

Weighted Argument Systems: Basic Definitions, Algorithms, and Complexity Results

Paul E. Dunne^a, Anthony Hunter^b, Peter McBurney^a,
Simon Parsons^c, Michael Wooldridge^a

^a*Dept. of Computer Science, University of Liverpool, Liverpool, UK*

^b*Dept. of Computer Science, University College London, London, UK*

^c*Brooklyn College, CUNY, Brooklyn, NY, USA*

Abstract

We introduce and investigate a natural extension of Dung's well-known model of argument systems in which attacks are associated with a *weight*, indicating the relative strength of the attack. A key concept in our framework is the notion of an *inconsistency budget*, which characterises how much inconsistency we are prepared to tolerate: given an inconsistency budget β , we would be prepared to disregard attacks up to a total weight of β . The key advantage of this approach is that it permits a much finer grained level of analysis of argument systems than unweighted systems, and gives useful solutions when conventional (unweighted) argument systems have none. We begin by reviewing Dung's abstract argument systems, and motivating weights on attacks (as opposed to the alternative possibility, which is to attach weights to arguments). We then present the framework of weighted argument systems. We investigate solutions for weighted argument systems and the complexity of computing such solutions, focussing in particular on weighted variations of grounded extensions. Finally, we relate our work to the most relevant examples of argumentation frameworks that incorporate strengths.

Key words: Argumentation, handling inconsistency, computational complexity

1. Introduction

Inconsistency between the beliefs and/or preferences of agents is ubiquitous in everyday life, and yet coping with inconsistency remains an essentially unsolved problem in artificial intelligence [12]. One of the key aims of *argumentation* research is to provide principled techniques for handling inconsistency [13, 46].

Email addresses: ped@csc.liv.ac.uk (Paul E. Dunne), a.hunter@cs.ucl.ac.uk
(Anthony Hunter), mcburney@liverpool.ac.uk (Peter
McBurney), parsons@sci.brooklyn.cuny.edu (Simon Parsons), mjl@csc.liv.ac.uk
(Michael Wooldridge)

Although there are several different perspectives on argumentation, a common view is that argumentation starts with a collection of statements, called *arguments*, which are related through the notions of *support* and *attack*. Typically, argument α_1 supporting argument α_2 would be grounds for accepting α_2 if one accepted α_1 , while argument α_1 attacking argument α_2 would be grounds for *not* accepting α_2 if one accepted α_1 . Now, if we allow arguments to attack one-another, then such collections of arguments may be inconsistent; and the key question then becomes how to obtain a rationally justifiable position from such an inconsistent argument set. Various solutions have been proposed for this problem, such as *admissible sets*, *preferred extensions*, and *grounded extensions* [19]. However, none of these solutions is without drawbacks. A common situation is that, while a solution may be guaranteed to give an answer, the answer may be the empty set. Conversely, several answers may be provided, with nothing to distinguish between them. These drawbacks limit the value of these solutions as argument analysis tools.

In part to overcome these difficulties, there is a trend in the literature on formalizations of argumentation towards considering the *strength* of arguments. In this work, which goes back at least as far as [30], it is recognized that not all arguments are equal in strength, and that this needs to be taken into account when finding extensions of a collection of arguments and counterarguments. We review this literature in Section 2.1, and we conclude that whilst it is clear that taking the strength of arguments into account is a valuable development, it is not just the strength of the arguments, *per se*, that is important. The strength of the attack that one argument (which may itself be very strong) makes on another, can be weak.

In this paper, we introduce, formalise, and investigate a natural extension of Dung’s well-known model of argument systems [19], in which attacks between arguments are associated with a numeric *weight*, indicating the relative strength of the attack, or, equivalently, how reluctant we would be to disregard the attack. For example, consider the following arguments:

- α_1 = The house is in a good location, it is large enough for our family and it is affordable: we should buy it.
- α_2 = The house suffers from subsidence, which would be prohibitively expensive to fix: we should not buy it.

These arguments are mutually attacking, and in the terminology of Dung’s abstract argument systems [19], both arguments are credulously accepted, neither is sceptically accepted, and the grounded extension is empty. Thus the conventional analysis is not very useful for this scenario. However, the representation we are using surely misses a key point: the attacks are *not of equal weight*. We would surely regard the attack of α_2 on α_1 as being much stronger than the attack of α_1 on α_2 , though both are very strong arguments in their own right. Our framework allows us to take these differing weights of attack into consideration. (We note that an alternative to our approach, which has to some extent already been considered in the literature, is to attach weights to arguments,

rather than attaching weights to attacks between arguments. A detailed discussion of the relative merits of the two possibilities is given in section 2.)

A key concept in our framework is the notion of an *inconsistency budget*. The inconsistency budget characterises how much inconsistency we are prepared to tolerate: given an inconsistency budget β , we would be prepared to disregard attacks up to a total weight of β . By increasing the inconsistency budget, we get progressively more solutions, and this in turn gives a preference ordering over solutions: we prefer solutions obtained with a smaller inconsistency budget. This approach permits a very fine-grained level of analysis, and gives useful, non-trivial solutions when conventional (unweighted) argument systems have none.

The remainder of this paper is structured as follows:

- We begin by reviewing Dung’s abstract argument systems, and providing motivation for the idea of extending Dung’s framework with weights. We discuss in particular the relative merits of attaching weights to arguments versus attaching weights to attacks.
- In section 3, we present the formal model of weighted argument systems that we work with throughout the remainder of the paper. After presenting the model, we discuss the semantics of weights, i.e., how weights can be interpreted, and how they might be derived for some different domains.
- In sections 4–6, we investigate solutions to weighted argument systems and the associated complexity of computing these solutions, focussing in particular on weighted variations of grounded extensions. We also consider the more general algorithmic and combinatorial properties of our framework.
- In section 7, we consider the relationship between our weighted argument systems and four other well-known related extensions of Dung’s argument framework. We show that, in a precise formal sense, our weighted argument systems are strictly more expressive: all of the other frameworks considered can be represented as weighted argument systems, but there exist weighted argument systems that have no equivalent representation in the alternative frameworks.
- We conclude with some issues for future work.

2. Background

Since weighted argument systems and their associated solutions generalise Dung’s well-known abstract argument systems model, we begin by recalling the key concepts from this model — note that further discussion of related work may be found in section 7.

Definition 1. A *Dung-style* abstract argument system is a pair $D = \langle \mathcal{X}, \mathcal{A} \rangle$ where $\mathcal{X} = \{\alpha_1, \dots, \alpha_k\}$ is a finite set of arguments, and $\mathcal{A} \subseteq \mathcal{X} \times \mathcal{X}$ is a binary attack relation on \mathcal{X} [19].¹

The next step is to define *solutions* for such argument systems, i.e., concepts of what constitutes a set of mutually compatible arguments from \mathcal{X} within a system $\langle \mathcal{X}, \mathcal{A} \rangle$. Typically, such subsets are defined via predicates $\sigma : 2^{\mathcal{X}} \rightarrow \{\top, \perp\}$ so that when $\sigma(S)$ holds of $S \subseteq \mathcal{X}$ in $\langle \mathcal{X}, \mathcal{A} \rangle$ the set S is viewed as acceptable with respect to the criteria defined by σ (see e.g., Baroni and Giacomin [8]). We review a number of proposals for such criteria in the following definition.

Definition 2. Given an argument system $D = \langle \mathcal{X}, \mathcal{A} \rangle$ and a set $S \subseteq \mathcal{X}$, we say that S is *initial* if $\nexists \alpha_1 \in S$ such that $\exists \alpha_2 \in \mathcal{X} : \langle \alpha_2, \alpha_1 \rangle \in \mathcal{A}$ using $initial(\langle \mathcal{X}, \mathcal{A} \rangle)$ to denote the (unique) maximal such set. A set $S \subseteq \mathcal{X}$ is *conflict free* if $\forall \alpha_1, \alpha_2 \in S \langle \alpha_1, \alpha_2 \rangle \notin \mathcal{A}$; an argument α is *acceptable w.r.t. S* if $\forall \beta \in \mathcal{X}$ such that $\langle \beta, \alpha \rangle \in \mathcal{A}$, $\exists \gamma \in S$ such that $\langle \gamma, \beta \rangle \in \mathcal{A}$; the characteristic function, $\mathcal{F}(S)$ reports the set of arguments that are acceptable to S ; S is *admissible* if it is both conflict-free and every $\alpha \in S$ is acceptable w.r.t. S . The grounded extension, $ge(\langle \mathcal{X}, \mathcal{A} \rangle)$, is formed by the fixed-point of \mathcal{F}^k where $\mathcal{F}^0 = \emptyset$ and $\mathcal{F}^{k+1} = \mathcal{F}(\mathcal{F}^k)$. It is shown in [19] that (for finite systems) this fixed point always exists. Finally S is a *preferred extension* if it is a maximal (w.r.t. \subseteq) admissible set. Given an acceptability semantics, $\sigma : 2^{\mathcal{X}} \rightarrow \{\top, \perp\}$, we denote by $\mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle)$ the set of subsets of \mathcal{X} :

$$\mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle) = \{ S \subseteq \mathcal{X} : \sigma(S) \}$$

Specific cases of interest are:

$$\begin{aligned} \mathcal{E}_{init}(\langle \mathcal{X}, \mathcal{A} \rangle) &= \{ S : S = initial(\langle \mathcal{X}, \mathcal{A} \rangle) \} \\ \mathcal{E}_{adm}(\langle \mathcal{X}, \mathcal{A} \rangle) &= \{ S : S \text{ is admissible in } \langle \mathcal{X}, \mathcal{A} \rangle \} \\ \mathcal{E}_{pr}(\langle \mathcal{X}, \mathcal{A} \rangle) &= \{ S : S \text{ is a preferred extension of } \langle \mathcal{X}, \mathcal{A} \rangle \} \\ \mathcal{E}_{gr}(\langle \mathcal{X}, \mathcal{A} \rangle) &= \{ S : S \text{ is the grounded extension of } \langle \mathcal{X}, \mathcal{A} \rangle \}. \end{aligned}$$

The initial arguments \mathcal{E}_{init} , i.e., those belonging to the set $initial(\langle \mathcal{X}, \mathcal{A} \rangle)$, form in some regards, the least problematic notion of “collection of compatible arguments”, however, while every argument system contains a unique initial set of arguments, (i.e., $|\mathcal{E}_{init}(\langle \mathcal{X}, \mathcal{A} \rangle)| = 1$), it may be that $\mathcal{E}_{init}(\langle \mathcal{X}, \mathcal{A} \rangle) = \{\emptyset\}$. Such trivial cases are typically unhelpful. If we do not have a non-empty initial set of arguments, (which is the more general case), then we might look at conflict-free sets (the so-called *naive semantics*), admissible sets, and preferred extensions. There will always be at least one preferred extension, although, again, this may be the empty set [19, p.327]. Note that non-empty

¹Note that Dung’s model does not assume any internal structure for arguments, nor give any concrete interpretation for them. The intended interpretation of the attack relation in Dung’s model is also not completely defined, but intuitively, $(\alpha_1, \alpha_2) \in \mathcal{A}$ means that if one accepts α_1 , then one should not completely accept α_2 . In other words, *it would be inconsistent to accept α_2 if one accepted α_1 .*

function $ge(\mathcal{X}, \mathcal{A})$ returns a subset of \mathcal{X}

1. $in \leftarrow out \leftarrow \emptyset$
2. **while** $initial(\langle \mathcal{X}, \mathcal{A} \rangle) \neq \emptyset$ **do**
3. $in \leftarrow in \cup initial(\langle \mathcal{X}, \mathcal{A} \rangle)$
4. $out \leftarrow out \cup \{\alpha \in \mathcal{X} : \exists \alpha' \in in \text{ s.t. } \langle \alpha', \alpha \rangle \in \mathcal{A}\}$
5. $\mathcal{X} \leftarrow \mathcal{X} \setminus (out \cup in)$
6. $\mathcal{A} \leftarrow \mathcal{A}$ restricted to \mathcal{X}
7. **end-while**
8. **return** in .

Figure 1: Computing the grounded extension of an (unweighted) argument system: the function $ge(\dots)$.

preferred extensions may exist in argument systems whose set of initial arguments is empty, and so we can usefully apply this solution in some situations where the initial arguments are not an applicable analytical concept.

It is easily seen that $\mathcal{F}^1 = \mathcal{F}(\emptyset) = initial(\langle \mathcal{X}, \mathcal{A} \rangle)$ so that initial arguments play a central role in a constructive algorithm for computing the *grounded extension*: starting from the set $initial(\langle \mathcal{X}, \mathcal{A} \rangle)$, eliminate arguments that these arguments attack; and iterate until we reach no change (the next set of “initial arguments” having removed $initial(\langle \mathcal{X}, \mathcal{A} \rangle)$ and those arguments attacked by it forms the set $\mathcal{F}^2 = \mathcal{F}(\mathcal{F}(\emptyset))$). The polynomial time algorithm to compute the grounded extension of an argument system is given in Figure 1.

Notice that, while all of these criteria are guaranteed to give some “answer”, it is possible that the only answer they give is the empty set, i.e., $\mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle) \neq \emptyset$ but we may have $\mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle) = \{\emptyset\}$. We view this as a key limitation of conventional systems.

Of course as well as ideas of criteria by which subsets of \mathcal{X} are distinguished as allowed or disqualified, we are concerned with properties of *arguments* themselves with respect to such properties. We focus on the notions of *credulous* and *sceptical* acceptability.

Definition 3. Given a semantics $\sigma : 2^{\mathcal{X}} \rightarrow \{\top, \perp\}$, an argument, $\alpha \in \mathcal{X}$ is credulously accepted w.r.t. σ in $\langle \mathcal{X}, \mathcal{A} \rangle$ if there is at least one $S \subseteq \mathcal{X}$ for which $S \in \mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle)$ and $\alpha \in S$. An argument, $\alpha \in \mathcal{X}$, is sceptically accepted w.r.t. σ in $\langle \mathcal{X}, \mathcal{A} \rangle$ if $\forall S \in \mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle)$ we have $\alpha \in S$.

The decision problems credulous acceptance w.r.t. σ (denoted CA_σ) and sceptical acceptance w.r.t. σ (denoted SA_σ) both take as instances an argument system $\langle \mathcal{X}, \mathcal{A} \rangle$ together with an argument $\alpha \in \mathcal{X}$. An instance, $\langle \langle \mathcal{X}, \mathcal{A} \rangle, \alpha \rangle$ is accepted, in the former case, if $\exists S \in \mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle) : \alpha \in S$; such an instance is accepted, in the latter case, if $\forall S \in \mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle) : \alpha \in S$.

Sceptical acceptance imposes stronger conditions on the acceptability status of an argument than credulous acceptance: determining whether a given set of arguments

is the initial set or admissible can be solved in polynomial time; however, determining whether a set of arguments is a preferred extension is coNP-complete, checking whether an argument is credulously accepted w.r.t. admissibility (and thus w.r.t. preferred extensions) is NP-complete, while checking whether an argument is sceptically accepted w.r.t. preferred extensions is Π_2^P -complete [17, 23].

2.1. Extending Dung's Framework

As we noted above, while Dung's framework is both natural and powerful, the conventional solutions considered with respect to this framework have obvious limitations: there can be multiple solutions (in which case which do we choose?) and it may be that the only solution is the emptyset. There have been a number of proposals for extending Dung's framework in order to allow for more sophisticated modelling and analysis of conflicting information, in order to overcome these limitations. A common theme among some of these proposals is the observation that not all arguments are equal, and that the relative strength of the arguments needs to be taken into account somehow.

The first such extension of Dung's work that we are aware of is [44], where priorities between rules are used to resolve conflicts². These priorities seem best interpreted as relating to the strength of the arguments — indeed the strength of arguments are inferred from the strengths of the rules from which the arguments are constructed. A similar notion is at the heart of the argumentation systems in [1, 2], though here there is a preference order over all an agent's beliefs, and an argument has a preference level equal to the minimum level of the beliefs from which it is constructed.

Another early development of Dung's proposal with weights was Value-based Argumentation Frameworks (VAFs) [10]. In the VAF approach, the strength of an argument depends on the social values that it advances, and determining whether the attack of one argument on another succeeds depends on the comparative strength of the values advanced by the arguments concerned. Furthermore, some arguments can be shown to be acceptable whatever the relative strengths of the values involved are. This means that the agents involved in the argumentation can concur on the acceptance of arguments, even when they differ as to which social values are more important. One of the interesting questions that arises from this proposal is whether the notion of argument strength can be generalised from representing social values to representing other notions, and if so in what ways can the strength be harnessed for analysing argument graphs.

In a sense, a more general approach to developing Dung's proposal is that of bipolar argumentation frameworks (BAFs) which takes into account two kinds of interaction between arguments: a positive interaction (an argument can support another argument) and a negative interaction (an argument can attack another argument) [15]. The BAF approach incorporates a gradual interaction-based valuation process in which the value

²The article [30], which as we noted above considered some notion of argument strength, was not based on Dung's framework.

of each argument α only depends on the value of the arguments which are directly interacting with α in the argumentation system. Various functions for this process are considered but each value obtained is only a function of the original graph. As a result, no extra information is made available with which to ascertain the strength of an argument.

Recently, a game-theoretic approach, based on the minimax theorem, has been developed for determining the degree to which an argument is acceptable given the counterarguments to it, and by recursion the counterarguments to the counterarguments [33]. So given an abstract argument system, this game-theoretic approach calculates the strength of each argument in such a way that if an argument is attacked, then its strength falls, but if the attack is in turn attacked, then the strength in the original argument rises. Furthermore, the process for this conforms to an interpretation of game theory for argumentation. Whilst this gives an approach with interesting properties, and appealing behaviour, the strength that is calculated is a function of the original graph, and so like the BAF approach, no extra information is made available with which to determine the strength of each argument.

Notice that all of the above frameworks extend conventional argument models with weights (or preferences, priorities, ...) that are attached to *arguments*. An alternative — which we explore in the remainder of the present paper — is to attach weights to the *attacks* between arguments. Let us pause for a moment to consider the relative merits of weights on arguments versus weights on attacks.

First, let us revisit the house-buying example from the introductory section, involving α_1 and α_2 . Suppose we want to analyse this scenario by using weights (or priorities) on arguments. We may consider each argument independently and say that each is a strong argument in the sense that the premises are true (or very likely to be true), and that the claim follows (with little doubt) from the premises. Then, we consider the arguments together, and we see that α_1 is not as strong (as a rationale for making the decision of whether to buy the house) as α_2 . Thus, in order to adjust the weight of α_1 , we need to consider it *relative to* α_2 , and so the weight of α_1 is a relative notion, which is modulated via the attack by α_2 . Therefore, the weight of α_1 determined from the weight of the attacks on α_1 . So in this example, which is a common situation, we suggest that the *weight of attack* is a primitive notion, and the *weight of the argument* is a derived notion.

Second, let us consider a quite different kind of example based on arguments presented using classical logic (see, e.g., [13] for a detailed introduction to deductive argumentation). We assume each argument is a pair $\langle \Phi, \alpha \rangle$ where Φ is a consistent set of formulae (called the support) that entails α (which is called the claim).

$$A_1 = \langle \{a \wedge b \wedge c, a \wedge b \wedge c \rightarrow d\}, d \rangle$$

$$A_2 = \langle \{a \wedge b \wedge \neg c\}, a \wedge b \wedge \neg c \rangle$$

$$A_3 = \langle \{\neg a \wedge \neg b \wedge \neg c\}, \neg a \wedge \neg b \wedge \neg c \rangle$$

We assume that an argument $\langle \Phi, \alpha \rangle$ undercuts $\langle \Psi, \beta \rangle$ when α entails the negation of the support Ψ (i.e., when $\{\alpha\} \vdash \neg(\wedge \Psi)$). Here, we see that A_2 undercuts A_1 and A_3 undercuts A_1 . However, we also see that the degree of undercut (i.e., measure of inconsistency arising between the claim of the attacker and the support of the attacked) is greater for the undercut by A_3 than for the undercut by A_2 . We can measure that degree of undercut as the distance between the nearest models of the claim of the attacker and the support of the attacked normalized by the number of propositional letters in the language (for more information see [13]). Here, we see the degree of undercut by A_2 is $1/3$ whereas the degree of undercut by A_3 is 1. This therefore is a simple and precise illustration of a weight on the attack relation that can only be indirectly defined as a weight (or priority) on arguments.

To conclude this brief discussion on the question of whether weights on attacks would be better represented by weights (or priorities) on arguments, we believe that there are some situations where the weight of argument is a primitive notion and some situations where the weight of attack is a primitive notion. The above examples illustrate our claim that weight of attack can be a primitive notion, and we will provide further discussion and examples (such as when the weight represents the number of votes in support of an attack) in Section 3.

Some other frameworks have taken the direction of this paper, that is, to consider the weights on attacks, and in the remainder of this section, we discuss this work.

The idea of explicitly adding weights to attacks was proposed in [9]. However, the emphasis of that work was on how weights can be changed dynamically and over time, rather than on extending Dung's framework for determining extensions.

In another recent proposal for developing Dung's model, extra information representing the relative strength of attack is incorporated [32]. This is the only other approach that we are aware of which distinguishes the *strength of attack* from the strength of an argument. In this proposal, which we refer to as varied-strength attacks (or VSA) approach, each arc is assigned a type, and there is a partial ordering over the types. As a simple example, consider the following argument graph conforming to Dung's proposal, where α_1 is attacked by α_2 which in turn is attacked by α_3 .

$$\alpha_3 \rightarrow \alpha_2 \rightarrow \alpha_1$$

Here, α_3 defends the attack on α_1 , and as a result $\{\alpha_3, \alpha_1\}$ is the preferred and grounded extension. Now, consider the following VSA version of the graph, where the attack by α_3 is of type i and the attack by α_2 is of type j .

$$\alpha_3 \rightarrow_i \alpha_2 \rightarrow_j \alpha_1$$

This gives us a finer grained range of defence depending on whether type j is higher, or lower, or equally, ranked than type i , or incomparable with it. Furthermore, this allows for a finer definition of acceptable extension that specifies the required level of the defence of any argument in the extension. For instance, it can be insisted in the VSA approach that every defence of an argument should be by an attack that is stronger, so

in the above graph that would mean that the type of \rightarrow_i needs to be stronger than the type of \rightarrow_j in order for $\{\alpha_3, \alpha_1\}$ to be the preferred, grounded extension.

Finally, an important mechanism that attempts to unify approaches based on disregarding attacks are the *extended argumentation frameworks* (EAFs) of Modgil [36, 35]. An EAF is specified as a triple $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$ where $\langle \mathcal{X}, \mathcal{A} \rangle$ describes a standard Dung-style argument system and $\mathcal{D} \subseteq \mathcal{X} \times \mathcal{A}$, i.e., $\langle \alpha_1, \langle \alpha_2, \alpha_3 \rangle \rangle \in \mathcal{D}$ indicates that the *argument* α_1 attacks the *attack* $\langle \alpha_2, \alpha_3 \rangle$ thus those agents endorsing α_1 may disregard the attack on α_3 by α_2 . Modgil [36] describes how EAFs may be configured to capture acceptability in both PAFs and VAFs. While the basic structure underpinning the semantics of EAFs elaborates Dung’s grounded extensions (as is also the case w.r.t. VAFs and PAFs) there are a number of technical subtleties arising in the formal definition. In Section 7 we consider distinguishing aspects of our approach compared to preference-based, extended argumentation frameworks and other methods in greater depth.

From these proposals for developing Dung’s original approach, there is a common theme that arguments, or attacks by arguments, have variable strength. Some of these proposals are restricted to determining that strength is based on the other arguments available in the graph, together with their connectivity, and so the strength of an argument is a function solely of the graph. Others, in particular the VAF approach [10] and the VSA approach [32], use explicit ranking information over the arguments or the attacks by arguments. This ranking information requires extra information to be given along with the set of arguments and the attack relation. So, whilst there is gathering momentum for representing and reasoning with the strength of arguments or their attacks, there is not a consensus on the exact notion of argument strength or how it should be used. Furthermore, for the explicit representation of extra information pertaining to argument strength, we see that the use of explicit numerical weights is under-developed. So, for these reasons, we would like to present weighted argument systems as a valuable new proposal that should further extend and clarify aspects of this trend towards considering strength, in particular the explicit consideration of strength of attack between arguments..

In some respects, the strength of attack approach offers a richer representation than the strength of argument approach: For a connected graph with n arguments, we have up to n^2 strength of attack values, whereas we will only have n strength of argument values. This allows for more finer-grained consideration of argument graphs since we are considering more information.

However, in a sense, strength of attack and strength of argument are different kinds of information. For example, consider arguments α and α' that attack each other, perhaps as part of a larger argument graph. By using strength of argument values, we effectively decide whether one or both arguments needs to be considered as unacceptable. Given that they attack each other, we cannot have an acceptable set of arguments containing both α and α' . In contrast, by using strength of attack, we can render the strength of attack so weak that neither argument forces the other argument to be unacceptable.

Another advantage of using the strength of attack approach is that it can be considered locally. For instance, if arguments α and α' do attack each other, and they are part of a larger argument graph, and we wanted α to be in an acceptable extension in preference to α' , then all we need to do is set the strength of attack to be greater for α against α' than α' against α . This is a local decision, and it is not affected by the values for any other attack. In contrast, if we use strength of argument, then we would choose a value for α to be stronger than α' , but this is a global decision, because we would need to check that the value for α was adjusted so that it defeated or was defeated as required by all the other arguments it was connected to.

3. Weighted Argument Systems

We now introduce our model of weighted argument systems, and the key solutions we use throughout the remainder of the paper. Weighted argument systems extend Dung-style abstract argument systems by adding numeric weights to every edge in the attack graph, intuitively corresponding to the strength of the attack, or equivalently, how reluctant we would be to disregard it. Formally,

Definition 4. A weighted argument system is a triple $W = \langle \mathcal{X}, \mathcal{A}, w \rangle$ where $\langle \mathcal{X}, \mathcal{A} \rangle$ is a Dung-style abstract argument system, and $w : \mathcal{A} \rightarrow \mathbb{R}_{>}$ is a function assigning real valued weights³ to attacks.

In what follows, when we say simply “argument system”, we mean “Dung-style (unweighted) abstract argument system”.

Notice that we require attacks to have a positive *non-zero* weight. There may be cases where it is interesting to allow zero-weight attacks, in which case some of the analysis of this paper does not go through. However, given our intuitive reading of weights (that they indicate the strength of an attack) allowing 0-weight attacks is perhaps counter-intuitive. Suppose by appealing to a particular 0-weight attack you were able to support some particular argument, then an opponent could discard the attack *at no cost*. So, we will assume attacks must have non-zero weight. We postpone discussion of where weights come from until section 3.2.

3.1. Inconsistency Budgets

A key idea in what follows is that of an *inconsistency budget*, $\beta \in \mathbb{R}_{\geq}$, which we use to characterise *how much inconsistency we are prepared to tolerate*. The intended interpretation is that, given an inconsistency budget β , we would be prepared to *disregard attacks up to a total weight of β* . Conventional abstract argument systems implicitly assume an inconsistency budget of 0. However, by relaxing this constraint, allowing larger inconsistency budgets, we can obtain progressively more solutions from an argument system.

³We let $\mathbb{R}_{>}$ denote the real numbers greater than 0, and \mathbb{R}_{\geq} denote the real numbers greater than or equal to 0.




Figure 2: Weighted argument system W_1 from Example 1.

Definition 5. Let $\langle \mathcal{X}, \mathcal{A}, w \rangle$ be a weighted argument system. Given $R \subseteq \mathcal{A}$,

$$wt(R, w) = \sum_{\langle \alpha_1, \alpha_2 \rangle \in R} w(\langle \alpha_1, \alpha_2 \rangle)$$

The function $sub(\cdot \cdot \cdot)$, which takes an attack relation \mathcal{A} , weight function $w : \mathcal{A} \rightarrow \mathbb{R}_{>}$, and inconsistency budget $\beta \in \mathbb{R}_{\geq}$, returns the set of subsets R of \mathcal{A} whose total weight does not exceed β , i.e.,

$$sub(\mathcal{A}, w, \beta) = \{R : R \subseteq \mathcal{A} \ \& \ wt(R, w) \leq \beta\}.$$

We now use inconsistency budgets to introduce weighted variants of the semantics introduced for abstract argument systems, above.

Definition 6. Given a weighted argument system $\langle \mathcal{X}, \mathcal{A}, w \rangle$, let $\sigma : 2^{\mathcal{X}} \rightarrow \{\top, \perp\}$. For $\beta \in \mathbb{R}_{\geq}$, the subset $\mathcal{E}_{\sigma}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$ of $2^{\mathcal{X}}$ is given as

$$\mathcal{E}_{\sigma}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta) = \{S \subseteq \mathcal{X} : \exists R \in sub(\mathcal{A}, w, \beta) \ \& \ S \in \mathcal{E}_{\sigma}(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)\}$$

A set $S \in \mathcal{E}_{\sigma}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$ will, subsequently, be termed to as a $\beta - \sigma$ set (extension), so that we refer to β -grounded extensions, β -admissible sets, β -preferred extensions etc.

It is immediate from this definition that for every weighted argument system, irrespective of how σ is defined $\mathcal{E}_{\sigma}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, 0) = \mathcal{E}_{\sigma}(\langle \mathcal{X}, \mathcal{A} \rangle)$, i.e., by not allowing any inconsistency, weighted argument systems recover Dung argument system and their associated semantics.

Example 1. Consider the weighted argument system W_1 , illustrated in Figure 2. The initial set of arguments in W_1 is empty; however, $\{\alpha_5\} \in \mathcal{E}_{init}^{WT}(W_1, 1)$ since we can delete (α_4, α_5) with $\beta = 1$. If $\beta = 2$, we have $\mathcal{E}_{init}^{WT}(W_1, 2) = \{\emptyset, \{\alpha_4\}, \{\alpha_5\}\}$. Table 1 shows initial sets (and other examples) for some increasing values of β .

Weighted argument systems have the following, (readily proved), property:

β	$\mathcal{E}_{ini}^{WT}(W_1, \beta)$	$\mathcal{E}_{pr}^{WT}(W_1, \beta)$	$\mathcal{E}_{ge}^{WT}(W_1, \beta)$
0	$\{\emptyset\}$	$\{\{\alpha_1, \alpha_2, \alpha_4, \alpha_6\}, \{\alpha_3, \alpha_5, \alpha_7, \alpha_8\}\}$	$\{\emptyset\}$
1	$\{\emptyset, \{\alpha_5\}\}$	$\{\{\alpha_1, \alpha_2, \alpha_4, \alpha_6\}, \{\alpha_3, \alpha_5, \alpha_7, \alpha_8\}\}$	$\{\emptyset, \{\alpha_3, \alpha_5, \alpha_7, \alpha_8\}\}$
2	$\{\emptyset, \{\alpha_4\}, \{\alpha_5\}\}$	$\{\{\alpha_1, \alpha_2, \alpha_4, \alpha_6\}, \{\alpha_3, \alpha_5, \alpha_7, \alpha_8\}\}$	$\{\emptyset, \{\alpha_3, \alpha_5, \alpha_7, \alpha_8\}, \{\alpha_1, \alpha_2, \alpha_4, \alpha_6\}\}$
3	$\{\emptyset, \{\alpha_4\}, \{\alpha_5\}, \{\alpha_4, \alpha_5\}\}$	$\{\{\alpha_1, \alpha_2, \alpha_4, \alpha_6\}, \{\alpha_3, \alpha_5, \alpha_7, \alpha_8\}, \{\alpha_1, \alpha_2, \alpha_4, \alpha_5, \alpha_7, \alpha_8\}\}$	$\{\emptyset, \{\alpha_3, \alpha_5, \alpha_7, \alpha_8\}, \{\alpha_1, \alpha_2, \alpha_4, \alpha_6\}, \{\alpha_1, \alpha_2, \alpha_4, \alpha_5, \alpha_7, \alpha_8\}\}$

Table 1: $\mathcal{E}_\sigma^{WT}(W_1, \beta)$ for some increasing values of β .

Proposition 1. *Let $\langle \mathcal{X}, \mathcal{A}, w \rangle$ be a weighted abstract argument system. For every set of arguments $S \subseteq \mathcal{X}$, $\exists \beta \in \mathbb{R}_{\geq}$ and $T \subseteq \mathcal{X}$ such that $S \subseteq T$ and $T \in \mathcal{E}_{ini}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$.*

Thus, intuitively, every set of arguments can be made initial at some cost, and the cost required to do this immediately gives us a preference ordering over sets of arguments: we prefer sets of arguments that require a smaller budget. Notice that a similar observation holds true for conflict-freeness, admissibility, preferred extensions, credulous acceptance, and sceptical acceptance.⁴

Now, consider how grounded extensions are generalised within weighted systems. The first observation to make is that while in unweighted argument systems the grounded extension is unique, this will not necessarily be the case in weighted argument systems: *in weighted systems there may be many β -grounded extensions*, i.e., while for every $\langle \mathcal{X}, \mathcal{A} \rangle$ we have $|\mathcal{E}_{gr}(\langle \mathcal{X}, \mathcal{A} \rangle)| = 1$ for unweighted systems and $\beta > 0$ it is possible that $|\mathcal{E}_{gr}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)| > 1$. We observe that other developments building on standard argument systems, most notably the resolution-based grounded semantics of Baroni and Giacomin [7] also result in a multiple status variant. Table 1 shows β -grounded extensions for some increasing values of β for system W_1 of Figure 2.

3.2. Where do Weights Come From?

Having introduced the basic framework of weighted argument systems, some obvious and important questions arise: What do weights *mean*? Where do they come from? If we want to represent a particular argumentation domain using weighted argument systems, how are we to compute the weights for this domain? We will not demand any specific interpretation of weights, and the technical treatment of weighted argument systems in this paper does not of course require any such interpretation, apart from assuming that the weights are real-valued and may be combined additively. However, from the point of view of motivation, it is important to consider this issue seriously (if only to convince the reader that weights are not a purely technical device). We here discuss three possible interpretations of weights on attacks: weights as measures

⁴It should be noted, however, that we may not be able always to guarantee $S \in \mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$ no matter how large a budget is used: only $S \subseteq T$ with $T \in \mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$.

of votes in support of attacks; weights as measures of the inconsistency of argument-pairs; and weights as rankings of different types of attack. We emphasize that these three examples by no means exhaust the possibilities for the meaning of weights on attacks.

Weighted Majority Relations

In a multi-agent setting, one natural interpretation is that a weight represents *the number of votes in support of the attack*. This interpretation makes a link between argumentation and *social choice theory* — the theory of voting systems and collective decision making [3, 45]. Related work in argumentation drawing on social choice theory is [50], which considers aggregation of multiple types of attack in argumentation frameworks.

The following example is based on the well-known *Condorcet paradox*, which concerns majority voting on three mutually exclusive options. The basic idea of the paradox is that, given a group of voters, it possible for a majority to vote for A over B , for a majority to vote for B over C , and for a majority to vote for C over A .

More formally, assume we have three mutually exclusive options to vote on, and a group of k voters. Let x be the proportion of voters who vote for A over B , y be the proportion of voters who vote for B over C , and z be the proportion of voters who vote for C over A , where $x, y, z > 0.5$. Hence, a majority vote for A over B , a majority vote for B over C , and a majority vote for C over A . Also assume that there are no abstentions (and therefore for each pair of options $O_i, O_j \in \{A, B, C\}$, if $O_i \neq O_j$, then each voter either votes for O_i over O_j or for O_j over O_i). We can represent this instance of the Condorcet paradox in the weighted graphs with arcs defined as follows.

- A attacks B with weight x ;
- B attacks A with weight $1 - x$;
- B attacks C with weight y ;
- C attacks B with weight $1 - y$;
- C attacks A with weight z ; and
- A attacks C with weight $1 - z$.

Using the inconsistency budget, it is simple and intuitive to get useful solutions. To illustrate, suppose also that $\text{Min}(x, y, z)$ is x , then we have that the minimum budget β for which we get a non-empty initial set is when β is $x + (1 - y) + (1 - z)$. For this, $\{B, C\}$ is β -initial, β -acceptable, etc.

Now suppose we have a fourth option D , but there is only information to compare it with B and C . Suppose a majority (represented by proportion v) vote for C over D , and a majority (represented by proportion w) vote for D over B . Hence, the argument graph now has the following extra edges:

- C attacks D with weight v ;
- D attacks C with weight $1 - v$;
- D attacks B with weight w ; and
- B attacks D with weight $1 - w$.

This extended graph can be regarded as capturing two instances of the Condorcet paradox. This more complex situation involves incompleteness concerning the relative preference of B and D , and conflicts arising from the two instances of the paradox, and the interactions between the two instances.

Using inconsistency budgets, as before, we have a straightforward way to resolve the situation. To illustrate, suppose that $Min(v, w, x, y, z)$ is x , then we have that the minimum budget β for which we get a non-empty initial set is when β is $(1 - v) + (1 - w) + x + (1 - y) + (1 - z)$, and for this, $\{B, C, D\}$ is β -initial, β -acceptable, etc.

Weights as Measures of Inconsistency

Because attacks are between two arguments, a simple interpretation is to use weights as a measure of the extent of inconsistency between pairs of arguments. In other words, a higher weight on the attack between two arguments denotes greater inconsistency between the arguments concerned than does a lower weight. For example, consider the following arguments that are based on classical predicate logic.

$$A_1 = \langle \{\forall x.p(x)\}, \forall x.p(x) \rangle$$

$$A_2 = \langle \{\neg\forall x.p(x)\}, \exists x.\neg p(x) \rangle$$

$$A_3 = \langle \{\neg\exists x.p(x)\}, \forall x.\neg p(x) \rangle$$

Here, both A_2 and A_3 are undercuts of A_1 since each has a claim that negates the support of A_1 . As with the propositional case, we can measure the degree of undercut in terms of the models: For arguments A and A' , where A undercuts A' , we can determine the distance between the nearest models of the claim of A and of the support of A' . However, for the first-order case, we need to consider first-order models. One option is to consider Herbrand models composed from a Herbrand universe. Here, suppose the Herbrand universe is of cardinality n ; then we may assign a degree of undercut for A_2 on A_1 of $1/n$, whereas we may assign a degree of undercut for A_3 on A_1 of n . In the former case, it is because just one element of the universe is involved in a contradiction concerning the predicate p , whereas in the latter case, it is because all elements of the universe are involved in a contradiction concerning the predicate p . There is a range of options for defining a measure of the degree of undercut (see Besnard and Hunter [13] for more discussion of this), and the measures selected can be used in assessing the incoherence arising in a set of arguments. Moreover, we see

that the degree of undercut is an evaluated notion associated with the attack (or pair of arguments) and not just with individual arguments.

A similar notion is explored in work by Pinyol and Sabater-Mir [42], using a weighted argumentation framework for multi-agent communications concerning reputation. Here the weight of an attack between two arguments measures the relative reliability, as assessed by a decision-maker, of the sources of the two arguments.

Weights as Ranking of Attacks

Attacks between arguments may be of different types, or may be expressions of different perspectives (perhaps by different people) over the sets of arguments. In these cases, we can use weights to rank the relative strength of the different attacks between arguments. Here, higher weight denotes a stronger attack, but the absolute weight assigned to an attack is not important, just the relative weight compared to the weights assigned to other attacks. In this interpretation, the weights may arise from objective considerations (such as logical measures of the inconsistency between pairs of arguments, as in the previous interpretation), or from subjective criteria, as when the attacks are assigned by different stakeholders. To illustrate this approach, we consider an example modified from Example 1 of [50]. Three arguments are presented for a patient undergoing medical diagnosis, as follows:

- A:** Symptoms x, y and z indicate the presence of disease d_1 , which would suggest therapy t_1 .
- B:** Symptoms x, w and z indicate the presence of disease d_2 , which would suggest therapy t_2 .
- C:** Symptoms x and z indicate the presence of disease d_3 , which would suggest therapy t_3 .

In comparing one argument with another, a number of different criteria may be applied. In Example 1 of [50], for instance, these arguments are compared in terms of the specificity of the argument, whether or not the two treatments may be applied jointly, and the degree to which the underlying symptoms are present. One could readily imagine additional criteria, such as the relative costs of the recommended therapies, their relative durations, and the relative extent of side effects. Each of these criteria could be used to define a relation between pairs of arguments, which, applied here, identifies attacks between arguments. For each comparison criterion, appropriate weights could be defined by a decision-maker for the argument pairs where an attack is present. For instance, consider the relative strength of symptoms as a criterion for comparison. Suppose symptoms w and z were strongly present in the patient, while the other two symptoms were only weakly present. The attack of argument B on argument A could then be assigned a weight of 2, since two of B's three required symptoms are strongly present while only 1 of A's three required symptoms is. Similarly, the attack of argument B on C could be given a weight of 1.34, since two of B's three required symptoms are strongly present while only one of C's two required symptoms is.

In a similar way, we could define weights on the attacks between arguments based on the other identified comparison criteria. The relative importance of these different types of attacks could then be denoted by relative weights on the corresponding attacks, as assessed by a decision-maker. For example, if the criterion of the relative presence of symptoms was considered a more important comparison criterion than (say) the relative costs of the recommend therapies, then attacks of the first type could be assigned a weight greater than those of the second type. As mentioned, here the absolute numbers are not important; rather the weights are aiming to capture the relative degree of importance of the underlying comparison criteria to a decision-maker. One could readily imagine applications where attacks are identified and ranked by multiple stakeholders, such as the three doctors of Example 1 in [50], each stakeholder comparing argument pairs according to their own chosen set of criteria.

An objection to this example may be that applications such as this could be represented by defining multiple preference relations over the set of arguments. As we demonstrate in Section 7, such a conclusion, although true in particular cases, is not warranted in general.

4. Complexity of Weighted Solutions

Prima facie, it appears that weighted argument systems offer some additional expressive power over unweighted argument systems. So, does this apparently additional power come with some additional computational cost? The β versions of the decision problems for admissibility, checking preferred extensions, sceptical, and credulous acceptance are in fact no harder (although of course no easier) than the corresponding unweighted decision problems — these results are easy to establish. However, the story for β -grounded extensions is more complicated, since there may be *multiple* β -grounded extensions. Since there are multiple β -grounded extensions, we can consider credulous and sceptical variations of the problem, as with preferred extensions.

Let us make this discussion more formal. Where D is an arbitrary decision problem, instances of which include a (standard) Dung-style argument system $\langle \mathcal{X}, \mathcal{A} \rangle$, we use WD to denote the same problem but with *weighted* systems $\langle \mathcal{X}, \mathcal{A}, w \rangle$. Thus, suppose $\sigma : 2^{\mathcal{X}} \rightarrow \{\top, \perp\}$ is some argumentation semantics. In [26] the “canonical” decision problems in argument systems (each of which may be instantiated with respect to particular σ) are defined as in Table 2.

Consider the cases of verifying whether a given set S is a β - σ set, i.e., the decision problem $WVER_{\sigma}(\langle \mathcal{X}, \mathcal{A}, w \rangle, S, \beta)$.

Proposition 2. *The decision problem $WVER_{\sigma}$ with instances $\langle \mathcal{X}, \mathcal{A}, w \rangle$, $S \subseteq \mathcal{X}$ and $\beta \in \mathbb{R}_{>}$ is polynomial time time decidable for $\sigma \in \{\text{init}, \text{adm}\}$.*

Table 2: Standard Decision Problems in Argument Systems.

Problem	Instance	Question
VER_σ	$\langle \mathcal{X}, \mathcal{A} \rangle; S \subseteq \mathcal{X}$	Is $S \in \mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle)$?
CA_σ	$\langle \mathcal{X}, \mathcal{A} \rangle; x \in \mathcal{X}$	Is there <i>any</i> $S \in \mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle)$ for which $x \in S$?
SA_σ	$\langle \mathcal{X}, \mathcal{A} \rangle; x \in \mathcal{X}$	Is x a member of <i>every</i> $T \in \mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle)$?
EX_σ	$\langle \mathcal{X}, \mathcal{A} \rangle$	Is $\mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle)$ <i>non-empty</i> ?
NE_σ	$\langle \mathcal{X}, \mathcal{A} \rangle$	Is there any $S \in \mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle)$ for which $S \neq \emptyset$?

Proof. The case for $\text{wVER}_{\text{init}}$ is obvious, so consider wVER_{adm} . For S to be admissible in $\langle \mathcal{X}, \mathcal{A} \rangle$, it must be conflict-free with each argument in S acceptable to S . So, we need to eliminate from \mathcal{A} any attacks between members of S , and any attacks on members of S from outside S that are not defended by members of S ; so simply compute these two sets of attacks, and check whether the sum of their weights is $\leq \beta$. \square

The cases for credulous and sceptical acceptance with respect to preferred extensions — i.e., $\text{wCA}_{\text{pr}}(\langle \langle \mathcal{X}, \mathcal{A}, w \rangle, x, \beta \rangle)$ and $\text{wSA}_{\text{pr}}(\langle \langle \mathcal{X}, \mathcal{A}, w \rangle, x, \beta \rangle)$ — require a little more thought. We can show that, as with checking admissibility, the analogous weighted problems, while hard, are no harder than the unweighted solutions.

Proposition 3.

- a. wCA_{pr} is NP-complete.
- b. wSA_{pr} is Π_2^p -complete.

Proof. Consider credulous acceptance and an instance $\langle \langle \mathcal{X}, \mathcal{A}, w \rangle, x, \beta \rangle$. For membership of NP, simply guess some $R \subseteq \mathcal{A}$, and $S \subseteq \mathcal{X} \setminus \{x\}$. Then verify that both $\sum_{\langle x, y \rangle \in R} w(\langle x, y \rangle) \leq \beta$ and that $S \cup \{x\}$ is admissible in $\mathcal{A} \setminus R$ (we do not need to check maximality: if an argument is contained in some β -admissible set, it is contained in some maximal β -admissible set). For hardness, observe that wCA_{pr} generalises CA_{pr} i.e., the unweighted version which is NP-complete [17]. Formally, given an instance $\langle \langle \mathcal{X}, \mathcal{A} \rangle, x \rangle$ of CA_{pr} create a weighted argument system $\langle \mathcal{X}, \mathcal{A}, w \rangle$ in which $w(\langle p, q \rangle) = 1$ for all $\langle p, q \rangle \in \mathcal{A}$, and ask whether $\langle \langle \mathcal{X}, \mathcal{A}, w \rangle, x, 0 \rangle$ is accepted as an instance of wCA_{pr} . The proof for wSA_{pr} is similar using the Π_2^p -hardness of SA_{pr} demonstrated in [23]. \square

Our principal interest is in the grounded extension, and so we now focus on the problem wCA_{gr} , instances of which are of the form $\langle \langle \mathcal{X}, \mathcal{A}, w \rangle, \beta, x \rangle$ where β is an inconsistency budget and $x \in \mathcal{X}$. An instance of wCA_{gr} is accepted if $\exists S \in \mathcal{E}_{\text{gr}}^{\text{WT}}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$ such that $x \in S$. The analogous *sceptical* form of wCA_{gr} (denoted by wSA_{gr}) has instances of the same form, but with these accepted if and only if $\forall S \in \mathcal{E}_{\text{gr}}^{\text{WT}}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$, we have $x \in S$.

Proposition 4. wCA_{gr} is NP-complete, and remains NP-complete even if the attack relation is planar and/or tripartite and/or has no argument which is attacked by more than two others.

Proof. For membership, a conventional “guess and check” approach suffices. For NP-hardness, we reduce from 3-SAT. Given an instance $\varphi(Z_n)$ of 3-SAT with m clauses C_j over propositional variables $Z_n = \{z_1, \dots, z_n\}$, form the weighted argument system $\langle X_\varphi, \mathcal{A}_\varphi, w_\varphi \rangle$, illustrated in Figure 3. Specifically, X_φ has $3n + m + 1$ arguments: an argument C_j for each clause of $\varphi(Z_n)$; arguments $\{z_i, \neg z_i, u_i\}$ for each variable of Z_n , and an argument φ . The relationship, \mathcal{A}_φ , contains attacks (C_j, φ) for each clause of φ , $(z_i, \neg z_i)$, $(\neg z_i, z_i)$, (z_i, u_i) , $(\neg z_i, u_i)$, and (u_i, φ) for each $1 \leq i \leq n$. Finally, \mathcal{A}_φ contains an attack (z_i, C_j) if z_i is a literal in C_j , and $(\neg z_i, C_j)$ if $\neg z_i$ occurs in C_j . The weighting function w_φ assigns weight 1 to each of the attacks $\{(z_i, \neg z_i), (\neg z_i, z_i)\}$ and weight $n+1$ to all remaining attacks. To complete the instance the available budget is set to n and the argument of interest to φ . We claim that $\varphi \in S$ for some $S \in \mathcal{E}_{gr}^{WT}(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi, w_\varphi \rangle, n)$ if and only if $\varphi(Z_n)$ is satisfiable. We first note that φ is credulously accepted in the (unweighted) system $\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \rangle$ if and only if $\varphi(Z_n)$ is satisfiable.⁵ We deduce that if $\varphi(Z_n)$ is satisfied by an assignment $\langle a_1, a_2, \dots, a_n \rangle$ of Z_n then φ is a member of the grounded extension of the acyclic system $\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus B \rangle$ in which B contains $(\neg z_i, z_i)$ (if $a_i = \top$) and $(z_i, \neg z_i)$ (if $a_i = \perp$). Noting that B has total weight n , and that the subset $\{y_1, y_2, \dots, y_n\}$ in which $y_i = z_i$ (if $a_i = \top$) and $\neg z_i$ (if $a_i = \perp$) is unattacked, it follows that from $\varphi(Z_n)$ satisfiable we may identify a suitable weight n set of attacks, B , to yield $\varphi \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus B \rangle)$.

On the other hand, suppose that $\varphi \in S$ for some $S \in \mathcal{E}_{gr}^{WT}(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi, w_\varphi \rangle, n)$. Consider the set of attacks, B , eliminated from \mathcal{A}_φ in order to form the system $\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus B_\varphi \rangle$ with grounded extension S . Since $\varphi \in S$, exactly one of $(z_i, \neg z_i)$ and $(\neg z_i, z_i)$ must be in B for every i . Otherwise, if for some i , neither attack is in B then $\{z_i, \neg z_i\} \cap S = \emptyset$, and thus φ has no defence to the attack by u_i , contradicting $\varphi \in S$; similarly if both attacks are in B then, from the fact B has total weight at most n , for some other variable, z_k , both $(z_k, \neg z_k)$ and $(\neg z_k, z_k)$ would be in $\mathcal{A}_\varphi \setminus B$. In total, from $S \in \mathcal{E}_{gr}^{WT}(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi, w_\varphi \rangle, n)$ and $\varphi \in S$ for each $1 \leq i \leq n$ we identify exactly one unattacked argument, y_i from $\{z_i, \neg z_i\}$, so that $S = \{\varphi, y_1, \dots, y_n\}$. That the assignment $z_i = \top$ (if $y_i = z_i$) and $z_i = \perp$ (if $y_i = \neg z_i$) satisfies $\varphi(Z_n)$ is immediate from [17].

The remaining cases (for planar, tripartite graphs, etc.) can be derived from the reductions from 3-SAT given in [20]. \square

Now consider the “sceptical” version of this problem.

Proposition 5. *WSA_{gr} is coNP-complete.*

Proof. Membership of coNP is immediate from the algorithm which checks for every $B \subseteq A$ that if $\sum_{e \in B} w(e) \leq \beta$ then $x \in ge(\langle \mathcal{X}, \mathcal{A} \setminus B \rangle)$. For coNP-hardness, we use a reduction from UNSAT, assuming w.l.o.g. that the problem instance is presented in CNF. Given an m -clause instance $\varphi(Z_n)$ of UNSAT, we construct a weighted argument system $\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi, w_\varphi \rangle$ as follows. The set \mathcal{X}_φ contains $4n + m + 3$ arguments: $\{\varphi, \psi, \chi\}$;

⁵This follows from [17] which uses a similar construction for which the u_i arguments and associated attacks do not occur.




Figure 3: The reduction used in Proposition 4.

$\{z_i, \neg z_i, u_i, v_i : 1 \leq i \leq n\}$; and $\{C_j : 1 \leq j \leq m\}$. The attack set \mathcal{A}_φ comprises: $\{(\varphi, \psi), (\chi, \varphi)\}$; $\{(v_i, z_i), (v_i, \neg z_i), (z_i, u_i), (\neg z_i, u_i), (u_i, \varphi)\}$ for each $1 \leq i \leq n$; $\{(C_j, \varphi) : 1 \leq j \leq m\}$; $\{(z_i, C_j) : z_i \in C_j\}$ and $\{(\neg z_i, C_j) : \neg z_i \in C_j\}$. The attacks are weighted so that $w_\varphi((\chi, \varphi)) = 1$; $w_\varphi((v_i, z_i)) = w_\varphi((v_i, \neg z_i)) = 1$; all remaining attacks have weight $n+2$. The instance is completed using ψ as the relevant argument and an inconsistency tolerance of $n+1$. See Figure 4 for an illustration of the construction.

Now, suppose that $\varphi(Z_n)$ is satisfied by an assignment $\alpha = \langle a_1, \dots, a_n \rangle$ of Z_n . Consider the subset B_α of \mathcal{A}_φ given by $\{(\chi, \varphi)\} \cup \{(v_i, z_i) : a_i = \top\} \cup \{(v_i, \neg z_i) : a_i = \perp\}$. The weight of B_α is $n+1$ and since α satisfies $\varphi(Z_n)$ it follows that $ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus B_\alpha \rangle)$ contains exactly the arguments $\{\chi, \varphi\} \cup \{v_1, \dots, v_n\} \cup \{z_i : a_i = \top\} \cup \{\neg z_i : a_i = \perp\}$. Hence $\psi \notin ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus B_\alpha \rangle)$ as required.

Conversely, suppose $\langle \langle \mathcal{X}_\varphi, \mathcal{A}_\varphi, w \rangle, \psi, n+1 \rangle$ is *not* accepted. We show that we may construct a satisfying assignment of $\varphi(Z_n)$ in such cases. Consider $B \subseteq \mathcal{A}_\varphi$ of weight at most $n+1$ for which $\psi \notin ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus B \rangle)$. It must be the case that $(\chi, \varphi) \in B$ for otherwise the attack by φ on ψ is defended so that ψ would belong to the grounded extension. The remaining elements of B must form a subset of the attacks $\{(v_i, z_i), (v_i, \neg z_i)\}$ (since all remaining attacks are too costly). Furthermore, exactly one of $\{(v_i, z_i), (v_i, \neg z_i)\}$ must belong to B for each $1 \leq i \leq n$: otherwise, some u_i will be in $ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus B \rangle)$, thus providing a defence to the attack on ψ by φ and contradicting the assumption $\psi \notin ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus B \rangle)$. Now consider the assignment α , with $a_i = \top$ if $(v_i, z_i) \in B$, $a_i = \perp$ if $(v_i, \neg z_i) \in B$. We now see that α must satisfy $\varphi(Z_n)$: in order for $\psi \notin ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus B \rangle)$ to hold, it must be the case that $\varphi \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus B \rangle)$, i.e., each of the C_j attacks on φ must be counterattacked by one of its constituent literal (arguments) y_i . Noting that v_i is always in $ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus B \rangle)$, if $a_i = \top$




Figure 4: The reduction used in Proposition 5.

clauses containing $\neg z_i$ cannot be attacked (since the attack $(v_i, \neg z_i)$ is still present). It follows that the assignment, α , attacks each clause so that $\varphi \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus B \rangle)$. In sum, if $\langle \langle \mathcal{X}_\varphi, \mathcal{A}_\varphi, w \rangle, \psi, n+1 \rangle$ is accepted as an instance of WSA_{gr} then $\varphi(Z_n)$ is unsatisfiable, so completing the proof. \square

Note that in some cases, considering sceptical grounded extensions is of limited value. Consider the set $initial(\langle \mathcal{X}, \mathcal{A} \rangle)$ then we have:

Proposition 6. *Let $\langle \mathcal{X}, \mathcal{A}, w \rangle$ be a weighted argument system and β be an inconsistency budget. Then $initial(\langle \mathcal{X}, \mathcal{A} \rangle) \neq \emptyset$ iff $(\bigcap_{Y \in \mathcal{E}_{gr}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)} Y) \neq \emptyset$.*

In some respects the intractability status of WCA_{gr} and WSA_{gr} in the general case, as illustrated by the results of Propositions 4 and 5 is unsurprising: intractability results hold for all proposed so-called multiple status semantics in Dung’s model (e.g., preferred extensions via the results of [17], semi-stable semantics [24, 27]). Within variants of standard argument systems whose basic semantics of interest is multiple status form of the grounded extensions, e.g., the resolution-based grounded semantics, the corresponding credulous and sceptical acceptance problems have also been shown to be intractable [6]. Even within “single-status” semantics (where there is a unique subset of arguments satisfying the specified condition), there are examples where the

credulous acceptance problem⁶ is intractable (e.g., ideal semantics [21, 22]).

There are, however, a number of cases for which these decision problems *can* be solved in polynomial time, in particular if the topology of the system $\langle \mathcal{X}, \mathcal{A} \rangle$ satisfies given constraints. Thus, all of the the decision problems in Table 2 are polynomial time decidable in preferred semantics in the following cases:

- a. If $\langle \mathcal{X}, \mathcal{A} \rangle$ is *acyclic* (as shown in Dung [19]).
- b. If $\langle \mathcal{X}, \mathcal{A} \rangle$ is *symmetric*, i.e., if for every (distinct) $x, y \in \mathcal{X}$, $\langle x, y \rangle \in \mathcal{A} \Leftrightarrow \langle y, x \rangle \in \mathcal{A}$, (Coste-Marquis et al. [16]).
- c. If $\langle \mathcal{X}, \mathcal{A} \rangle$ is *bipartite*, i.e., \mathcal{X} may be partitioned into two sets \mathcal{Y} and \mathcal{Z} , both of which are conflict-free. (Dunne [20]).

Given this, we might hope that if our weighted argument system falls into one of these three classes then the related decision questions have feasible algorithms. In fact, for two of these cases, this hope is disappointed.

Proposition 7. *WCA_{gr} is NP-complete even if $\langle \mathcal{X}, \mathcal{A} \rangle$ is acyclic.*

Proof. To prove NP-hardness, we use a reduction from 3-SAT. Given an instance $\varphi(Z_n)$ of 3-SAT form the weighted argument system $\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi, w \rangle$ with

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

The weighting function has

$$w(\langle x, y \rangle) = \begin{cases} 1 & \text{if } \langle x, y \rangle = \langle \neg z_i, z_i \rangle \\ n + 1 & \text{otherwise} \end{cases}$$

The environment for a typical variable x of the CNF φ is shown in Figure 5. The full instance is $\langle \langle \mathcal{X}_\varphi, \mathcal{A}_\varphi, w \rangle, n, \varphi \rangle$. Noting that $\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \rangle$ is, indeed, acyclic, we claim this instance is accepted if and only if $\varphi(Z_n)$ is accepted as an instance of 3-SAT.

First suppose that there is some $\alpha \in \{\top, \perp\}^n$ that satisfies $\varphi(Z_n)$. Consider the subset R_α of \mathcal{A}_φ given by $\{ \langle \neg z_i, z_i \rangle : \alpha_i = \top \}$. Then $wt(R, w) \leq n$ and, since α satisfies φ , for each clause C_j we can identify a literal $y_j \in C_j$ that is set to \top under α . We have two cases:

Case 1: $y_j = z_i$. Then the argument $z_i \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R_\alpha \rangle)$ and, hence, $C_j \notin ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R_\alpha \rangle)$.

⁶In single status semantics this, of course, is equivalent to sceptical acceptance.




Figure 5: Acyclic variable setting device in reduction from 3 – SAT.

Case 2: $y_j = \neg z_i$. In this case, $z_i \notin ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R_\alpha \rangle)$ and $v_i \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R_\alpha \rangle)$ (from the fact that $\neg z_i$ is *always* in $ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R_\alpha \rangle)$, however, the attack $\langle \neg z_i, z_i \rangle$ is now in $\mathcal{A}_\varphi \setminus R_\alpha$ so precluding z_i). Since $\langle v_i, C_j \rangle \in \mathcal{A}_\varphi \setminus R_\alpha$ it follows that $C_j \notin ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R_\alpha \rangle)$.

Thus, for each C_j , $C_j \notin ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R_\alpha \rangle)$ and it follows that $\varphi \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R_\alpha \rangle)$ as required.

On the other hand suppose there is some $R \subseteq \mathcal{A}_\varphi$ for which $wt(R, w) \leq n$ and $\varphi \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$. It must be the case that $R \subseteq \{ \langle \neg z_i, z_i \rangle : 1 \leq i \leq n \}$ for otherwise $wt(R, w) > n$. From $\varphi \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$, it follows that $\forall C_j, C_j \notin ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$ and hence $u_j \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$ for some $\langle u_j, C_j \rangle \in \mathcal{A}_\varphi \setminus R$. Again we have two cases,

Case 1: $u_j = v_i$. In which case $z_i \notin ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$ (otherwise $\langle z_i, v_i \rangle \in \mathcal{A}_\varphi \setminus R$).

Case 2: $u_j = z_i$. In which event we must have $\langle \neg z_i, z_i \rangle \in R$.

Now, choosing the assignment α_R via $\alpha_i := \top$ if $z_i \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$ and $\alpha_i := \perp$ if $v_i \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$ then α_R satisfies $\varphi(Z_n)$ (every clause has at least one literal assigned \top). \square

As a technical aside, note that the only previously known case for which acyclic systems prove to be intractable (in fact, with the specific sub-class of *binary trees*) are for the canonical decision problems on VAFs.

Corollary 1. WCA_{gr} is NP-complete even when $\langle \mathcal{X}, \mathcal{A} \rangle$ is required to satisfy all of the following properties: $\langle \mathcal{X}, \mathcal{A} \rangle$ is acyclic; $\langle \mathcal{X}, \mathcal{A} \rangle$ is tripartite; $\langle \mathcal{X}, \mathcal{A} \rangle$ is planar; every $x \in \mathcal{X}$ is attacked by and attacks at most two other arguments in \mathcal{X} .

Proof. The system constructed in Proposition 7 is tripartite — use three colours $\{B, W, R\}$ and assign φ, z_j arguments the colour B ; C_j arguments the colour W ; v_j and $\neg z_j$ the colour R . The extension to planar and bounded number of attacks follows from the construction of Dunne [20]. \square

We obtain a similar result for the sceptical version, so that as with Proposition 5

Proposition 8. WSA_{gr} is coNP-complete even if $\langle \mathcal{X}, \mathcal{A} \rangle$ is acyclic.

Proof. The coNP-hardness proof uses a near identical construction to that of Proposition 7: given an instance $\varphi(Z_n)$ of *unsatisfiability*, form the weighted system $\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi, w \rangle$ described in the proof of Proposition 7 adding an additional argument, ψ , whose sole attacker is φ . Fix $w(\langle \varphi, \psi \rangle) = n + 1$. Then it is easily seen that $\forall R \subseteq \mathcal{A} : wt(R, w) \leq n \Rightarrow \psi \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$ if and only if $\varphi(Z_n)$ is unsatisfiable. \square

By definition *symmetric* argument systems contain cycles. For these cases, however, we still have

Proposition 9. WCA_{gr} is NP-complete even if $\langle \mathcal{X}, \mathcal{A} \rangle$ is symmetric.

Proof. We use a reduction from 3-SAT restricted to instances $\varphi(Z_n)$ in which *every* variable z_i occurs in *exactly* three clauses of φ (without loss of generality, exactly two such occurrences are in positive form and exactly one occurrence in negated form). For a variable z_i let $C_{i,1}^+$ and $C_{i,2}^+$ denote the two clauses in which the literal z_i occurs and $C_{i,3}^-$ the (unique) clause in which $\neg z_i$ features. The “variable setting” component is based on that of Figure 5 and is shown in Figure 6. The full system, $\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi, w \rangle$, formed from $\varphi(Z_n)$ has arguments

$$\mathcal{X}_\varphi = \{z_i, \neg z_i : 1 \leq i \leq n\} \cup \{C_{i,1}^+, C_{i,2}^+, C_{i,3}^- : 1 \leq i \leq n\} \cup \{\varphi\} \cup \{v_i : 1 \leq i \leq n\}$$

and attacks

$$\begin{aligned} \mathcal{A}_\varphi = & \{ \langle C_{i,1}^+, \varphi \rangle, \langle C_{i,2}^+, \varphi \rangle, \langle C_{i,3}^-, \varphi \rangle : 1 \leq i \leq n \} \cup \\ & \{ \langle \varphi, C_{i,1}^+ \rangle, \langle \varphi, C_{i,2}^+ \rangle, \langle \varphi, C_{i,3}^- \rangle : 1 \leq i \leq n \} \cup \\ & \{ \langle \neg z_i, z_i \rangle, \langle z_i, v_i \rangle : 1 \leq i \leq n \} \cup \\ & \{ \langle z_i, \neg z_i \rangle, \langle v_i, z_i \rangle : 1 \leq i \leq n \} \cup \\ & \{ \langle v_i, C_{i,3}^- \rangle, \langle C_{i,3}^-, v_i \rangle : 1 \leq i \leq n \} \cup \\ & \{ \langle z_i, C_{i,1}^+ \rangle, \langle z_i, C_{i,2}^+ \rangle : 1 \leq i \leq n \} \cup \\ & \{ \langle C_{i,1}^+, z_i \rangle, \langle C_{i,2}^+, z_i \rangle : 1 \leq i \leq n \} \end{aligned}$$




Figure 6: Symmetric variable setting device in reduction from 3 – SAT.

Finally the weighting function has

$$w(\langle x, y \rangle) = \begin{cases} 0.25 & \text{if } \langle x, y \rangle \in \{ \langle \neg z_i, z_i \rangle, \langle v_i, z_i \rangle, \langle C_{i,j}^+, z_i \rangle \} \\ 0.5 & \text{if } \langle x, y \rangle \in \{ \langle z_i, \neg z_i \rangle, \langle C_{i,3}^-, v_i \rangle \} \\ n+1 & \text{otherwise} \end{cases}$$

We note that the system $\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \rangle$ is symmetric. We claim that $\varphi(Z_n)$ is satisfiable if and only $\langle \langle \mathcal{X}_\varphi, \mathcal{A}_\varphi, w \rangle, n, \varphi \rangle$ is accepted as an instance of WCA_{gr} .

Suppose that $\alpha = \langle a_1, a_2, \dots, a_n \rangle$ defines a satisfying assignment of φ . Set R_α to be the subset of \mathcal{A}_φ containing

$$\bigcup_{a_i=\top} \{ \langle \neg z_i, z_i \rangle, \langle C_{i,1}^+, z_i \rangle, \langle C_{i,2}^+, z_i \rangle, \langle v_i, z_i \rangle \} \cup \bigcup_{a_i=\perp} \{ \langle z_i, \neg z_i \rangle, \langle C_{i,3}^-, v_i \rangle \}$$

Then $wt(R_\alpha, w) = \sum_{a_i=\top} (4 \times 0.25) + \sum_{a_i=\perp} (2 \times 0.5) = n$. In addition, $\varphi \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R_\alpha \rangle)$. To see this, notice that since α satisfies φ , every clause C has some literal y assigned \top . As before consider the two possible cases:

Case 1: $y = \neg z_i$ (so that $C = C_{i,3}^-$)

Then $\neg z_i \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R_\alpha \rangle)$ since R_α contains all attacks on $\neg z_i$. Thus, $z_i \notin ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R_\alpha \rangle)$ and $v_i \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R_\alpha \rangle)$ (since R_α contains the attack $\langle C_{i,3}^-, v_i \rangle$). It follows that the attack by $C_{i,3}^-$ on φ is defended by v_i .

Case 2: $y = z_i$ (so that $C \in \{C_{i,1}^+, C_{i,2}^+\}$)

In this case $z_i \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R_\alpha \rangle)$ (since, again, R_α contains all attacks on z_i). Consequently the attacks by $C_{i,1}^+$ and $C_{i,2}^+$ on φ are defended by z_i .

In summary, every attacker of φ is counterattacked by some argument in $ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R_\alpha \rangle)$ and it follows that $\varphi \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R_\alpha \rangle)$ as a result.

On the other hand suppose that $R \subset \mathcal{A}_\varphi$ is such that $wt(R, w) \leq n$ and $\varphi \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$. Then for each C some attacker u of C must also be in $ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$. Consider the assignment $\langle a_1, \dots, a_n \rangle$ of Z_n defined by the two possibilities:

Case 1: $u = v_i$ (so that $v_i \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$ and $C = C_{i,3}^-$)

From $v_i \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$ and $wt(R, w) \leq n$ we have $\langle z_i, v_i \rangle \notin R$ thus $z_i \notin ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$. Furthermore, in order to defend the attack by z_i on v_i some attacker p of z_i is in $ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$. Certainly $p \notin \{C_{i,1}^+, C_{i,2}^+\}$ for this contradicts $\varphi \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$ or $wt(R, w) \leq n$. We must, therefore have $\neg z_i \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$ (so that $\langle z_i, \neg z_i \rangle \in R$). In this case we set $a_i = \perp$.

Case 2: $u = z_i$ (so that $z_i \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$ and $C \in \{C_{i,1}^+, C_{i,2}^+\}$) and we can set $a_i = \top$.

It is easily seen that this assignment is well-defined, i.e., no variable of Z_n is assigned both \top and \perp since we cannot have *both* $z_i \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$ and $v_i \in ge(\langle \mathcal{X}_\varphi, \mathcal{A}_\varphi \setminus R \rangle)$. In addition the assignment ensures that every clause of $\varphi(Z_n)$ has at least one true literal, i.e., $\varphi(Z_n)$ is satisfiable. \square

To the best of our knowledge, this is the first known intractability result for a canonical decision problem in symmetric systems. In VAFs, the main decision problems are trivial in symmetric systems on account of the normal “consistency” assumption applied.

Noting that $ge(\langle \mathcal{X}, \mathcal{A} \rangle) = \emptyset$ whenever $\langle \mathcal{X}, \mathcal{A} \rangle$ is symmetric, the problem WSA_{gr} is trivial in such cases (with an inconsistency budget of 0 the grounded extension will be empty). We cannot, therefore, obtain an exact analogue of Proposition 5 for symmetric systems. If, however, we consider an alternative formulation – $WSA_{gr}^>$ whose instances are $\langle \langle \mathcal{X}, \mathcal{A}, w \rangle, x, \beta, \gamma \rangle$ with $\gamma > \beta \geq 0$ accepted if $\forall R \subseteq \mathcal{A}$ ($\beta \leq wt(R, w) \leq \gamma \Rightarrow x \in ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$) holds, then

Proposition 10. $WSA_{gr}^>$ is coNP-complete even if $\langle \mathcal{X}, \mathcal{A} \rangle$ is symmetric.

5. How Much Inconsistency Do We Need?

An obvious question now arises: Suppose we have a weighted argument system $\langle \mathcal{X}, \mathcal{A}, w \rangle$ and a set of arguments S . Then *what is the smallest amount of inconsistency β we would need to tolerate in order to obtain $T \in \mathcal{E}_\sigma^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$ with $S \subseteq T$?* Now, when considering initial, conflict-free and admissible sets, the answer is easy: we know exactly which attacks we would have to disregard to make a set of arguments satisfy the specific criteria imposed – we have no choice in the matter. However, when considering grounded extensions, the answer is not so easy. As we saw above, there may be multiple ways of getting a set of arguments into a weighted extension, each with

potentially different costs; we are thus typically interested in solving the following problem:

$$\begin{aligned} &\text{minimise } \beta^* \text{ such that} \\ &\exists Y \in \mathcal{E}_{gr}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta^*) : S \subseteq Y \end{aligned} \quad (1)$$

What can we say about (1)? Consider the following problem. We are given a weighted argument system $\langle \mathcal{X}, \mathcal{A}, w \rangle$ and an inconsistency budget $\beta \in \mathbb{R}_{\geq}$, and asked whether β is *minimal*, i.e., whether $\forall \beta' < \beta$ and $\forall Y \in \mathcal{E}_{gr}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta')$, we have that $S \not\subseteq Y$. (This problem does *not* require that S is contained in a some β -grounded extension of $\langle \mathcal{X}, \mathcal{A}, w \rangle$.)

Proposition 11. *Given a weighted argument system $\langle \mathcal{X}, \mathcal{A}, w \rangle$, set of arguments $S \subseteq \mathcal{X}$, and inconsistency budget β , checking whether β is minimal w.r.t. $\langle \mathcal{X}, \mathcal{A}, w \rangle$ and S is coNP-complete.*

Proof. Consider the complement problem, i.e., the problem of checking whether $\exists \beta' < \beta$ and $\exists Y \in \mathcal{E}_{gr}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta')$ such that $S \subseteq Y$. Membership in NP is immediate. For NP-hardness, we can reduce SAT, using essentially the same construction for the weighted argument system as Proposition 7; we ask whether $n + 1$ is not minimal for argument set $\{\varphi\}$. \square

This leads very naturally to the following question: is β the *smallest* inconsistency budget required to ensure that S is contained in some β -grounded extension. We refer to this problem as *checking whether β is the minimal budget for S* .

Proposition 12. *Given a weighted argument system $\langle \mathcal{X}, \mathcal{A}, w \rangle$, set of arguments $S \subseteq \mathcal{X}$, and inconsistency budget β , checking whether β is the minimal budget for S is D^p -complete.*

Proof. For membership of D^p , we must exhibit two languages L_1 and L_2 such that $L_1 \in \text{NP}$, $L_2 \in \text{coNP}$, and $L_1 \cap L_2$ is the set of instances accepted by the minimal budget problem. Language L_1 is given by Proposition 7, while language L_2 is given by Proposition 11. For hardness, we reduce the Critical Variable Problem (CVP) — for a proof that CVP is D^p -complete, see, e.g., [11]. An instance of CVP is given by a propositional formula φ in CNF, and a variable z from φ . We are asked if, under the valuation $z = \top$ the formula φ is satisfiable, while under the valuation $z = \perp$ it is unsatisfiable. We proceed to create an instance of the minimal budget problem by using essentially the same construction as Proposition 7, except that the attack $(z, \neg z)$ is given a weight of 0.5. Now, in the resulting system, n is the minimal budget for $\{\varphi\}$ iff z is a critical variable in φ . \square

The problem considered in Proposition 12 only deals with the question of *deciding* if a specified budget is, indeed, minimal. In practice we would actually wish to *compute* what this minimal budget is or even to compute a subset of \mathcal{A} that achieves this minimal budget, i.e., given $\langle \mathcal{X}, \mathcal{A}, w \rangle$ and $x \in \mathcal{X}$ we have two function problems relating to minimal budget computations.

- a. MIN-BUDGET Determine the *value* of the function, $\beta\text{-opt}(\langle \mathcal{X}, \mathcal{A}, w \rangle, x)$ defined as,

$$\min \{k \in \mathbb{R}_{\geq} : \exists R \subseteq \mathcal{A} \text{ with } wt(R, w) \leq k \text{ and } x \in ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)\}$$

- b. MIN-WITNESS, Report *any* subset R of \mathcal{A} for which $wt(R, w)$ is minimal among those sets with $x \in ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$.

It should be noted that the function defined by MIN-WITNESS, which we denote by $\beta\text{-witness}$ will, typically, be *multi-valued* whereas that defined by MIN-BUDGET is *single valued*. In addition any *lower* bound on MIN-BUDGET will, trivially, be a lower bound on the complexity of $\beta\text{-witness}$ since given a witnessing subset, R , from the computation of $\beta\text{-witness}$, computing MIN-BUDGET simply requires the (polynomial time) calculation $\sum_{\langle p, q \rangle \in R} w(\langle p, q \rangle)$.

We now consider the complexity of these function problems.

Proposition 13. MIN-BUDGET is FP^{NP} -complete.

Proof. Observe that

$$0 \leq \beta\text{-opt}(\langle \mathcal{X}, \mathcal{A}, w \rangle, x) \leq \sum_{y : \langle y, x \rangle \in \mathcal{A}} w(\langle y, x \rangle) = wt(\mathcal{A}, w) \leq |\mathcal{A}| \times \max_{\langle y, z \rangle \in \mathcal{A}} w(\langle y, z \rangle)$$

Hence, if the instance $\langle \langle \mathcal{X}, \mathcal{A}, w \rangle, x \rangle$ is encoded using t bits, the value of $\beta\text{-opt}(\langle \mathcal{X}, \mathcal{A}, w \rangle, x)$ requires (at most) t^2 bits. We obtain an FP^{NP} method to compute $\beta\text{-opt}(\langle \mathcal{X}, \mathcal{A}, w \rangle, x)$ using a standard binary search technique involving $O(t^2)$ calls to an NP oracle for WCA_{gr} , e.g., as illustrated in Algorithm 1. To show MIN-BUDGET is FP^{NP} -hard, we use a reduction from the following function problem:

LEX-MIN SAT : Given $\varphi(Z_n)$ a CNF formula $L_{\min}(\varphi)$ reports the $n + 1$ bit value corresponding to

$$\begin{cases} 1^{n+1} & \text{if } \varphi(Z_n) \text{ is } \textit{unsatisfiable} \\ 0a_1a_2 \dots a_n & \text{where } a_1a_2 \dots a_n \text{ is the } \textit{lexicographically minimal} \text{ satisfying assignment of } \varphi. \end{cases}$$

[For distinct bit sequences $a_1a_2 \dots a_n$, $b_1b_2 \dots b_n$ we say $a_1 \dots a_n$ lexicographically precedes $b_1 \dots b_n$ if $a_1 = 0$ and $b_1 = 1$ or $(a_1 = b_1$ and $a_2 \dots a_n$ lexicographically precedes $b_2 \dots b_n$).]

The function L_{\min} has been shown to be FP^{NP} -hard by Krentel [31]. Given an instance $\varphi(Z_n)$ of L_{\min} form the argument system $\langle \mathcal{X}_{\varphi}, \mathcal{A}_{\varphi} \rangle$ of Proposition 7. We add two arguments $\{\chi, \psi\}$ to this together with attacks $\{\langle \varphi, \chi \rangle, \langle \chi, \psi \rangle\}$ letting $\mathcal{Y}_{\varphi} = \mathcal{X}_{\varphi} \cup \{\chi, \psi\}$ and $\mathcal{B}_{\varphi} = \mathcal{A}_{\varphi} \cup \{\langle \varphi, \chi \rangle, \langle \chi, \psi \rangle\}$. The weighting function in the instance is,

$$w(\langle x, y \rangle) = \begin{cases} 2^{i-1} & \text{if } \langle x, y \rangle = \langle \neg z_i, z_i \rangle \\ 2^n & \text{if } \langle x, y \rangle = \langle \chi, \psi \rangle \\ 2^n + 1 & \text{otherwise} \end{cases}$$

Algorithm 1 Computing $\text{MIN-BUDGET}(\langle \mathcal{X}, \mathcal{A}, w \rangle, x)$.

```

1: function SEARCH( $\langle \mathcal{X}, \mathcal{A}, w \rangle, x, low, high$ )
2:  $mid := (low + high)/2$ ;
3: if ( $|high - low| < \delta$ ) then
4:   if ( $\text{WCA}_{gr}(\langle \mathcal{X}, \mathcal{A}, low \rangle, low, x)$ ) then
5:     return  $low$ ;
6:   else
7:     return  $high$ ;
8:   end if
9: end if
10: if ( $\text{WCA}_{gr}(\langle \mathcal{X}, \mathcal{A}, w \rangle, mid, x)$ ) then
11:   return SEARCH( $\langle \mathcal{X}, \mathcal{A}, w \rangle, x, low, mid$ );
12: else
13:   return SEARCH( $\langle \mathcal{X}, \mathcal{A}, w \rangle, x, mid, high$ );
14: end if
15: end
16:  $\delta := \min_{\langle x, y \rangle \in \mathcal{A}} w(\langle x, y \rangle)$ ;
17: return SEARCH( $\langle \mathcal{X}, \mathcal{A}, w \rangle, x, 0, \sum_{\langle y, x \rangle \in \mathcal{A}} w(\langle y, x \rangle)$ );

```

To complete the instance of MIN-BUDGET we set the argument of interest to be ψ . First notice that (via a similar argument to that of Proposition 7) $\varphi(Z_n)$ is unsatisfiable if and only if the least expensive budget is formed by setting $R = \{\langle \chi, \psi \rangle\}$ at a cost of 2^n . Hence $\beta\text{-opt}$ is found to be 2^n if and only if the bit string 1^{n+1} is reported as the value of $L_{\min}(\varphi)$. If $\varphi(Z_n)$ is satisfiable, let $\alpha = a_1 a_2 \dots a_n$ be the corresponding lexicographically minimal satisfying assignment, and $R_\alpha = \{\langle \neg z_i, z_i \rangle : a_i = 1\}$. Via an identical argument to that of Proposition 7 we have $\varphi \in \text{ge}(\langle \mathcal{Y}_\varphi, \mathcal{B}_\varphi \setminus R_\alpha \rangle)$ thus $\psi \in \text{ge}(\langle \mathcal{Y}_\varphi, \mathcal{B}_\varphi \setminus R_\alpha \rangle)$. It is easy to see that $\text{wt}(R_\alpha, w) = \sum_{i: a_i=1} 2^{i-1} \leq 2^n - 1$. Furthermore, if $R_\beta \subset \mathcal{B}_\varphi$ satisfied $\text{wt}(R_\beta, w) < \text{wt}(R_\alpha)$ then the lowest index attack $\langle \neg z_i, z_i \rangle$ in R_β must be less than the lowest indexed attack $\langle \neg z_i, z_i \rangle$ in R_α : this would contradict the definition of α as lexicographically minimal. On the other hand, if $\beta\text{-opt}(\langle \langle \mathcal{Y}_\varphi, \mathcal{B}_\varphi, w \rangle, \psi \rangle) = b \leq 2^n - 1$, then letting $b_1 \dots b_n$ be the n -bit binary representation of b , i.e., $b = \sum_{i: b_i=1} 2^{i-1}$, defines the lexicographically minimal satisfying assignment of φ so that $L_{\min}(\varphi) = 0b_1 \dots b_n$. □

Corollary 2. MIN-WITNESS is FP^{NP} -complete.

Proof. That MIN-WITNESS is FP^{NP} -hard is immediate from Proposition 13 and the remarks preceding its statement. The upper bound $\text{MIN-WITNESS} \in \text{FP}^{\text{NP}}$ follows by, for example the “prefix search” method described in [29]. □

It may be noted that the reduction in Proposition 13 uses weights whose value is exponential in $|\mathcal{X}|$, e.g., the weighting of $w(\langle \neg z_n, z_n \rangle) = 2^{n-1}$ and $|\mathcal{X}| = 2n + m + 2 =$

$O(n^3)$. This motivates the question of whether restricting the size of weights, e.g., insisting that $\max_{(x,y) \in \mathcal{A}} \log w(\langle x, y \rangle) = O(\log |\mathcal{X}|)$ so that the reduction of Proposition 13 no longer applies, could result in more efficient algorithms for MIN-BUDGET and MIN-WITNESS. Notice this question is well motivated in the sense that in practice we would not expect edge weights to be exponentially large with respect to the number of arguments, e.g., in moderate size systems of, say 1000 arguments, it seems unlikely that individual attacks with cost 2^{1000} would be meaningful or that (using $\mathbb{R}_>$) a level of precision needed to distinguish 1000 bit values was required.

We say a weighted argument system $\langle \mathcal{X}, \mathcal{A}, w \rangle$ is *reasonable* if

$$\max_{(x,y) \in \mathcal{A}} \log_2(w(\langle x, y \rangle)) \leq \lceil \log_2(|\mathcal{X}|) \rceil$$

For a reasonable weighted argument system it is immediate that

$$\forall x \in \mathcal{X} \quad \beta - \text{opt}(\langle \mathcal{X}, \mathcal{A}, w \rangle, x) \leq (|\mathcal{X}| - 1) \times 2|\mathcal{X}| \leq 2|\mathcal{X}|^2$$

and, thus, the size of any optimal budget can be represented in $O(\log |\mathcal{X}|)$ bits.

Proposition 14. *For $\langle \mathcal{X}, \mathcal{A}, w \rangle$ restricted to reasonable weighted argument systems,*

- a. MIN-BUDGET $\in \text{FP}^{\text{NP}}[O(\log |\mathcal{X}|)] \subseteq \text{FP}_{\parallel}^{\text{NP}}$.
- b. MIN-WITNESS is $\text{FP}_{\parallel}^{\text{NP}}$ -hard.

*Proof.*⁷ For part (a) it suffices to note that in reasonable systems, $\beta - \text{opt}(\langle \mathcal{X}, \mathcal{A}, w \rangle, x)$ requires at most $2 \log |\mathcal{X}| + 1$ to encode. Hence using a similar binary search approach to Algorithm 1 will compute $\beta - \text{opt}(\langle \mathcal{X}, \mathcal{A}, w \rangle, x)$ in $O(\log |\mathcal{X}|)$ calls to an oracle for WCA_{gr} .

For part (b), consider the following function problem defined on instances $\varphi(Z_n)$ of CNF-SAT

$$\mathcal{F}_{zero}(\varphi) = \begin{cases} 1^{n+1} & \text{if } \varphi \text{ is unsatisfiable} \\ 0\alpha & \text{where } \alpha \text{ is any maximum number of 0s satisfying assignment} \end{cases}$$

In [14], \mathcal{F}_{zero} is shown to be $\text{FP}_{\parallel}^{\text{NP}}$ -hard, so in order to prove (b) we show that we can (in polynomial time) transform instances of \mathcal{F}_{zero} , i.e., CNF formulae φ , to instances of MIN-WITNESS, i.e., reasonable weighted argument systems $\langle \mathcal{X}_{\varphi}, \mathcal{A}_{\varphi}, w \rangle$ and a specified argument of these, in such a way that solutions to MIN-WITNESS($\langle \mathcal{X}_{\varphi}, \mathcal{A}_{\varphi}, w \rangle, x$)

⁷Recall that $\text{FP}^{\text{NP}[\log]}$ is the class of functions computable by polynomial time algorithms that are allowed to make $O(\log |I|)$ queries to an NP oracle on instances, I , encoded in $|I|$ bits; $\text{FP}_{\parallel}^{\text{NP}}$ is the class of functions computable by polynomial time algorithms that make (polynomially many) *non-adaptive* queries to an NP oracle, so that all queries can be viewed as performed *simultaneously* (in parallel). For the analogous decision classes it is known that $\text{P}^{\text{NP}[\log]} = \text{P}_{\parallel}^{\text{NP}}$; however, only the relationship $\text{FP}^{\text{NP}[\log]} \subseteq \text{FP}_{\parallel}^{\text{NP}}$ has been proven for the function case ([28, Thm. 2.2]) and it is conjectured that the containment is strict, cf. the series of results from [28, pp. 184–190].

(of weight less than $n + 1$) have a 1-1 correspondence with *satisfying* solutions of $\mathcal{F}_{zero}(\varphi)$.

Given $\varphi(Z_n)$ in CNF, construct the argument system $\langle \mathcal{Y}_\varphi, \mathcal{B}_\varphi \rangle$ described in Proposition 13. We now use the weighting function:

$$w(\langle x, y \rangle) = \begin{cases} 1 & \text{if } \langle x, y \rangle = \langle \neg z_i, z_i \rangle \\ n + 1 & \text{if } \langle x, y \rangle = \langle \chi, \psi \rangle \\ n + 2 & \text{otherwise} \end{cases}$$

The system $\langle \mathcal{Y}_\varphi, \mathcal{B}_\varphi, w \rangle$ is clearly reasonable — $|\mathcal{Y}_\varphi| = 3n + m + 3$ and the maximum weight of $n + 2$ requires at most $\lceil \log_2(n + 2) \rceil + 1$ bits. The complete instance of MIN-WITNESS is $\langle \langle \mathcal{Y}_\varphi, \mathcal{B}_\varphi, w \rangle, \psi \rangle$. Now consider any $n + 1$ binary sequence $\alpha = ba_1a_2 \dots a_n$ defining a solution of $\mathcal{F}_{zero}(\varphi)$. First notice that if $\alpha = 1^{n+1}$ then φ is unsatisfiable and, as argued in Proposition 13, it follows that

$$\text{MIN-WITNESS}(\langle \mathcal{Y}_\varphi, \mathcal{B}_\varphi, w \rangle, \psi) = \{ \langle \chi, \psi \rangle \}$$

Similarly it follows that if $\{ \langle \chi, \psi \rangle \}$ describes the minimal set, R , witnessing $\psi \in ge(\langle \mathcal{Y}_\varphi, \mathcal{B}_\varphi \setminus R \rangle)$ then φ must be unsatisfiable and hence $\mathcal{F}_{zero}(\varphi) = 1^{n+1}$.

Now consider any solution of $\mathcal{F}_{zero}(\varphi)$ in the event that $\varphi(Z_n)$ is satisfiable, so that $b = 0$ and $a_1a_2 \dots a_n$ has a maximum number of 0 values. Let $R_\alpha = \{ \langle \neg z_i, z_i \rangle : a_i = 1 \}$. Then $\psi \in ge(\langle \mathcal{Y}_\varphi, \mathcal{B}_\varphi \setminus R_\alpha \rangle)$ and $wt(R_\alpha, w)$ is minimal amongst subsets of \mathcal{B}_φ with this property. Similarly given $R \subset \mathcal{B}_\varphi$ with $wt(R, w) \leq n$ and R a solution of $\text{MIN-WITNESS}(\langle \mathcal{Y}_\varphi, \mathcal{B}_\varphi, w \rangle, \psi)$, the $n + 1$ bit sequence $0a_1a_2 \dots a_n$ with $a_i = 1$ if and only if $\langle \neg z_i, z_i \rangle \in R$ defines a solution of $\mathcal{F}_{zero}(\varphi)$. \square

We noted above that one way of deriving a preference order over sets of arguments is to consider the minimal inconsistency budget required to make a set of arguments a solution. A related idea is to *count* the number of weighted extensions that an argument set appears in, for a given budget: we prefer argument sets that appear in more weighted grounded extensions. Formally, we denote the *rank* of an argument set S , given a weighted argument system $\langle \mathcal{X}, \mathcal{A}, w \rangle$ and inconsistency budget β , by $\rho(S, \langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$:

$$\rho(S, \langle \mathcal{X}, \mathcal{A}, w \rangle, \beta) = |\{ Y \in \mathcal{E}_{gr}^{\text{WT}}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta) : S \subseteq Y \}|.$$

Proposition 15. *Given a weighted argument system $\langle \mathcal{X}, \mathcal{A}, w \rangle$, argument set $S \subseteq \mathcal{X}$, and inconsistency budget β , computing $\rho(S, \langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$ is #P-complete.*

Proof. (Outline) For membership, consider a non-deterministic Turing machine that guesses some subset R of \mathcal{A} , and verifies that both $\sum_{\langle p, q \rangle \in R} w(\langle p, q \rangle) \leq \beta$ and $S \subseteq ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$. The number of accepting computations of this machine will be $\rho(S, \langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$. For hardness, we can reduce #SAT, using the construction of Proposition 7. It is straightforward to see that the reduction is parsimonious. \square

5.1. Computing Minimal Budgets via Mixed Integer Linear Programming

FP^{NP} is perhaps the paradigm complexity class for combinatorial optimisation problems such as the travelling salesman problem [39, p.416]. The fact that MIN-WITNESS is FP^{NP} -complete suggests that it might therefore be fruitful to apply techniques developed for combinatorial optimisation to MIN-WITNESS. In this sub-section, we will show how MIN-WITNESS can be formulated as a *mixed integer linear program* (MILP) [40].

The specific problem we address is as follows. We are given a weighted argument system $\langle \mathcal{X}, \mathcal{A}, w \rangle$ and some *target argument set* $Y \subseteq \mathcal{X}$. We then want to find the subgraph R of \mathcal{A} such that S is contained in the grounded extension of $\langle \mathcal{X}, \mathcal{A} \setminus R \rangle$ and $\text{wt}(R, w)$ is minimised. The MILP we produce for this is shown in Figure 7, and makes use of the following variables:

- for all $\alpha \in \mathcal{X}$, $\text{in}_\alpha \in \{0, 1\}$ and $\text{out}_\alpha \in \{0, 1\}$ are variables that indicate the final status of argument α : $\text{in}_\alpha = 1$ ($\text{out}_\alpha = 0$) will mean that argument α is in, while $\text{in}_\alpha = 0$ ($\text{out}_\alpha = 1$) will mean that α is out (of course there is some redundancy here, we do not need both sets of variables, but it makes the MILP easier to comprehend);
- for all $\langle p, q \rangle \in \mathcal{A}$, $\text{dis}_{\langle p, q \rangle} \in \{0, 1\}$ are variables that indicate whether an attack $\langle p, q \rangle$ has been “disabled” ($\text{dis}_{\langle p, q \rangle} = 1$) or not ($\text{dis}_{\langle p, q \rangle} = 0$), i.e., whether we are choosing disregarding the attack or not — if we disregard an attack, then we incur its cost;
- for all $\langle p, q \rangle \in \mathcal{A}$, $\text{dead}_{\langle p, q \rangle}$ will indicate whether the attack is active ($\text{dead}_{\langle p, q \rangle} = 0$) or not ($\text{dead}_{\langle p, q \rangle} = 1$): an attack $\langle p, q \rangle$ will be active if both $\text{dis}_{\langle p, q \rangle} = 0$ and $\text{in}_p = 1$.

The MILP makes use of one subsidiary definition: we let $\text{attacks}(\alpha)$ denote the set of arguments attacking α :

$$\text{attacks}(\alpha) = \{\alpha' : \langle \alpha', \alpha \rangle \in \mathcal{A}\}.$$

The MILP may be understood as follows:

- constraints (3), (4), and (6) say that variables $\text{dead}_{\langle p, q \rangle}$, $\text{dis}_{\langle p, q \rangle}$, and in_α take a value of either 0 or 1;
- constraint (5) says that all the variables in the target set Y must be “in”;
- constraint (7) states the relationship between in_α and out_α (and hence ensures that out_α variables take 0, 1 values);
- constraint (8) ensures that arguments with no attackers are “in”;

- constraint (9) ensures that variable $dead_{\langle\alpha, \alpha'\rangle}$ takes the value 1 if argument α is out, while constraint (10) ensures that variable $dead_{\langle\alpha, \alpha'\rangle}$ takes the value 1 if the attack $\langle\alpha, \alpha'\rangle$ is disabled;
- constraint (11) ensure that $dead_{\langle\alpha, \alpha'\rangle}$ takes the value 0 if argument α is in and the attack $\langle\alpha, \alpha'\rangle$ is enabled;
- constraint (12) ensures that if all attacks on an argument α are dead, then α is “in”;
- constraint (13) ensures that if any attack $tuple\alpha', \alpha$ is not dead, then α is “out”.

Note that the $w(\langle\alpha, \alpha'\rangle)$ values will be constants in the MILP we create, which is clearly of size polynomial in the inputs. In any solution to the MILP, the variables $dis_{\langle\alpha, \alpha'\rangle}$ such that $dis_{\langle\alpha, \alpha'\rangle} = 1$ will be the attacks to disregard: the fact that there is no smaller cost set is immediate from the objective function (2).

5.2. Positive results

Although credulous and sceptical acceptance are hard for acyclic frameworks, we can describe one class of frameworks for which these problems (and the budget, witnessing functions mentioned earlier are polynomial time computable).

We first note a straightforward simplification that can be made to the structure $\langle\mathcal{X}, \mathcal{A}\rangle$.

Lemma 1. *Let $\langle\mathcal{X}, \mathcal{A}, w\rangle$ be an arbitrary weighted argument system. We can construct in polynomial time a system $\langle\mathcal{X} \cup \mathcal{Y}, \mathcal{B}, w'\rangle$ for which*

- For any $R \subseteq \mathcal{A}$ for which $S = ge(\langle\mathcal{X}, \mathcal{A} \setminus R\rangle)$ there is a subset R' of \mathcal{B} (identifiable in polynomial time) such that $wt(R', w) \leq wt(R, w)$ and $S = ge(\langle\mathcal{X} \cup \mathcal{Y}, \mathcal{B} \setminus R'\rangle) \cap \mathcal{X}$.*
- For any $R' \subseteq \mathcal{B}$ with $S = ge(\langle\mathcal{X} \cup \mathcal{Y}, \mathcal{B} \setminus R'\rangle)$ there is a subset R of \mathcal{A} (identifiable in polynomial time) with $wt(R, w) \leq wt(R', w)$ and $S \cap \mathcal{X} = ge(\langle\mathcal{X}, \mathcal{A} \setminus R\rangle)$.*
- $\forall z \in \mathcal{X} \cup \mathcal{Y}, |\{u : \langle u, z \rangle \in \mathcal{B}\}| \leq 2$ and $|\{u : \langle z, u \rangle \in \mathcal{B}\}| \leq 2$*

In simplified terms Lemma 1 asserts that any weighted argument system may be “simulated” by one in which no argument attacks or is attacked by more than 2 arguments.

Proof. Let $\langle\mathcal{X}, \mathcal{A}, w\rangle$ a weighted argument system. Consider any $z \in \mathcal{X}$ for which $|\{y : \langle y, z \rangle \in \mathcal{A}\}| \geq 3$ and let

$$\{y : \langle y, z \rangle \in \mathcal{A}\} = \{y_1, y_2, \dots, y_k\}$$

Modify $\langle\mathcal{X}, \mathcal{A}, w\rangle$ as illustrated in the Figure 8 (it is *not* assumed that $\langle y_i, y_j \rangle \notin \mathcal{A}$).

The system of Figure 8, has arguments $\mathcal{X} \cup \{p, q\}$ and attacks

$$\mathcal{B} = \mathcal{A} \setminus \{\langle y_i, z \rangle : 2 \leq i \leq k\} \cup \{\langle p, z \rangle, \langle q, p \rangle\} \cup \{\langle y_i, q \rangle : 2 \leq i \leq k\}$$

minimise:

$$\sum_{\langle \alpha, \alpha' \rangle \in \mathcal{A}} w(\langle \alpha, \alpha' \rangle) \cdot dis_{\langle \alpha, \alpha' \rangle} \quad (2)$$

subject to constraints:

$$dead_{\langle \alpha, \alpha' \rangle} \in \{0, 1\} \quad \forall \langle \alpha, \alpha' \rangle \in \mathcal{A} \quad (3)$$

$$dis_{\langle \alpha, \alpha' \rangle} \in \{0, 1\} \quad \forall \langle \alpha, \alpha' \rangle \in \mathcal{A} \quad (4)$$

$$in_\alpha = 1 \quad \forall \alpha \in Y \quad (5)$$

$$in_\alpha \in \{0, 1\} \quad \forall \alpha \in \mathcal{X} \setminus Y \quad (6)$$

$$out_\alpha = 1 - in_\alpha \quad \forall \alpha \in \mathcal{X} \quad (7)$$

$$in_\alpha = 1 \quad \forall \alpha \in \mathcal{X} : attacks(\alpha) = \emptyset \quad (8)$$

$$dead_{\langle \alpha, \alpha' \rangle} \geq out_\alpha \quad \forall \langle \alpha, \alpha' \rangle \in \mathcal{A} \quad (9)$$

$$dead_{\langle \alpha, \alpha' \rangle} \geq dis_{\langle \alpha, \alpha' \rangle} \quad \forall \langle \alpha, \alpha' \rangle \in \mathcal{A} \quad (10)$$

$$out_\alpha + dis_{\langle \alpha, \alpha' \rangle} \geq dead_{\langle \alpha, \alpha' \rangle} \quad \forall \langle \alpha, \alpha' \rangle \in \mathcal{A} \quad (11)$$

$$out_\alpha \leq \sum_{\alpha' \in attacks(\alpha)} (1 - dead_{\langle \alpha', \alpha \rangle}) \quad \forall \alpha \in \mathcal{X} : attacks(\alpha) \neq \emptyset \quad (12)$$

$$out_\alpha \geq 1 - dead_{\langle \alpha', \alpha \rangle} \quad \forall \alpha : attacks(\alpha) \neq \emptyset : \forall \alpha' \in attacks(\alpha) \quad (13)$$

Figure 7: MILP for computing least weight grounded extensions.

The additional attacks have

$$w(\langle p, z \rangle) = \sum_{i=2}^k w(\langle y_i, z \rangle) ; w(\langle y_i, q \rangle) = w(\langle y_i, z \rangle) ; w(\langle q, p \rangle) = \sum_{\langle r, s \rangle \in \mathcal{A}} w(\langle r, s \rangle)$$

It is clear that the construction reduces the number of attacks on z to two and creates one additional argument (q) but which is attacked by only $k - 1$ (rather than k) arguments. The number of attacks on any other argument is unchanged. It thus suffices to show the following

- a. For any $R \subseteq \mathcal{A}$ with $S = ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$ there is some $R' \subseteq \mathcal{B}$ for which $S = ge(\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \setminus R' \rangle) \cap \mathcal{X}$ and $wt(R', w) \leq wt(R, w)$
- b. For any $R' \subseteq \mathcal{B}$ with $S' = ge(\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \setminus R' \rangle)$ there is some $R \subseteq \mathcal{A}$ for which $S' \cap \mathcal{X} = ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$ and $wt(R, w) \leq wt(R', w)$.




Figure 8: Reducing the number of attacks on arguments.

For (a) consider any set $R \subseteq \mathcal{A}$ and $S = ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$. Let $R_y = R \cap \{\langle y_i, z \rangle : 2 \leq i \leq k\}$. Now consider the subset, R' of \mathcal{B} defined via

$$R' = R \setminus R_y \cup \{\langle y_i, q \rangle : \langle y_i, z \rangle \in R_y\}$$

Certainly $wt(R', w) \leq wt(R, w)$, since

$$wt(R', w) = wt(R, w) - \sum_{y_i \in S_y} w(\langle y_i, z \rangle) + \sum_{y_i \in S_y} w(\langle y_i, q \rangle) = wt(R, w)$$

Furthermore

$$S = ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle) = ge(\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \setminus R' \rangle) \cap \mathcal{X} = S'$$

For suppose this were not the case. Then either there is some $x \in S$ for which $x \notin S'$ or some $x \in S'$ for which $x \notin S$. Consider the first of these — $\exists x \in S$ and $x \notin S'$. Since $x \in \mathcal{X}$ it follows that $x \notin ge(\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \setminus R' \rangle)$ and thus there is some argument $u \in \mathcal{X} \cup \{p, q\}$ for which $\langle u, x \rangle \in \mathcal{B} \setminus R'$ and $\langle u, x \rangle \notin \mathcal{A} \setminus R$. The only possible choice for $\{u, x\}$ is that of $x = z$ and $u = p$. It suffices to show that $z \in S$ implies $q \in S'$ (and hence, $p \notin S'$) from which $z \in S'$ as required. This, however, follows from $z \in S$ and the definition of R' as $R \setminus R_y \cup \{\langle y_i, q \rangle : \langle y_i, z \rangle \in R_y\}$ so that every attack on z in $\langle \mathcal{X}, \mathcal{A} \rangle$ is either defended (hence the corresponding attack on q is also defended in $\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \rangle$) or is an element of R so that the corresponding attack on q is in R' .

We are left only with the second possibility — $\exists x \in S'$ for which $x \notin S$. Noting that $S' \subseteq \mathcal{X}$ in order for $x \notin S$ to hold, there must be some $u \in \mathcal{X}$ with $\langle u, x \rangle \in \mathcal{A} \setminus R$ and $\langle u, x \rangle \notin \mathcal{B} \setminus R'$. Again the only possible choices are $x = z$ and $u \in \{y_2, \dots, y_k\}$. From $z \in S'$ and $\langle p, z \rangle \notin R'$ it follows that $q \in ge(\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \setminus R' \rangle)$; from $\langle u, z \rangle \in \mathcal{A} \setminus R$ it follows that $\langle u, q \rangle \in \mathcal{B} \setminus R'$. It must therefore be the case that the attack by u on q is defended in $ge(\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \setminus R' \rangle)$ and an identical defence to the attack by u on z (in $\langle \mathcal{X}, \mathcal{A} \setminus R \rangle$) is available, i.e., $z \in ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$ as required.

For part (b) again let $R' \subseteq \mathcal{B}$ and $S' = ge(\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \setminus R' \rangle)$. Consider the set $S' \cap \mathcal{X}$. We wish to identify $R \subseteq \mathcal{A}$ with $wt(R, w) \leq wt(R', w)$ and $S' \cap \mathcal{X} = ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$. Define $R \subseteq \mathcal{A}$ to be

$$\begin{cases} R' \cap \mathcal{A} \cup \{\langle y_i, z \rangle : \langle y_i, q \rangle \in R'\} & \text{if } \langle p, z \rangle \notin R' \\ R' \cap \mathcal{A} \cup \{\langle y_i, z \rangle : 2 \leq i \leq k\} & \text{if } \langle p, z \rangle \in R' \end{cases}$$

It is clear that $wt(R, w) \leq wt(R', w)$ and

$$ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle) \cap \mathcal{X} \setminus \{z\} = ge(\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \setminus R' \rangle) \cap \mathcal{X} \setminus \{z\}$$

so it remains only to show that

$$z \in ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle) \Leftrightarrow z \in ge(\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \setminus R' \rangle)$$

It is easy to see (from the definition of R) that this is the case whenever $\langle p, z \rangle \in R'$, so we may assume $\langle p, z \rangle \notin R'$. From $z \in ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$ it must be the case that $y_j \in ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$ implies $\langle y_j, z \rangle \in R$ so that $\langle y_j, q \rangle \in R'$. We deduce that $q \in ge(\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \setminus R' \rangle)$ since for each $\langle y_j, q \rangle \notin R'$ we have $y_j \notin ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle) \cup ge(\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \setminus R' \rangle)$. In consequence $z \in ge(\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \setminus R' \rangle)$ as required.

On the other hand suppose that $z \notin ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$. Then there is some y_j for which $y_j \in ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$ and $\langle y_j, z \rangle \notin R$: if $y_j = y_1$ then $\langle y_1, z \rangle \notin R'$ and $y_1 \in ge(\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \setminus R' \rangle)$ so that $z \notin ge(\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \setminus R' \rangle)$ as required. If $y_j \in \{y_2, \dots, y_k\}$ then $\langle y_j, z \rangle \in R$ implies $\langle y_j, q \rangle \in R'$ so that, since $y_j \in ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle) \cap ge(\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \setminus R' \rangle)$, we get $q \notin ge(\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \setminus R' \rangle)$ and hence $z \notin ge(\langle \mathcal{X} \cup \{p, q\}, \mathcal{B} \setminus R' \rangle)$ completing the proof of (b).

It is clear that the construction described can be repeated until all arguments satisfy the specified conditions. A similar construction deals with the case $|\{y : \langle z, y \rangle \in \mathcal{A}\}| \geq 3$. \square

We have shown, in the previous section, that deciding credulous acceptance with respect to a given budget and the computation of a minimal budget are computationally intractable not only in the case of general weighted argument systems but also when the supporting system $\langle \mathcal{X}, \mathcal{A} \rangle$ is acyclic or symmetric. These raise the question of the existence of “non-trivial” restrictions on the topology $\langle \mathcal{X}, \mathcal{A} \rangle$ that admit efficient decision procedures. We now consider one such restriction.

Recall that an argument system $\langle \mathcal{X}, \mathcal{A} \rangle$ is *bipartite* if \mathcal{X} may be partitioned into two sets \mathcal{Y} and \mathcal{Z} — i.e., $\mathcal{Y} \cup \mathcal{Z} = \mathcal{X}$ and $\mathcal{Y} \cap \mathcal{Z} = \emptyset$ — both of which are *conflict-free*.

Proposition 16. *If $\langle \mathcal{X}, \mathcal{A}, w \rangle$ is such that $\langle \mathcal{X}, \mathcal{A} \rangle$ is bipartite and acyclic⁸ then MIN-BUDGET and MIN-WITNESS are polynomial time computable.*

Proof. Let $\langle \mathcal{Y} \cup \mathcal{Z}, \mathcal{A}, w \rangle$ be a weighted bipartite acyclic argument system (where $\langle \mathcal{Y}, \mathcal{Z} \rangle$ is the partition of arguments into two conflict-free sets) and suppose $y \in \mathcal{Y}$ is an argument for which we wish to determine a minimum cost budget and associated witnessing subset, R of \mathcal{A} so that $x \in ge(\langle \mathcal{Y} \cup \mathcal{Z}, \mathcal{A} \setminus R \rangle)$. Noting the construction of Lemma 1 preserves both the properties of bipartition and acyclicity, without loss of generality we may assume no argument of $\mathcal{Y} \cup \mathcal{Z}$ attacks or is attacked by three or more arguments. Consider the subset of \mathcal{A} defined by

$$B_y = \{ \langle p, q \rangle \in \mathcal{A} : \text{There is a directed path from } q \text{ to } y \text{ in } \langle \mathcal{Y} \cup \mathcal{Z}, \mathcal{A} \rangle \}$$

If R_{\min} is a minimal cost subset of \mathcal{A} for which $y \in ge(\langle \mathcal{Y} \cup \mathcal{Z}, \mathcal{A} \setminus R_{\min} \rangle)$ then it must be the case that $R_{\min} \subseteq B_y$: $\langle \mathcal{Y} \cup \mathcal{Z}, \mathcal{A} \rangle$ is acyclic and so only attacks in B_y can influence whether or not $y \in ge(\langle \mathcal{Y} \cup \mathcal{Z}, \mathcal{A} \setminus R \rangle)$. Now consider the following labelling, $\theta : \mathcal{Y} \cup \mathcal{Z} \rightarrow \{\mathbf{in}, \mathbf{fail}\}$ of $\mathcal{Y} \cup \mathcal{Z}$

$$\theta(x) = \begin{cases} \mathbf{in} & \text{if } x = y \text{ or } \langle x, u \rangle \in \mathcal{A} \text{ and } \theta(u) = \mathbf{fail} \\ \mathbf{fail} & \text{if } \langle x, u \rangle \in \mathcal{A} \text{ and } \theta(u) = \mathbf{in} \end{cases}$$

This labelling is consistent since (with $y \in \mathcal{Y}$) only arguments in \mathcal{Z} could be labelled **fail** and only arguments in \mathcal{Y} could be labelled **in** (it is, of course, possible that not all arguments will be labelled, however, these will make no difference to the minimum budget). The idea behind the algorithm uses the fact that, in order for $y \in ge(\langle \mathcal{Y} \cup \mathcal{Z}, \mathcal{A} \setminus R \rangle)$, some defence to the (at most two) attacks, $\langle z_1, y \rangle$ and $\langle z_2, y \rangle$ is required. Such a defence can be formed either by including $\langle z_i, y \rangle \in R_{\min}$ or by ensuring that at least one attacker, $y_{i,j}$ of z_i also belongs to $ge(\langle \mathcal{Y} \cup \mathcal{Z}, \mathcal{A} \setminus R_{\min} \rangle)$. Since $\langle \mathcal{Y} \cup \mathcal{Z}, \mathcal{A} \rangle$ is acyclic we can continue reasoning thus until the case of arguments with no attackers is reached. In summary we can define the cost $c_{in}(x)$ of ensuring that x (with $\theta(x) = \mathbf{in}$ and attackers (a subset of) $\{u_1, u_2\}$) is included in $ge(\langle \mathcal{Y} \cup \mathcal{Z}, \mathcal{A} \setminus R \rangle)$, as follows. The enumeration below refers to the cases illustrated in Figure 9.

- a. If $\{y\}^- = \emptyset$ then

$$c_{in}(y) = 0$$

To avoid repetition, in each of the subsequent cases $\{y\}^- \neq \emptyset$.

- b. If $\{y\}^- \subseteq \{u_1, u_2\}$ with both $\{u_1\}^- = \{u_2\}^- = \emptyset$ then

$$c_{in}(y) = w(\langle u_1, y \rangle) + w(\langle u_2, y \rangle)$$

- c. If $\{y\}^- \subseteq \{u_1, u_2\}$, $\{u_1\}^- = \emptyset$, $\{u_2\}^- = \{y_1, y_2\}$ then

$$c_{in}(y) = w(\langle u_1, y \rangle) + \min\{c_{in}(y_1), c_{in}(y_2), w(\langle u_2, y \rangle)\}$$

⁸It is well known that testing if $\langle \mathcal{X}, \mathcal{A} \rangle$ satisfies both conditions can be carried out by polynomial time methods.




Figure 9: Cases in definition of $c_{in}(y)$.

d. If $\{y\}^- = \{u_1, u_2\}$, $\{u_1\}^- = \{y_1, y_2\}$, $\{u_2\}^- = \{y_2, y_3\}$

$$c_{in}(y) = \min \begin{cases} w(\langle u_1, y \rangle) + w(\langle u_2, y \rangle) \\ w(\langle u_1, y \rangle) + c_{in}(y_3) \\ c_{in}(y_1) + w(\langle u_2, y \rangle) \\ c_{in}(y_1) + c_{in}(y_3) \\ c_{in}(y_2) \end{cases}$$

e. If $\{y\}^- \subseteq \{u_1, u_2\}$, $\{u_1\}^- = \{y_1, y_2\}$, $\{u_2\}^- = \{y_2\}$

$$c_{in}(y) = \min \begin{cases} w(\langle u_1, y \rangle) + w(\langle u_2, y \rangle) \\ c_{in}(y_1) + w(\langle u_2, y \rangle) \\ c_{in}(y_2) \end{cases}$$

f. If $\{y\}^- \subseteq \{u_1, u_2\}$, $\{u_1\}^- = \{u_2\}^- = \{y_1, y_2\}$

$$c_{in}(y) = \min \begin{cases} w(\langle u_1, y \rangle) + w(\langle u_2, y \rangle) \\ c_{in}(y_1) \\ c_{in}(y_2) \end{cases}$$

g. If $\{y\}^- \subseteq \{u_1, u_2\}$, $\{u_1\}^- = \{y_1, y_2\}$, $\{u_2\}^- = \{y_3, y_4\}$ and $\{u_1\}^- \cap \{u_2\}^- = \emptyset$

$$c_{in}(y) = \min \begin{cases} w(\langle u_1, y \rangle) + w(\langle u_2, y \rangle) \\ w(\langle u_1, y \rangle) + c_{in}(y_3) \\ w(\langle u_1, y \rangle) + c_{in}(y_4) \\ c_{in}(y_1) + w(\langle u_2, y \rangle) \\ c_{in}(y_2) + w(\langle u_2, y \rangle) \\ c_{in}(y_1) + c_{in}(y_3) \\ c_{in}(y_1) + c_{in}(y_4) \\ c_{in}(y_2) + c_{in}(y_3) \\ c_{in}(y_2) + c_{in}(y_4) \end{cases}$$

Notice that the values $c_{in}(x)$ (for $\theta(x) = in$) are easily computed in polynomial time for each $x \in \mathcal{Y}$. To complete the proof, define $\delta(x)$ when $\theta(x) = in$ via

$$\delta(x) = \begin{cases} 0 & \text{if } \{x\}^- = \emptyset \\ 1 + \max\{\delta(y_i) : y_i \in \{\{x\}^-\}^-\} & \text{otherwise} \end{cases}$$

An easy inductive argument on $\delta(x)$ establishes that $c_{in}(x)$ is the minimal inconsistency budget for x . \square

6. Combinatorial Properties

We first note the following “extremal” result. Consider the function, $\mu(\langle \mathcal{X}, \mathcal{A}, w \rangle)$ defined by

$$\mu(\langle \mathcal{X}, \mathcal{A}, w \rangle) = \min\{\beta : \exists R \subseteq \mathcal{A} \text{ such that } wt(R, w) \leq \beta \text{ and } ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle) \neq \emptyset.\}$$

i.e., $\mu(\langle \mathcal{X}, \mathcal{A}, w \rangle)$ is the minimal budget needed to form a non-empty grounded extension within the system $\langle \mathcal{X}, \mathcal{A} \rangle$. Suppose we now define $\sigma : \mathbb{N} \rightarrow \mathbb{R}_{\geq}$ as

$$\tau(n) = \max_{\langle \mathcal{X}, \mathcal{A}, w \rangle : |\mathcal{X}|=n} \frac{\mu(\langle \mathcal{X}, \mathcal{A}, w \rangle)}{wt(\mathcal{A}, w)}$$

The following result gives an exact bound on $\tau(n)$.

Proposition 17. $\tau(n) = \frac{1}{n}$

Proof. To see that $\tau(n) \geq 1/n$, consider the system $\langle \mathcal{X}, \mathcal{A}, w \rangle$ of n arguments $\{x_1, \dots, x_n\}$ with attacks $\{\langle x_1, x_j \rangle : 2 \leq j \leq n\}$ each of weight 1; and $\{\langle x_j, x_1 \rangle : 2 \leq j \leq n\}$ each of weight $1/(n-1)$. This system has an empty grounded extension and in order to create a system with a non-empty grounded extension, we can either remove all attacks on x_1 (total cost $(n-1)/(n-1) = 1$) or the single attack on any x_j ($j \neq 1$) at cost 1. Hence for $\mu(\langle \mathcal{X}, \mathcal{A}, w \rangle) = 1$ and since $wt(\mathcal{A}, w) = (n-1)(1 + 1/(n-1)) = n$ the lower bound follows.

For the upper bound first observe that, for all $\langle \mathcal{X}, \mathcal{A}, w \rangle$

$$\mu(\langle \mathcal{X}, \mathcal{A}, w \rangle) \leq \min_{x \in \mathcal{X}} \sum_{y : \langle y, x \rangle \in \mathcal{A}} w(\langle y, x \rangle)$$

i.e., $\mu(\langle \mathcal{X}, \mathcal{A}, w \rangle)$ cannot be greater than the minimal inconsistency budget required to make a single argument *unattacked*, so that

$$\tau(n) \leq \max_{\langle \mathcal{X}, \mathcal{A}, w \rangle : |\mathcal{X}|=n} \frac{\min_{x \in \mathcal{X}} \sum_{y : \langle y, x \rangle \in \mathcal{A}} w(\langle y, x \rangle)}{wt(\mathcal{A}, w)}$$

Now for any $\langle \mathcal{X}, \mathcal{A}, w \rangle$

$$\sum_{x \in \mathcal{X}} \sum_{y : \langle y, x \rangle \in \mathcal{A}} w(\langle y, x \rangle) = wt(\mathcal{A}, w)$$

and, thus

$$\min_{x \in \mathcal{X}} \sum_{y : \langle y, x \rangle \in \mathcal{A}} w(\langle y, x \rangle) \leq \frac{wt(\mathcal{A}, w)}{|\mathcal{X}|}$$

with equality *if and only if* for all $x, x' \in \mathcal{X}$,

$$\sum_{y : \langle y, x \rangle \in \mathcal{A}} w(\langle y, x \rangle) = \sum_{y : \langle y, x' \rangle \in \mathcal{A}} w(\langle y, x' \rangle)$$

Hence,

$$\mu(\langle \mathcal{X}, \mathcal{A}, w \rangle) \leq \frac{wt(\mathcal{A}, w)}{|\mathcal{X}|} \times \frac{1}{wt(\mathcal{A}, w)} = \frac{1}{|\mathcal{X}|}$$

so that $\tau(n) \leq 1/n$. □

Recall that the precise formulation of WCA_{gr} asks of an instance $\langle \langle \mathcal{X}, \mathcal{A}, w \rangle, \beta, x \rangle$ if there is some $R \subseteq \mathcal{A}$ of cost at most β for which $x \in ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$. If, instead of a *single* argument x we consider a *set* of arguments S , certainly the associated decision problem is as hard (but no harder). Related to this question we have the following,

POSITION:

Instance: $\langle \mathcal{X}, \mathcal{A}, w \rangle; S \subseteq \mathcal{X}$.

Question: Is there a subset R of \mathcal{A} for which $S = ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$?

Notice the difference between POSITION and WCA_{gr} is that latter asks for a subset, R of \mathcal{A} , for which S is a *subset* of $ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$, whereas POSITION asks about the existence of a subset for which S is *identical to* $ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$. Now, although POSITION does not use budgets in its specification, one has an obvious function form and associated optimization problem:

$$OPT-POSN(\langle \mathcal{X}, \mathcal{A}, w \rangle, S) = \begin{cases} \perp & \text{if } \neg \text{POSITION}(\langle \mathcal{X}, \mathcal{A}, w \rangle, S) \\ \min\{ wt(R, w) : S = ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle) \} & \text{otherwise} \end{cases}$$

The *strongly-connected component* (SCC) decomposition of $\langle \mathcal{X}, \mathcal{A} \rangle$ partitions \mathcal{X} according to the equivalence classes induced by the relation $\rho(x, y)$ defined over $\mathcal{X} \times \mathcal{X}$ 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 $\langle \mathcal{X}, \mathcal{A} \rangle$. We will denote the set of strongly connected components of $\langle \mathcal{X}, \mathcal{A} \rangle$ as $\{\mathcal{C}_1, \dots, \mathcal{C}_k\}$, with $\mathcal{C}_i = \langle \mathcal{X}_i, \mathcal{A}_i \rangle$. 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 $\{\mathcal{C}_1, \dots, \mathcal{C}_k\}$ is defined as $(\mathcal{C}_i \prec \mathcal{C}_j) \Leftrightarrow (i \neq j)$ and $\exists p \in \mathcal{C}_i, q \in \mathcal{C}_j$ such that there is a directed path from p to q .

Proposition 18. POSITION is polynomial time decidable.

Proof. Let $\langle \langle \mathcal{X}, \mathcal{A}, w \rangle, S \rangle$ be an instance of POSITION. Consider the argument system formed by removing all attacks, $R_S = A_S \cup B_S$, where

$$\begin{aligned} A_S &= \{ \langle x, y \rangle : y \in S \} \\ B_S &= \{ \langle y, z \rangle \in \mathcal{A} \setminus A_S : y \in S \text{ and } |\mathcal{C}(z)| \geq 2 \} \end{aligned}$$

where $\mathcal{C}(x)$ is the component containing x in the strongly-connected component decomposition of $\langle \mathcal{X}, \mathcal{A} \setminus A_S \rangle$. If $S = ge(\langle \mathcal{X}, \mathcal{A} \setminus R_S \rangle)$ it is immediate that $\langle \langle \mathcal{X}, \mathcal{A} \rangle, S \rangle$ is accepted as an instance of POSITION. Thus suppose $S \neq ge(\langle \mathcal{X}, \mathcal{A} \setminus R_S \rangle)$. Noting that every $y \in S$ has $\{y\}^- = \emptyset$ in $\langle \mathcal{X}, \mathcal{A} \setminus R_S \rangle$ (since A_S contains every attack on S), it is certainly the case that $S \subseteq ge(\langle \mathcal{X}, \mathcal{A} \setminus R_S \rangle)$. Thus if these sets differ then $ge(\langle \mathcal{X}, \mathcal{A} \setminus R_S \rangle) \setminus S \neq \emptyset$. It is easy to see that in such cases we must have an argument $x \in \mathcal{X} \setminus S$ for which $\{x\}^- \subseteq S^+$, i.e., in the system $\langle \mathcal{X}, \mathcal{A} \setminus R_S \rangle$ every attacker of x is counterattacked by some argument of S . It therefore follows that any subset T of \mathcal{A} for which $S = ge(\langle \mathcal{X}, \mathcal{A} \setminus T \rangle)$ must be such that there is no argument $x \in \mathcal{X} \setminus S$ for which $\{x\}^- \subseteq S^+$ (with respect to $\langle \mathcal{X}, \mathcal{A} \setminus T \rangle$). Now suppose, to the contrary that $\langle \langle \mathcal{X}, \mathcal{A}, w \rangle, S \rangle$ is a positive instance of POSITION but $S \neq ge(\langle \mathcal{X}, \mathcal{A} \setminus R_S \rangle)$. Let $T \subseteq \mathcal{A}$ be such that $S = ge(\langle \mathcal{X}, \mathcal{A} \setminus T \rangle)$. Notice that, without loss of generality, we may assume $A_S \subseteq T$: since $S = ge(\langle \mathcal{X}, \mathcal{A} \setminus T \rangle)$, hence S counterattacks any attacker of S , i.e. $y \in S^-$ implies $y \in S^+$ (w.r.t. $\langle \mathcal{X}, \mathcal{A} \setminus T \rangle$). Thus $T \cup A_S$ also provides a witnessing subset of \mathcal{A} for $\langle \langle \mathcal{X}, \mathcal{A}, w \rangle, S \rangle$ being accepted. Now consider the strongly connected component decomposition of $\langle \mathcal{X}, \mathcal{A} \setminus A_S \rangle$ and, in particular, any argument $x \in \mathcal{X} \setminus S$ for which $\{x\}^- \subseteq S^+$ (in $\langle \mathcal{X}, \mathcal{A} \setminus A_S \rangle$). From the premise that $S \neq ge(\langle \mathcal{X}, \mathcal{A} \setminus R_S \rangle)$ we know that such an argument exists. First observe that $|\mathcal{C}(x)| = 1$: x cannot be in the same component as any argument of S in $\langle \mathcal{X}, \mathcal{A} \setminus A_S \rangle$ since every argument in S is unattacked in this system. Where it the case that $|\mathcal{C}(x)| \geq 2$ this would contradict $\{x\}^- \subseteq S^+$ (since x must be involved in a mutual attack with some argument not in S). In total, from $|\mathcal{C}(x)| = 1$ and $\{x\}^- \subseteq S^+$ it not possible to extend A_S to some subset T in such a way that x is kept out of $ge(\langle \mathcal{X}, \mathcal{A} \setminus T \rangle)$: adding any subset of the attacks on x to T will not affect $\{x\}^- \subseteq S^+$. Similarly ensuring that some $z \in S^+$ (in $\langle \mathcal{X}, \mathcal{A} \setminus A_S \rangle$) no longer belongs to S^+ (in $\langle \mathcal{X}, \mathcal{A} \setminus T \rangle$) will then force $z \in ge(\langle \mathcal{X}, \mathcal{A} \setminus T \rangle)$ Figure 10. \square




Figure 10: Relationship between S , S^+ , A_S , $z \in S^+$ and y with $\{y\}^- \subseteq S^+$.

7. Related Work

In this section, we formally investigate the relationship of our framework to other related argumentation frameworks in the literature. Before we present our technical results, we note that there are several other frameworks that should be mentioned in passing.

First, there are other interesting developments of abstract argumentation such as a framework for defeasible reasoning about preferences that provides a context dependent mechanism for determining which argument is preferred to which [34, 37]. This also offers a valuable solution to dealing with multiple extensions, but conceptually and formally the proposal is complementary to ours. Also of interest are the proposals for introducing information about how the audience views each argument [11]. In addition, our work is also related to work on possibilistic truth-maintenance systems [18] where assumptions are weighted, conclusions based on the assumptions inherit the weights, and consistent “environments” are established. What is particularly reminiscent about the work in [18] is that, again in our terms, it makes use of inconsistency budget — this is exactly the weight with which the inconsistency \perp can be inferred. Anything that can be inferred with a greater weight than \perp is then taken to hold, anything with a lesser weight is taken to be unreliable, which is broadly the effect of the inconsistency budget in our work.

We note that there has been a good deal of work on incorporating numerical and non-numerical strengths (though not strengths of attack) into argumentation systems that are not based on Dung’s work. In [30], to take the earliest example, the use of probability measures and beliefs in the sense of Shafer’s theory of evidence [47], is described. An argumentation system that uses weights which represent qualitative abstractions of probability values is presented in [41], while in [49], the weights are infinitesimal probabilities in the sense of [48]. There is also much work on combinations of logic and probability such as [4], [38] and [43], which, while not explicitly taking the form of argumentation, has much in common with it.

Now, let us turn to the issue of the formal relationship between weighted argument

systems and other argumentation frameworks. The most closely related research with respect to the themes of the current article concern the variety of proposals building directly on Dung’s original argument system model, in particular those approaches formalising settings under which attacks in \mathcal{A} are disregarded and their associated semantics. In this regard the main contributions of interest are:

- the *preference-based frameworks* (PAFs) of Amgoud and Cayrol [1];
- the *value-based argumentation frameworks* (VAFs) of Bench-Capon [10];
- *resolution-based semantics* of Baroni and Giacomin [7, 6]; and
- the *extended argumentation frameworks* (EAFs) recently proposed and analysed by Modgil [36, 35].

With regard to these formalisms, the two key questions of interest with respect to weighted argument systems are the following:

- Q1: Is it possible (by suitable choice of weighting function and inconsistency budget) always to define a weighted argument system that, in some meaningful sense, encapsulates the characteristics of instances of these systems?
- Q2: Secondly, are there examples whereby, e.g., the set $\mathcal{E}_{gr}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$, either cannot be described within such systems or can only be so described using significantly larger frameworks?

Before addressing these questions, it is useful to review some common characteristics shared by all of these proposals. One immediate such feature is that all propose a semantics predicated on *removing attacks* from a standard Dung-style argument system $\langle \mathcal{X}, \mathcal{A} \rangle$. What differentiates the proposals is the manner in which an attack $\langle x, y \rangle \in \mathcal{A}$ may be disregarded is specified. A second feature common to all of these formalisms is in the nature of those subsets of \mathcal{X} which are viewed as “collectively justified” following elimination of attacks in $\langle \mathcal{X}, \mathcal{A} \rangle$. Although it is sometimes possible to view these semantics as “parametric”⁹ (as is the case for EAFs, resolution-based semantics and weighted argument systems), it turns out that the most useful solution concept derives from the *grounded extension*. Hence, in keeping with the related approaches, our principal focus in the preceding pages has been on properties of $\mathcal{E}_{gr}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$; in preference and value-based frameworks (having discarded attacks, $R \subseteq \mathcal{A}$ which do not succeed) the subset of \mathcal{X} in $\langle \mathcal{X}, \mathcal{A} \setminus R \rangle$ of interest is exactly $\mathcal{E}_{gr}(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$, [1, Defn. 3.6, p. 205], [10, Thm. 6.4, p. 438].¹⁰ In both [7, 6] and [36, 35] it has been

⁹i.e., Given an arbitrary Dung argumentation semantics $\sigma : 2^{\mathcal{X}} \rightarrow \{\top, \perp\}$ a related semantics based on σ is well-defined for the modified structures.

¹⁰Although [10] refers to “preferred extensions” in $\langle \mathcal{X}, \mathcal{A} \setminus R \rangle$, since this is acyclic it satisfies $\mathcal{E}_{gr}(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle) = \mathcal{E}_{pr}(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$. In both PAF and VAF approaches a key property of the preference (resp. value ordering) relations considered is that these have the effect of eliminating some set of attacks, \mathcal{R} , with the result that $\langle \mathcal{X}, \mathcal{A} \setminus \mathcal{R} \rangle$ becomes *acyclic*.




Figure 11: $\mathcal{E}_{gr}^{WT}(W_3, 2) = \{\emptyset, \{y\}, \{x, z\}\}$

the reformulation of grounded extensions within the augmented systems that has given rise to interesting issues: although other extension-based semantics are well-defined within these approaches (as, indeed, is the case with weighted systems), the fact that known complexity lower bounds, e.g., for credulous and sceptical preferred semantics, trivially, continue to hold renders detailed consideration of such to be uninteresting. In contrast, given that computational problems concerning grounded extensions are tractable in (Dung style) AFs, the extent to which computationally undesirable properties arise within analogues from augmented systems is of some importance.

In terms of how such mechanisms related to weighted argument systems the questions highlighted above can now be more formally phrased as follows: given $AUG \in \{PAF, VAF, RES, EAF\}$ with gr -AUG denoting the variant of grounded extensions proposed, what is the relationship (if any) between \mathcal{E}_{gr}^{AUG} and \mathcal{E}_{gr}^{WT} ?

7.1. Preference-based and Value-based frameworks

A *preference-based argumentation framework* (PAF) is a triple $\langle \mathcal{X}, \mathcal{A}, \succ \rangle$ where $\langle \mathcal{X}, \mathcal{A} \rangle$ is a Dung-style argumentation framework and $\succ \subseteq \mathcal{X} \times \mathcal{X}$ is a *preorder relation* over \mathcal{X} . Fixing $R = \{\langle x, y \rangle \in \mathcal{A} : y \succ x\}$ the acceptable arguments of $\langle \mathcal{X}, \mathcal{A}, \succ \rangle$ are exactly those belonging to $ge(\langle \mathcal{X}, \mathcal{A} \setminus R \rangle)$. Notice that once \succ is fixed, gr -PAF is a unique status semantics, i.e., $|\mathcal{E}_{gr}^{PAF}(\langle \mathcal{X}, \mathcal{A}, \succ \rangle)| = 1$. As we have already seen, in weighted argument systems, $\mathcal{E}_{gr}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$ may contain a number of distinct sets, thus one immediately finds examples wherein

$$\exists w, \beta \forall \succ \mathcal{E}_{gr}^{PAF}(\langle \mathcal{X}, \mathcal{A}, \succ \rangle) \neq \mathcal{E}_{gr}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$$

As a simple example of this, we have the system W_3 of Fig. 11: given an inconsistency budget $\beta = 2$ yields three β -grounded extensions: the empty set (removing no attacks); $\{y\}$ by removing the unit cost attacks $\{\langle x, y \rangle, \langle z, y \rangle\}$ and $\{x, z\}$ by removing either of the attacks $\langle y, z \rangle$ or $\langle y, x \rangle$ of weight 2.

In contrast we have,

Proposition 19. *Let $\langle \mathcal{X}, \mathcal{A}, \succ \rangle$ be any PAF. There is a weight function $w : \mathcal{A} \rightarrow \mathbb{R}_{>}$ and inconsistency budget β for which $\mathcal{E}_{gr}^{PAF}(\mathcal{X}, \mathcal{A}, \succ) \subseteq \mathcal{E}_{gr}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$.*

Proof. Given $\langle \mathcal{X}, \mathcal{A}, \succ \rangle$, fix R to be

$$\min\{ S \subseteq \mathcal{A} : \forall \langle x, y \rangle \in S \ y \succ x \ \& \ ge(\langle \mathcal{X}, \mathcal{A} \setminus S \rangle) = ge(\langle \mathcal{X}, \mathcal{A} \setminus \{ \langle p, q \rangle : q \succ p \} \rangle)\}$$

where \min is w.r.t. \subseteq . Notice that R is well-defined since the set $\{ \langle p, q \rangle : q \succ p \}$ has the required property. Now use the weight function under which $w(\langle x, y \rangle) = 1$ for any $\langle x, y \rangle \in R$ and $w(\langle x, y \rangle) = |\mathcal{X}|^2 + 1$ for all other attacks. Finally set $\beta = |R|$. It is immediate, from the choice of S , that $\mathcal{E}_{gr}^{\text{PAF}}(\mathcal{X}, \mathcal{A}, \succ) \subseteq \mathcal{E}_{gr}^{\text{WT}}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$. \square

A similar construction may be used for VAFs w.r.t. a given audience $R \subset \mathcal{V} \times \mathcal{V}$: recall that an *audience*, R , for a VAF, $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle$, is a (possibly partial) preorder \succ_R of \mathcal{V} . Given an audience, R , and VAF, $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle$ attacks $\langle x, y \rangle \in \mathcal{A}$ are disregarded should it be the case that $\eta(y) \succ_R \eta(x)$. When R yields a total ordering of \mathcal{V} (so-called *specific* audiences), the reduced framework formed by eliminating $\langle x, y \rangle \in \mathcal{A}$ for which $\eta(y) \succ_R \eta(x)$ is acyclic. In this case weighted argument systems can mirror the semantics of specific audiences in VAFs by applying a weighting function $w(\langle x, y \rangle) = 1$ for every $\langle x, y \rangle \in \mathcal{A}$ for which $\eta(y) \succ_R \eta(x)$ and $w(\langle x, y \rangle) = |\mathcal{X}|^2 + 1$ otherwise.

In general it cannot be guaranteed that w and β may be chosen to give $\mathcal{E}_{gr}^{\text{PAF}}(\mathcal{X}, \mathcal{A}, \succ) = \mathcal{E}_{gr}^{\text{WT}}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$ (and similarly for audiences in VAFs): the reason being that $\mathcal{E}_{gr}^{\text{WT}}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$ may contain subsets of \mathcal{X} resulting by using an inconsistency budget, γ , which is strictly smaller than that allowed by β in the constructions described. With the construction described, however, since R is chosen to be *minimal* we obtain,

$$\mathcal{E}_{gr}^{\text{PAF}}(\mathcal{X}, \mathcal{A}, \succ) = \mathcal{E}_{gr}^{\text{WT}}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta) \setminus \mathcal{E}_{gr}^{\text{WT}}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta - 1)$$

We conclude with a brief comment regarding the extent to which sets in $\mathcal{E}_{gr}^{\text{WT}}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$ can be captured (using a suitable preference order) within $\mathcal{E}_{gr}^{\text{PAF}}(\mathcal{X}, \mathcal{A}, \succ)$. As a basic observation we note that since \succ must be transitive it is easy to construct examples in which $S \in \mathcal{E}_{gr}^{\text{WT}}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$ for some choice of β but $S \notin \mathcal{E}_{gr}^{\text{PAF}}(\mathcal{X}, \mathcal{A}, \succ)$, e.g., choose $\mathcal{X} = p, q, r$, $\mathcal{A} = \{ \langle p, q \rangle, \langle q, r \rangle, \langle r, p \rangle \}$ with all attacks having weight 1; $\{ p, q, r \} \in \mathcal{E}_{gr}^{\text{WT}}(\langle \mathcal{X}, \mathcal{A}, w \rangle, 3)$ but cannot form the the grounded extension arising from any preference ordering of $\{ p, q, r \}$ (since any such ordering would have to satisfy $p \succ r \succ q \succ p$). While, superficially the presence of (potentially) $|\mathcal{X}|!$ possible preference orderings would suggest there is sufficient scope to capture all $2^{|\mathcal{X}|}$ subsets of $|\mathcal{X}|$ (or at least those expressible as weighted solutions), we note the following property of “almost all” PAFs¹¹.

Proposition 20. *For almost all preference orderings, \succ , in almost all AFs $\langle \mathcal{X}, \mathcal{A} \rangle$, $ge(\langle \mathcal{X}, \mathcal{A}, \succ \rangle) = \emptyset$, i.e.*

$$\lim_{|\mathcal{X}| \rightarrow \infty} \frac{|\{ \langle \mathcal{X}, \mathcal{A}, \succ \rangle : ge(\langle \mathcal{X}, \mathcal{A}, \succ \rangle) \neq \emptyset \}|}{|\{ \langle \mathcal{X}, \mathcal{A}, \succ \rangle : \langle \mathcal{X}, \mathcal{A}, \succ \rangle \text{ is a PAF} \}|} = 0$$

¹¹We say a property, P , holds for *almost all* PAFs if the proportion of n argument PAFs with the property approaches 1 as $n \rightarrow \infty$.




Figure 12: $\{\{x, y, u\}, \{x, u, z\}\} \subset \mathcal{E}_{gr}^{WT}(W_4, 3) \setminus (\mathcal{E}_{gr}^{RES}(W_4) \cup \mathcal{E}_{gr}^{WT}(W_4, 2))$

Proof. (Outline) An easy counting argument based on the facts that a PAF is a directed graph structure and the fraction of all directed graphs with at least one initial arguments becomes vanishingly small as $|\mathcal{X}|$ increases. \square

In other words, in a precise technical sense, *almost all* PAFs have an empty grounded extension.

7.2. Resolution-based grounded semantics — GR*

Given $\langle \mathcal{X}, \mathcal{A} \rangle$, Baroni and Giacomin [7] introduce resolution-based semantics in terms of eliminating *mutual attacks* between pairs of arguments. The set of mutually attacking arguments $M(\langle \mathcal{X}, \mathcal{A} \rangle)$ consists of those pairs $\{x, y\} \subseteq \mathcal{X}$

$$\{\{x, y\} : x \neq y \ \& \ \{\langle x, y \rangle, \langle y, x \rangle\} \subseteq \mathcal{A}\}$$

A (full) *resolution* of $\langle \mathcal{X}, \mathcal{A} \rangle$ is a subset R of \mathcal{A} for which R contains *exactly one* of the attacks $\langle x, y \rangle$ or $\langle y, x \rangle$ for each $\{x, y\} \in M(\langle \mathcal{X}, \mathcal{A} \rangle)$. The set of *resolution-based grounded extensions* (\mathcal{E}_{gr}^{RES}) of $\langle \mathcal{X}, \mathcal{A} \rangle$ is given as

$$\min \{ge(\mathcal{X}, \mathcal{A} \setminus R) : R \text{ is a full resolution of } M(\langle \mathcal{X}, \mathcal{A} \rangle).\}$$

where, again, min is w.r.t. \subseteq .

Given $\langle \mathcal{X}, \mathcal{A} \rangle$, by assigning $w(\langle x, y \rangle) = w(\langle y, x \rangle) = 1$, if $\{x, y\} \in M(\langle \mathcal{X}, \mathcal{A} \rangle)$, and $w(\langle x, y \rangle) = |\mathcal{X}|^2 + 1$ whenever $\{x, y\} \notin M(\langle \mathcal{X}, \mathcal{A} \rangle)$ with $\beta = |M(\langle \mathcal{X}, \mathcal{A} \rangle)|$ it is easily seen that $\mathcal{E}_{gr}^{RES}(\langle \mathcal{X}, \mathcal{A} \rangle) \subseteq \mathcal{E}_{gr}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$. Typically, however, $\mathcal{E}_{gr}^{RES}(\langle \mathcal{X}, \mathcal{A} \rangle)$ will be a *strict* subset of $\mathcal{E}_{gr}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$ even when a budget of exactly β is used. For example, consider Fig. 12 for which \mathcal{E}_{gr}^{RES} is $\{\{u, y\}, \{x, z\}, \{x, u\}\}$.

In this example, $\mathcal{E}_{gr}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, 3) \setminus \mathcal{E}_{gr}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, 2) \neq \mathcal{E}_{gr}^{RES}(\langle \mathcal{X}, \mathcal{A} \rangle)$ since

$$\{\{x, y, u\}, \{x, u, z\}\} \subset \mathcal{E}_{gr}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, 3) \setminus \mathcal{E}_{gr}^{WT}(\langle \mathcal{X}, \mathcal{A}, w \rangle, 2)$$

7.3. Extended Argumentation Frameworks

Recently, Modgil [36, 35] has described a development of the basic directed graph formalism in Dung's systems to encompass the notion of attacks being themselves attacked. Formally,

Definition 7. An extended argumentation framework (EAF) is defined by a triple $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$ wherein $\langle \mathcal{X}, \mathcal{A} \rangle$ is a standard Dung argument framework and $\mathcal{D} \subseteq \mathcal{X} \times \mathcal{A}$ describes an attack relation whereby arguments are the source of attacks on attacks.¹² The relation \mathcal{D} must be such that: if $\{\langle v, \langle x, y \rangle \rangle, \langle u, \langle z, y \rangle \rangle\} \subseteq \mathcal{D}$ then $\{\langle v, u \rangle, \langle u, v \rangle\} \subseteq \mathcal{A}$.

A set $S \subseteq \mathcal{X}$, is conflict-free (within $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$) if $\forall x, y \in S$ should $\langle x, y \rangle \in \mathcal{A}$ then there is some $z \in S$ for which $\langle z, \langle x, y \rangle \rangle \in \mathcal{D}$. An attack $\langle x, y \rangle \in \mathcal{A}$ succeeds w.r.t $S \subseteq \mathcal{X}$ (denoted $x \rightarrow^S y$) if for every $z \in S$ $\langle z, \langle x, y \rangle \rangle \notin \mathcal{D}$.

The concept “ $x \in \mathcal{X}$ is acceptable w.r.t. $S \subseteq \mathcal{X}$ ” for EAFs (which underpins notions of the grounded extension) involves a number of subtleties: recall from Defn 2 that in standard AFs, x is acceptable to S if every attack — $\langle y, x \rangle$ — on x is countered by some attack $\langle z, y \rangle$ with $z \in S$. In EAFs, given S and x the standard definition of “ x is acceptable w.r.t. S ” fails to consider the possibility that the defence $\langle z, y \rangle$ to $\langle y, x \rangle$ may itself be attacked by some *argument* u , i.e., $\langle u, \langle y, z \rangle \rangle$ is an element of \mathcal{D} and, thus, not only must y be countered but also, if the defence provided by z is to be used, the argument u *must also be attacked* by some member of S . In order formally to capture the property that “ x is acceptable w.r.t. S ”, [36] introduces the concept of *reinstatement sets* for a (successful w.r.t. S) attack $v \rightarrow^S w$.

Definition 8. Given $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$ let $S \subseteq \mathcal{X}$ and $v \rightarrow^S w$ succeed w.r.t. S . A subset $R_S = \{y_1 \rightarrow^S z_1, y_2 \rightarrow^S z_2, \dots, y_r \rightarrow^S z_r\}$ of those attacks which succeed w.r.t. S is called a *reinstatement set* for the attack $v \rightarrow^S w$ if R_S satisfies all of the following conditions:

R1. $v \rightarrow^S w \in R_S$.

R2. $\forall 1 \leq i \leq r$ $y_i \in S$.

R3. For every $y \rightarrow^S z \in R_S$ and every $\langle z', \langle y, z \rangle \rangle \in \mathcal{D}$ there is some $y' \rightarrow^S z' \in R_S$.

An argument $x \in \mathcal{X}$ is acceptable w.r.t $S \subseteq \mathcal{X}$ in the EAF $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$ if for every successful attack $z \rightarrow^S x$ there is a successful attack $y \rightarrow^S z$ and a reinstatement set R_S for $y \rightarrow^S z$.

The grounded extension of the EAF $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$ is formed as

$$ge_{\text{EAF}}(\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle) = \bigcup_{k=0}^{\infty} \mathcal{F}_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle}^k$$

where $\mathcal{F}_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle}^0 = \emptyset$, $\mathcal{F}_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle}^{i+1} = \mathcal{F}_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle}(\mathcal{F}_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle}^{i+1})$ and for a conflict-free subset S of \mathcal{X} , $\mathcal{F}_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle}(S)$ is the set of arguments which are acceptable w.r.t. S .

A subset, S is admissible (in the EAF $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$) if every $x \in S$ is acceptable w.r.t. S ; S is preferred if it is a maximal admissible set.

As observed in [36], the grounded extension is well-defined, but also has a number of distinctive properties. Firstly, in contrast to Dung style AFs, it is *not* always the

¹²An alternative model of “attacks being attacked” is also described in Baroni *et al.* [5].

case that $S \in \mathcal{E}_{pr}^{\text{EAF}}(\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle)$ implies $ge_{\text{EAF}}(\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle) \subseteq S$, i.e., there are EAFs whose grounded extension contains arguments that do not belong to *every* preferred extension. Secondly, (and, again, contrasting with standard systems) there are EAFs within which one may identify subsets S and T and arguments x such that $S \subset T$, x is acceptable w.r.t S but x is *not* acceptable w.r.t. T .

Recent results, reported in Dunne *et al.* [25], have established that, in common with other developments of Dung's frameworks such as those we have considered above, the problem of determining if $x \in \mathcal{X}$ is acceptable w.r.t. $S \subseteq \mathcal{X}$ in the EAF $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$ is polynomial time decidable (and thus computation of the grounded extension as well as verifying S as an admissible set are also polynomial time problems). In contrast, however, [25] also establish that deciding if a given EAF exhibits the behaviours mentioned above (i.e., non-sceptically accepted members of the grounded extension, non-monotonic behaviour of argument acceptability) are likely to be computationally intractable.

Turning to the question of the extent to which EAFs can be modelled within weighted systems and *vice-versa* we note that, again, $|\mathcal{E}_{gr}^{\text{EAF}}(\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle)| = 1$. Given $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$ we obtain a weighted system $\langle \mathcal{X}, \mathcal{A}, w \rangle$ and inconsistency budget β with $\mathcal{E}_{gr}^{\text{EAF}}(\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle) \subseteq \mathcal{E}_{gr}^{\text{WT}}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$ as follows.

Observe that in progressing from $\mathcal{F}_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle}^0$ to $ge_{\text{EAF}}(\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle)$ as described by the process of Defn 8, typically a number of attacks in \mathcal{A} will be eliminated, e.g. if $S_1 = \mathcal{F}_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle}^1 \neq \emptyset$ then all attacks for which $\neg(x \rightarrow^{S_1} y)$ can be removed from \mathcal{A} . Define the set $R \subseteq \mathcal{A}$ to contain precisely those attacks which are discarded in the process of building $ge_{\text{EAF}}(\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle)$, i.e., $R^0 = \emptyset$, $R^i = \{\langle x, y \rangle : \neg(x \rightarrow^{S_i} y)\}$ where $S_i = \mathcal{F}_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle}^i$. Defining the weighted argument system $\langle \mathcal{X}, \mathcal{A}, w \rangle$ obtained from $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$ to have weight function $w(\langle x, y \rangle) = 1$ if $\langle x, y \rangle \in R$ and $w(\langle x, y \rangle) = |\mathcal{X}|^2 + 1$ otherwise, we immediately obtain (using $\beta = |R|$) that

$$ge_{\text{EAF}}(\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle) \in \mathcal{E}_{gr}^{\text{WT}}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$$

and, in fact,

$$\mathcal{E}_{gr}^{\text{EAF}}(\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle) = \mathcal{E}_{gr}^{\text{WT}}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta) \setminus \mathcal{E}_{gr}^{\text{WT}}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta - 1)$$

In summary given an EAF, $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$ a weighted system $\langle \mathcal{X}, \mathcal{A}, w \rangle$ and budget β can be specified in which the grounded extension of $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$ is recovered. We are left with the question of whether $\mathcal{E}_{gr}^{\text{WT}}(\langle \mathcal{X}, \mathcal{A}, w \rangle, \beta)$ can be described within a related EAF.

Consider the very basic system of Figure 13.

With an inconsistency budget of $\beta = 1$, we obtain $\mathcal{E}_{gr}^{\text{WT}}(W_5, 1) = \{\emptyset, \{x\}, \{y\}\}$. While a seemingly natural choice of EAF corresponding to W_5 is to add arguments (A and B) attacking each weighted attack (such arguments potentially representing as cases for disposing of either attack), because of the conditions on arguments involved in attacking *mutual attacks* the resulting EAF will have an empty grounded extension.

We note that similar behaviour will be observed when we have mutually attacking arguments for which the weights of such attacks are equal.




Figure 13: Weighted (W_5) vs. Extended Argument Frameworks

In summary, we have considered four approaches — PAFs, VAFs, resolution-based, and EAFs — each of which offers a formal basis under which attacks within a standard argument system may be discounted and which have promoted as the relevant semantics of interest concepts based on the grounded extension. In this article we have introduced another approach — weighted argument systems — offering a rationale for discounting attacks and focusing on grounded extensions of the reduced systems. For each of the alternative mechanisms we have argued that the \mathcal{E}_{gr}^{AUG} can be described by a weighted version of the underlying Dung-style framework and suitable inconsistency budget. In contrast, however, we can describe quite basic weighted systems and budgets for which \mathcal{E}_{gr}^{WT} is not encapsulated by replacing weights and budget by an alternative simplifying mechanism.

8. Conclusions

Several possibilities suggest themselves for future research. The first is to investigate specific interpretations for weights, as suggested in the paper. Another is to investigate the framework experimentally, to obtain a better understanding of the way the approach behaves. One obvious issue here is to look for “discontinuities” as the inconsistency budget grows, i.e., large increases in the number of accepted arguments for only a small increase in inconsistency budget. Finally, it would be interesting to consider further the relationship between our framework and the various other argumentation frameworks and semantics that have been proposed in the literature, as discussed in section 7.

Acknowledgments

An earlier version of this paper was published at AAMAS-2009.

References

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

- [2] L. Amgoud, N. Maudet, and S. Parsons. Modelling dialogues using argumentation. In E. Durfee, editor, *Proceedings of the Fourth International Conference on Multi-Agent Systems*, pages 31–38, Boston, MA, USA, 2000. IEEE Press.
- [3] K. J. Arrow, A. K. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare Volume 1*. Elsevier Science Publishers B.V.: Amsterdam, The Netherlands, 2002.
- [4] F. Bacchus, A. J. Grove, J. Y. Halpern, and D. Koller. From statistical knowledge bases to degrees of belief. *Artificial Intelligence*, 87:75–143, 1996.
- [5] P. Baroni, F. Cerutti, M. Giacomin, and G. Guida. Encompassing attacks to attacks in abstract argumentation frameworks. In *Proc. ECSQARU*, volume 5590 of *LNAI*, pages 83–94. Springer-Verlag: Berlin, Germany, 2009.
- [6] P. Baroni, P. E. Dunne, and M. Giacomin. Computational properties of resolution-based grounded semantics. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI-09)*, pages 683–689, Pasadena, CA, 2009.
- [7] P. Baroni and M. Giacomin. Resolution-based argumentation semantics. In *Proc. 2nd COMMA, FAIA 172*. IOS Press, 2008.
- [8] P. Baroni and M. Giacomin. Semantics of abstract argument systems. In I. Rahwan and G. R. Simari, editors, *Argumentation in Artificial Intelligence*, pages 25–44. Springer, Berlin, 2009.
- [9] H. Barringer, D. M. Gabbay, and J. Woods. Temporal dynamics of support and attack networks: From argumentation to zoology. In *Mechanizing Mathematical Reasoning (LNCS Volume 2605)*, pages 59–98. Springer-Verlag: Berlin, Germany, 2005.
- [10] T. J. M. Bench-Capon. Persuasion in practical argument using value based argumentation frameworks. *Journal of Logic and Computation*, 13(3):429–448, 2003.
- [11] T. J.M. Bench-Capon, S. Doutre, and P. E. Dunne. Audiences in argumentation frameworks. *Artificial Intelligence*, 171(1):42–71, 2007.
- [12] L. Bertossi, A. Hunter, and T. Schaub, editors. *Inconsistency Tolerance (LNCS Volume 3300)*. Springer-Verlag: Berlin, Germany, 2004.
- [13] Ph. Besnard and A. Hunter. *Elements of Argumentation*. MIT Press, 2008.
- [14] H. Burhrman, J. Kadin, and T. Thierauf. Functions computable with nonadaptive queries to NP. *Th. Comput. Systems*, 31:77–92, 1998.

- [15] C. Cayrol and M.-C. Lagasquie-Schiex. Gradual valuation for bipolar argumentation frameworks. In *Proceedings of the 8th European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, volume 3571 of *LNCS*, pages 366–377. Springer, 2005.
- [16] S. Coste-Marquis, C. Devred, and P. Marquis. Symmetric argumentation frameworks. In *Proc. 8th ECSQARU, (LNAI Volume 3571)*. Springer-Verlag: Berlin, Germany, 2005.
- [17] Y. Dimopolous and A. Torres. Graph theoretical structures in logic programs and default theories. *Theoretical Computer Science*, 170:209–244, 1996.
- [18] D. Dubois, J. Lang, and H. Prade. A possibilistic truth-maintenance system with uncertain justifications, and its application to belief revision. In J. P. Martins and M. Reinfrank, editors, *Truth Maintenance Systems*, number 515 in *Lecture Notes in Computer Science*, pages 87–106. Springer Verlag, 1991.
- [19] P. M. Dung. On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n -person games. *Artificial Intelligence*, 77:321–357, 1995.
- [20] P. E. Dunne. Computational properties of argument systems satisfying graph-theoretic constraints. *Artificial Intelligence*, 171:701–729, 2007.
- [21] P. E. Dunne. The computational complexity of ideal semantics I: abstract argumentation frameworks. In *Proc. 2nd COMMA, FAIA 172*, pages 147–158. IOS Press, 2008.
- [22] P. E. Dunne. The computational complexity of ideal semantics. *Artificial Intelligence*, 173:1559–1591, 2009.
- [23] P. E. Dunne and T. Bench-Capon. Coherence in finite argument systems. *Artificial Intelligence*, 141:187–203, 2002.
- [24] P. E. Dunne and M. Caminada. Computational complexity of semi-stable semantics in abstract argumentation frameworks. In *Proc. 11th JELIA, (LNAI Volume 5292)*. Springer-Verlag: Berlin, Germany, 2008.
- [25] P. E. Dunne, S. Modgil, and T. Bench-Capon. Computation in extended argumentation frameworks. In *Proc. ECAI 2010 (to appear)*, 2010.
- [26] P. E. Dunne and M. Wooldridge. Complexity of abstract argumentation. In I. Rahwan and G. Simari, editors, *Argumentation in Artificial Intelligence*. Springer-Verlag: Berlin, Germany, 2009.
- [27] W. Dvorak and S. Woltran. Complexity of semi-stable and stage semantics in argumentation frameworks. *Inf. Proc. Letters*, 110:425–430, 2010.

- [28] B. Jenner and J. Toran. Computing functions with parallel queries in NP. *Theoretical Computer Science*, 141:175–193, 1995.
- [29] B. Jenner and J. Toran. The complexity of obtaining solutions for problems in np and nl. In *Complexity theory retrospective II*, pages 155–178. Springer-Verlag: Berlin, Germany, 1997.
- [30] P. Krause, S. Ambler, M. Elvang-Gøransson, and J. Fox. A logic of argumentation for reasoning under uncertainty. *Computational Intelligence*, 11 (1):113–131, 1995.
- [31] M. Krentel. The complexity of optimization problems. *Journal of Computer and System Sciences*, 36(3):761–767, 1988.
- [32] D. Martinez, A. Garcia, and G. Simari. An abstract argumentation framework with varied-strength attacks. In *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR'08)*, 2008.
- [33] P. Matt and F. Toni. A game-theoretic measure of argument strength for abstract argumentation. In *Proceedings of 11th European Conference on Logics in Artificial Intelligence (JELIA'08)*, volume 5293 of *LNAI*, pages 285–297. Springer, 2008.
- [34] S. Modgil. An abstract theory of argumentation that accommodates defeasible reasoning about preferences. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, volume 4724 of *LNCS*. Springer, 2007.
- [35] S. Modgil. Labellings and games for extended argumentation frameworks. In *Proc. 21st IJCAI*, pages 873–878, 2009.
- [36] S. Modgil. Reasoning about preferences in argumentation frameworks. *Artificial Intelligence*, 173(9–10):901–934, 2009.
- [37] S. Modgil and T. Bench-Capon. Integrating object and meta-level value based argumentation. In *Computational Models of Argument: Proceedings of COMMA 2008*, pages 240–251. IOS Press, 2008.
- [38] E. Neufeld and D. Poole. Towards solving the multiple extension problem: combining defaults and probability. In L. N. Kanal, T. S. Levitt, and J. F. Lemmer, editors, *Uncertainty in Artificial Intelligence 3*, pages 35–44. Elsevier Science Publishers, Amsterdam, The Netherlands, 1989.
- [39] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley: Reading, MA, 1994.
- [40] C. H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization*. Prentice Hall International: Hemel Hempstead, England, 1982.

- [41] S. Parsons. On precise and correct qualitative probabilistic reasoning. *International Journal of Approximate Reasoning*, 35:111–135, 2004.
- [42] I. Pinyol and J. Sabater-Mir. An argumentation-based dialog for social evaluations exchange. In M. J. Wooldridge, editor, *Nineteenth European Conference on Artificial Intelligence (ECAI 2010)*, Lisbon, Portugal, August 2010. ECCAI.
- [43] D. Poole. Probabilistic horn abduction and Bayesian networks. *Artificial Intelligence*, 64:81–129, 1993.
- [44] H. Prakken and G. Sartor. Argument-based logic programming with defeasible priorities. *Journal of Applied Non-classical Logics*, 7:25–75, 1997.
- [45] I. Rahwan and K. Larson. Pareto optimality in abstract argumentation. In *Proceedings of 23rd Conference on Artificial Intelligence (AAAI'08)*. AAAI Press, 2008.
- [46] I. Rahwan and G. R. Simari, editors. *Argumentation in Artificial Intelligence*. Springer-Verlag: Berlin, Germany, 2009.
- [47] G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, NJ, 1976.
- [48] W. Spohn. A general non-probabilistic theory of inductive reasoning. In R. D. Shachter, T. S. Levitt, L. N. Kanal, and J. F. Lemmer, editors, *Uncertainty in Artificial Intelligence 4*, pages 149–158. Elsevier Science Publishers, Amsterdam, The Netherlands, 1990.
- [49] V. Tamma and S. Parsons. Argumentation and qualitative reasoning with kappa calculus. In *Proceedings of the 6th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, volume 2143 of *LNCS*, pages 680–691. Springer, 2001.
- [50] F. A. Tohmé, G. A. Bodanza, and G. R. Simari. Aggregation of attack relations: a social-choice theoretical analysis. In S. Hartmann and G. Kern-Isberner, editors, *Foundations of Information and Knowledge Systems (FoIKS 2008)*, pages 8–23, Pisa, Italy, 2008.