(S) = \{P \in \mathrm{SCC}(\mathrm{AF}) \mid P \neq S \wedge \exists \alpha \in P, \beta \in S : \alpha \rightarrow \beta\}$*, and* $\mathrm{parents}^*(S) = \{\alpha \in \mathcal{A} \mid \alpha \notin S \wedge \exists \beta \in S : \alpha \rightarrow \beta\}$*.*

Definition 11. *Let* $\mathrm{AF} =< \mathcal{A}, \rightarrow>$ *be an argumentation framework, and let* S *be a set* $S \subseteq \mathcal{A}$*. The* restriction *of* AF *to* S *is the argumentation framework* $\mathrm{AF}{\downarrow_S} =< S, \rightarrow \cap (S \times S) >$.

Let $\mathcal{SG}$ be the graph obtained by considering strongly connected components as single nodes, i.e. $\mathcal{SG} =< \mathrm{SCC}(\mathrm{AF}), R^* >$ where $(S_i, S_j) \in R^*$ iff $S_i \in \mathrm{parents}(S_j)$. It is easy to see that $\mathcal{SG}$ is acyclic: this justifies the idea of computing a particular extension E from the frontier towards the inside of the defeat graph. Basically, we start from the strongly connected components S_i that are initial in $\mathcal{SG}$, including in E a maximal conflict-free set of $\mathrm{AF}{\downarrow}_{S_i}$ for each S_i. Then, we proceed by considering an $S \in \mathrm{SCC}(\mathrm{AF})$ such that every $P \in \mathrm{parents}(S)$ is initial. Of course, E should not include those nodes of S that are attacked by nodes previously included in E. The question is how to proceed with the set S^U made up of the other nodes of S. If there is just a single node in S^U, the indications provided by Example 1 suggest to include it in E. If, on the other hand, $|S^U| > 1$, a tentative solution would be to include in E a maximal independent set of S^U. However, a simple example reveals that this option does not constrain enough the set of the extensions that can be identified from maximal conflict-free sets.

Example 7. Considering the argumentation framework AF_9 shown Fig. 5, we have that $\mathrm{SCC}(\mathrm{AF}_9) = \{S_1, S_2\}$, where $S_1 = \{\alpha\}$ and $S_2 = \{\beta_1, \beta_2, \beta_3, \beta_4\}$. S_1 is initial, and its unique maximal conflict-free set is $\{\alpha\}$ itself. This in turn excludes β_1 from all the extensions, leading to select a maximal conflict-free set of the subgraph $\mathrm{AF}_9{\downarrow}_{S_2 \setminus \{\beta_1\}}$. It turns out that $\mathcal{FI}(\mathrm{AF}_9{\downarrow}_{S_2 \setminus \{\beta_1\}}) = \{\{\beta_2, \beta_4\}, \{\beta_3\}\}$, therefore we get the two extensions $E_1 = \{\alpha, \beta_2, \beta_4\}$ and $E_2 = \{\alpha, \beta_3\}$, yielding α undefeated, β_1 defeated and $\beta_2, \beta_3, \beta_4$ provisionally defeated. However, in order to get the same outcome as the grounded (and preferred) semantics, only E_1 should be identified as an extension, while E_2 should be excluded.

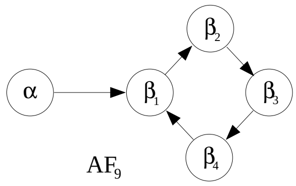

Fig. 5. An example supporting a recursive definition of extensions

In order to overcome this difficulty, in the example above $S_2 \setminus \{\beta_1\}$ should be treated in the same way as an ordinary graph, i.e. proceeding again from the frontier towards the inside. This suggests the alternative option that we choose, i.e. to define extensions recursively.

Definition 12. *Given an argumentation framework* $\mathrm{AF} =< \mathcal{A}, \rightarrow >$*, a set* $E \subseteq \mathcal{A}$ *and a strongly connected component* $S \in \mathrm{SCC}(\mathrm{AF})$*, we define:*

- $S^D(E) = \{\alpha \in S \mid \exists \beta \in \mathrm{parents}^*(S) : \beta \in E \wedge \beta \rightarrow \alpha\}$
- $S^U(E) = S \setminus S^D(E)$

In our proposal, the set of extensions, denoted as $\mathcal{FM}(\mathrm{AF})$, is defined as follows:

Definition 13. *Given an argumentation framework* $\mathrm{AF} = <\mathcal{A}, \rightarrow>$ *and a set* $E \subseteq \mathcal{A}$, *we have that* $E \in \mathcal{FM}(\mathrm{AF})$ *iff* $\forall S \in \mathrm{SCC}(\mathrm{AF})$

1. $S^D(E) \cap E = \emptyset$; *and*
2. $S^U(E) \cap E \begin{cases} \in \mathcal{FI}(\mathrm{AF}{\downarrow}_{S^U(E)}) & \textit{if } |\mathrm{SCC}(\mathrm{AF}{\downarrow}_{S^U(E)})| = 1 \\ \in \mathcal{FM}(\mathrm{AF}{\downarrow}_{S^U(E)}) & \textit{otherwise} \end{cases}$

Following the usual multiple-status approach, the defeat status of arguments is identified by the sets $\mathrm{U}_{\mathcal{M}}(\mathrm{AF})$, $\mathrm{D}_{\mathcal{M}}(\mathrm{AF})$, $\mathrm{P}_{\mathcal{M}}(\mathrm{AF})$, defined as in Definition 7 with reference to $\mathcal{FM}(\mathrm{AF})$ instead of $\mathcal{FP}(\mathrm{AF})$. In order to better understand Definition 13, let us show that, differently from the preferred semantics, it gives the right outcome to Example 6.

Example 8. We have that $\mathrm{SCC}(\mathrm{AF}_8) = \{S_1, S_2\}$, where $S_1 = \{\alpha, \beta, \gamma\}$ and $S_2 = \{\delta_1, \delta_2\}$. Taking into account Definition 12 and the fact that parents$(S_1) = \emptyset$, for any E $S_1^D(E) = \emptyset$ and $S_1^U(E) = S_1$. Thus, from Definition 13 a generic extension E must satisfy $(S_1 \cap E) \in \mathcal{FI}(\mathrm{AF}_8{\downarrow}_{S_1^U(E)})$, i.e. $(S_1 \cap E) \in \{\{\alpha\}, \{\beta\}, \{\gamma\}\}$. In case $(S_1 \cap E) = \{\alpha\}$, taking into account that parents$^*(S_2) = \{\gamma\}$ we get $S_2^D(E) = \emptyset$ and $(S_2 \cap E) \in \{\{\delta_1\}, \{\delta_2\}\}$: thus, we identify the extensions $E_1 = \{\alpha, \delta_1\}$ and $E_2 = \{\alpha, \delta_2\}$. Reasoning in a similar way in the case $(S_1 \cap E) = \{\beta\}$, we identify the extensions $E_3 = \{\beta, \delta_1\}$ and $E_4 = \{\beta, \delta_2\}$. Finally, if $(S_1 \cap E) = \{\gamma\}$ then $S_2^D(E) = \{\delta_1\}$, entailing by the first point of Definition 13 that $\delta_1 \notin E$. Moreover, $S_2^U(E) = \{\delta_2\}$, yielding $\delta_2 \in E$ and thus identifying the extension $E_5 = \{\gamma, \delta_2\}$. In sum, $\mathcal{FM}(\mathrm{AF}_8) = \{E_1, E_2, E_3, E_4, E_5\}$, therefore all the arguments are provisionally defeated.

It can be seen that all other examples considered above are handled correctly by the proposed semantics. In particular, in the argumentation frameworks AF_1, AF_2, AF_7 and AF_9 where preferred and grounded semantics agree, our semantics gives the same results ($\mathcal{FM}() = \mathcal{FP}()$ in all cases). Also in the argumentation framework AF_3 $\mathcal{FM}(\mathrm{AF}_3) = \mathcal{FP}(\mathrm{AF}_3)$, therefore our semantics correctly agrees with preferred semantics. Finally, $\mathcal{FM}(\mathrm{AF}_6) = \{\{\alpha, \delta\}, \{\beta, \delta\}, \{\gamma\}\}$, therefore our semantics agrees with grounded semantics as desired.

5 Relationships with Grounded and Preferred Semantics

After having validated our proposal by means of examples, in this section we show that it maintains some relationships with both the grounded and preferred semantics. First, we consider the fourth requirement stated in previous section, i.e. the agreement with the grounded semantics upon the status of undefeated and defeated arguments (proofs are not given due to space limitations). This result relies on two properties of the grounded semantics, which relate the defeat status assignment prescribed by the grounded semantics to the strongly connected components of the defeat graph. In particular, the first considers a non-trivial strongly connected component whose external attackers (if any) are all defeated or provisionally defeated, establishing that, in this case, all of its nodes are provisionally defeated.

Proposition 1. *Given* $\mathrm{AF} = <\mathcal{A}, \rightarrow>$*, let* $S \in \mathrm{SCC}(\mathrm{AF})$ *be such that* $|S| > 1$ *and* $\forall \gamma \in \mathrm{parents}^*(S)$ $\gamma \in (\mathrm{D}_{\mathcal{G}}(\mathrm{AF}) \cup \mathrm{P}_{\mathcal{G}}(\mathrm{AF}))$*. Then,* $S \subseteq \mathrm{P}_{\mathcal{G}}(\mathrm{AF})$*.*

The subsequent property introduces two subsets related to $S^D(E)$ and $S^U(E)$ in Definition 12, and establishes some relationships between the status assigned by the grounded semantics to their nodes and the constraints on E stated in Definition 12.

Proposition 2. *Let us consider* $\mathrm{AF} = <\mathcal{A}, \rightarrow>$ *and* $S \in \mathrm{SCC}(\mathrm{AF})$*. Let* $S^D \subseteq S$ *be a subset of* S *such that*

1. $S^D \supseteq \{\alpha \in S \mid \exists \beta \in \mathrm{parents}^*(S), \beta \rightarrow \alpha, \beta \in \mathrm{U}_{\mathcal{G}}(\mathrm{AF})\}$*; and*
2. $S^D \subseteq \{\alpha \in S \mid \exists \beta \in \mathrm{parents}^*(S), \beta \rightarrow \alpha, \beta \in (\mathrm{U}_{\mathcal{G}}(\mathrm{AF}) \cup \mathrm{P}_{\mathcal{G}}(\mathrm{AF}))\}$

and let $S^U = S \setminus S^D$*. Then, we have that* $\forall \alpha \in S^D$ $\alpha \in (\mathrm{D}_{\mathcal{G}}(\mathrm{AF}) \cup \mathrm{P}_{\mathcal{G}}(\mathrm{AF}))$*, and* $\forall \gamma \in S^U$*:*

- *if* $\gamma \in \mathrm{U}_{\mathcal{G}}(\mathrm{AF})$*, then* $\gamma \in \mathrm{U}_{\mathcal{G}}(\mathrm{AF}{\downarrow}_{S^U})$*;*
- *if* $\gamma \in \mathrm{D}_{\mathcal{G}}(\mathrm{AF})$*, then* $\gamma \in \mathrm{D}_{\mathcal{G}}(\mathrm{AF}{\downarrow}_{S^U})$*;*

The following theorem, exploiting the above properties, proves the agreement with grounded semantics upon the status of defeated and undefeated arguments.

Theorem 1. *Given* $\mathrm{AF} = <\mathcal{A}, \rightarrow>$*, we have that* $\forall E \in \mathcal{FM}(\mathrm{AF})$ $\mathrm{U}_{\mathcal{G}}(\mathrm{AF}) \subseteq E \wedge \mathrm{D}_{\mathcal{G}}(\mathrm{AF}) \subseteq (\mathcal{A} \setminus E)$*.*

Proof (Sketch). Referring to Definition 13, we consider a generic $E \in \mathcal{FM}(\mathrm{AF})$, and we assume recursively that, $\forall S_i \in \mathrm{SCC}(\mathrm{AF})$,

- $\forall P \in \mathcal{FM}(\mathrm{AF}{\downarrow}_{S_i^U(E)})$ $\mathrm{U}_{\mathcal{G}}(\mathrm{AF}) \subseteq P$
- $\forall P \in \mathcal{FM}(\mathrm{AF}{\downarrow}_{S_i^U(E)})$ $\mathrm{D}_{\mathcal{G}}(\mathrm{AF}) \subseteq (\mathcal{A} \setminus P)$

Then, reasoning by induction on the strongly connected components of AF, we prove that $\mathrm{U}_{\mathcal{G}}(\mathrm{AF}) \subseteq E$ and that $\mathrm{D}_{\mathcal{G}}(\mathrm{AF}) \subseteq (\mathcal{A} \setminus E)$. In particular, Proposition 1 is exploited to show that, if $|\mathrm{SCC}(\mathrm{AF}{\downarrow}_{S^U(E)})| = 1$, then all the nodes of $S^U(E)$ are provisionally defeated, so that there is nothing to prove for them (this happens for instance for initial strongly connected components). On the other hand, the main roles of Proposition 2 concern the first point of Definition 13 and the case $|\mathrm{SCC}(\mathrm{AF}{\downarrow}_{S^U(E)})| > 1$, where it is exploited to prove that if $S^U(E) \cap E \in \mathcal{FM}(\mathrm{AF}{\downarrow}_{S^U(E)})$, then the claim is satisfied for nodes of $S^U(E)$.

As far as preferred semantics is concerned, given a generic $\mathrm{AF} = <\mathcal{A}, \rightarrow>$ it is possible to prove that any preferred extension is included in one of our extensions, i.e. $\forall P \in \mathcal{FP}(\mathrm{AF})$ $\exists E \in \mathcal{FM}(\mathrm{AF}) : P \subseteq E$. This, in turn, entails that preferred semantics agrees upon the status of arguments that are defeated according to ours.

6 Conclusions

In this paper, we have proposed a novel argumentation semantics that, while maintaining the same capability of discriminating floating arguments as preferred semantics, correctly deals with the semantic problems arising from odd-length cycles and satisfies a set of requirements intuitively appealing. The symmetry assumptions that underly our work are related to an interpretation of argumentation as a framework for defeasible reasoning, following e.g. [1], while in other approaches that consider argumentation as a branch of dialogue [11] it may be the case that a different treatment of odd and even-length cycles is appropriate.

Our proposal builds on intuitions coming from both grounded and preferred semantics and, in a sense, combines the advantages of both of them, agreeing in several problematic examples with that among the two semantics which is closer to intuition. As for future work, we will investigate the relationship between our semantics and the notions of attack and defence lying at the heart of preferred semantics.

Acknowledgments. We thank the referees for their helpful comments.

References

1. Pollock, J.L.: How to reason defeasibly. Artificial Intelligence **57** (1992) 1–42
2. Parsons, S., Sierra, C., Jennings, N.: Agents that reason and negotiate by arguing. Journal of Logic and Computation **8** (1998) 261–292
3. Baroni, P., Giacomin, M., Guida, G.: Extending abstract argumentation systems theory. Artificial Intelligence **120** (2000) 251–270
4. 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
5. Prakken, H., Vreeswijk, G.A.W.: Logics for defeasible argumentation. In Gabbay, D.M., Guenthner, F., eds.: Handbook of Philosophical Logic, Second Edition. Kluwer Academic Publishers, Dordrecht (2001)
6. Pollock, J.L.: Justification and defeat. Artificial Intelligence **67** (1994) 377–407
7. Vreeswijk, G.A.W.: Abstract argumentation systems. Artificial Intelligence **90** (1997) 225–279
8. Makinson, D., Sclechta, K.: Floating conclusions and zombie paths: Two deep difficulties in the 'directly skeptical' approach to defeasible inheritance networks. Artificial Intelligence **48** (1991) 199–209
9. Schlechta, K.: Directly sceptical inheritance cannot capture the intersection of extensions. Journal of Logic and Computation **3** (1993) 455–467
10. Pollock, J.L.: Defeasible reasoning with variable degrees of justification. Artificial Intelligence **133** (2001) 233–282
11. Walton, D., Krabbe, E.: Commitment in Dialogue: Basic concept of interpersonal reasoning. State University of New York Press, Albany NY (1995)