

A generalization of Dung’s Abstract Framework for Argumentation

Arguing with Sets of Attacking Arguments

Søren Holbech Nielsen¹ and Simon Parsons²

¹ Department of Computer Science
Aalborg University, Aalborg
Denmark
holbech@cs.aau.dk

² Department of Computer and Information Science
Brooklyn College, City University of New York
Brooklyn, 11210 NY, USA
parsons@sci.brooklyn.cuny.edu

Abstract. One of the most widely studied systems of argumentation is the one described by Dung in a paper from 1995. Unfortunately, this framework does not allow for joint attacks on arguments, which we argue must be required of any truly abstract argumentation framework. A few frameworks can be said to allow for such interactions among arguments, but for various reasons we believe that these are inadequate for modelling argumentation systems with joint attacks. In this paper we propose a generalization of the framework of Dung, which allows for sets of arguments to attack other arguments. We extend the semantics associated with the original framework to this generalization, and prove that all results in the paper by Dung have an equivalent in this more abstract framework.

1 Introduction

In the last fifteen years or so, there has been much interest in argumentation systems within the artificial intelligence community³. This interest spreads across many different sub-areas of artificial intelligence. One of these is non-monotonic reasoning [10, 19], which exploits the fact that argumentation systems can handle, and resolve, inconsistencies [12, 13] and uses it to develop general descriptions of non-monotonic reasoning [8, 18]. This line of work is summarised in [28]. Another area that makes use of argumentation is reasoning and decision making under uncertainty [5, 16, 17], which exploits the dependency structure one can infer from arguments in order to correctly combine evidence. Much of this work is covered in [9]. More recently [23, 26], the multi-agent systems community has begun to make use of argumentation, using it to develop a notion of rational interaction [4, 20].

One very influential system of argumentation was that introduced by Dung [11]. This was, for instance, the basis for the work in [8], was the system extended by Amgoud in [1, 2], and subsequently as the basis for the dialogue systems in [3, 24]. In [11],

³ There were AI researchers who were interested in argumentation before this, for example [6, 7, 14, 21], but this interest was very localized.

Dung presents a very abstract framework for argumentation and a series of semantics for this framework. He goes on to prove a series of relationships between his framework and different varieties of formal logics, including a proof that logic programming can be seen as a special case of his framework. As a last result of the paper he provides a method for encoding systems of the argumentation framework as logic programs. The importance of Dung's results is mainly due to the fact that his framework abstracts away from details of language and argumentation rules, that the presented semantics therefore are clear and intuitive, and that relationships among arguments can be analysed in isolation from other (e.g. implicational) relationships. Furthermore, the results can easily be transferred to any other argumentation framework, by identifying that framework's equivalent of an attack. It is this generality, we believe, that has contributed to the popularity of the work, and we see it as a prime contender for becoming an established standard for further investigations into the nature of arguments and their interaction.

However, even though Dung tried to abstract away from the underlying language and structure of arguments, he did not succeed in doing so completely. In fact if his framework is expected to be able to model all possible kinds of attack, there is an implicit assumption that the underlying language contains a logical "and" connective. This hidden assumption arises from that fact that Dung's attack relation is a simple binary relation from one argument to another, rather than a relation mapping sets of arguments to other sets of arguments.

While not explicitly analyzing the fundamental problem of Dung's framework, some previous works, most notably the efforts of Verheij, have allowed for sets of attacking arguments, although mostly as side effects. We do not find these solutions fully satisfying, and none of them can be said to be conservative generalizations of the framework of [11], that is a generalization that makes the minimum changes to the Dung framework necessary to allow it to handle sets of attacking arguments. We elaborate further on this throughout the paper.

In this paper we analyze Dung's framework, and point out the hidden assumption on the underlying language. We present a generalization of Dung's framework, keeping as close to his ideas as possible, which frees the underlying language from being closed under some logical "and" connective. We do this by allowing sets of arguments to attack single arguments, and provide new definitions and proofs mirroring Dung's results for this more general framework. We also argue why allowing sets of arguments to attack other sets of arguments does not provide further flexibility, and provide an automated encoding of systems of the new framework in Prolog, mirroring Dung's encoding of his systems as logic programs.

The paper is organized as follows: In Sect. 2 we present the essentials of Dung's framework, and then through examples illustrate how a more general attack relation is needed for a truly abstract framework. Then, in Sect. 3 we present our generalization of Dung's framework, complete with definitions, proofs, and a Prolog encoding method. Following this, in Sect. 4, we review other works on argumentation systems where sets of arguments can attack other arguments, and relate them to the approach presented in this paper. Finally, we conclude on the work presented here. Throughout the paper we use the term argumentation *system*, where [11] uses argumentation *framework*, to denote the actual mathematical structures we work with. The term *framework* we reserve

for denoting the overall approaches to describing and reasoning about the argumentation systems, such as the one represented by [11] and the ones reviewed in Sect. 4.

2 Dung’s Framework

Dung [11] defines an *argumentation system* as a pair (A, \triangleright) , where A is a set of *arguments*, and $\triangleright \subseteq A \times A$ is an *attack relation*. If for two arguments A and B we have $A \triangleright B$, then we say that A *attacks* B , and that B is *attacked by* A . As examples, we might consider the following as arguments:

- E_1 “Joe does not like Jack”,
- E_2 “There is a nail in Jack’s antique coffee table”,
- E_3 “Joe hammered a nail into Jack’s antique coffee table”,
- E_4 “Joe plays golf, so Joe has full use of his arms”, and
- E_5 “Joe has no arms, so Joe cannot use a hammer, so Joe did not hammer a nail into Jack’s antique coffee table”.

As can be seen it is not required of an argument that it follows the “if X then conclude Y” pattern for reasoning, or, for that matter, that it represents sound reasoning.

As examples of attacks, we could have that $E_5 \triangleright E_3$, $E_3 \triangleright E_5$, and $E_4 \triangleright E_5$. Intuitively, and in any common-sense argumentation system, we would expect that $A \triangleright B$ if the validity of the argument A is somehow obstructing B from being valid.

Without loss of generality, we will assume that the arguments are members of some underlying language L . This assumption is necessary if any kind of meaning is to be extracted from an argumentation system. For instance, in our example, L would necessarily include the strings represented by E_1 to E_5 .

It seems reasonable that sometimes a number of arguments can interact and constitute a stronger attack on one or more of the other arguments. For instance, the two arguments E_1 and E_2 would jointly (but not separately) provide a case for the conclusion that Joe has struck a nail into Jack’s antique coffee table, and thus provide a joint attack on argument E_5 , which has the opposite conclusion. The principle of synergy among arguments is not new, and has previously been debated in connection to “accrual of arguments” (see e.g. [25, 27, 29]). The difference between that discussion and the issue addressed here is we (and Dung) do not consider arguments as having a numerical strength, and a set of defeated arguments thus cannot accrue to become undefeated, unless that set is explicitly specified to defeat each argument defeating its individual members.

Going back to the example, if this synergy is to be modelled under Dung’s limitations, somehow there must be a new argument:

E_6 : “Joe does not like Jack *and* there is a nail in Jack’s antique coffee table”,

which attacks E_5 . If this is taken to be a general solution, it is obviously required that the underlying language L is closed under some “and”-connective.

Furthermore, what we meant to state was that E_1 and E_2 jointly attacked E_5 and the solution does not quite suffice: It may turn out that \triangleright is defined in such a manner that one (or both) of E_1 and E_2 is attacked by another valid argument, while E_6 is

not. That would mean that “Joe does not like Jack *and* there is a nail in Jack’s coffee table” is a valid argument, whereas, say, “Joe does not like Jack” is not. Clearly this is nonsense, and in order to ensure that nonsense conclusions cannot arise, \triangleright would have to be restricted accordingly. This muddles the clear distinction between arguments and attacks, which was the very appeal of Dung’s framework.

These underlying consistency relations between arguments would seemingly be good candidates for encoding in a logical language (for example $E_1 \wedge E_2 \Rightarrow E_6$ and $E_6 \Rightarrow E_1$), and in fact an underlying logical language employing standard negation could be used to model sets of attacking arguments (i.e. $E_1 \wedge E_2 \Rightarrow \neg \text{conclusion}(E_5)$) with attack relations $\neg \text{conclusion}(A) \triangleright A$ for all arguments A), but we chose not to go this route for a number of reasons. Primarily, it adds a another level of interdependencies between arguments, which makes it hard to survey the effects of one set of argument on others and calls for more specialized formalisms for analysis than Dung’s. Moreover, examples of joint undercutting attacks seems to be inherently argumentative in nature, and only obscurely encoded in an implicative manner. Consider for instance the following arguments:

- F_1 “The Bible says that God is all good, so God is all good”,
- F_2 “The Bible was written by human beings”, and
- F_3 “Humans beings are not infallible”.

F_2 and F_3 attacks the validity of F_1 , but clearly it makes no sense to encode this as $F_2 \wedge F_3 \Rightarrow \neg \text{conclusion}(F_1)$ as the facts that human beings are not to be considered infallible and that some of them just happened to write the Bible, do not entail that God is not all good. To capture the intended meaning of the attack, one would have to add an explicit presumption, like “The Bible can be trusted on all matters” to F_1 , and allow for such assumptions to be targets of attacks, which — besides requiring identification of all such implicit assumptions — can hardly be said to be as elegant as allowing attacks at the argumentative level⁴

Having argued for the necessity of allowing a set of arguments to attack another argument, we now examine settings where an entire set of arguments is attacked by either a single argument or another set of arguments. Without loss of generality, we assume that what is needed is an attack

$$\{A_1, \dots, A_n\} \triangleright \{B_1, \dots, B_m\},$$

such that the validity of all the A -arguments prevents the B -arguments from being valid. There are two distinct manners in which this can be interpreted:

1. Either the validity of the A -arguments means that each B_i cannot be valid, no matter the validity of the other B -arguments, or
2. the validity of the A -arguments mean that not all of the B -arguments can be valid at the same time.

⁴ Those swayed more by practical considerations than examples should note that the original motivation for this work was to allow arguments about Bayesian networks, in which sets of attacking arguments very naturally occur.

[29] refers to these as “collective” and “indeterministic defeat”, respectively — a terminology we adopt in this text.

As an example consider the following twist on the story about Jack, Joe, and the antique coffee table:

E_7 “Jack has been telling lies about Joe to Jill”

E_8 “Jack is a rabbit”

E_9 “Joe loves all animals”

If E_8 is a valid argument, then none of the arguments in the set $\{E_3, E_7\}$ can be valid: E_3 because rabbits do not own antique coffee tables, and E_7 because rabbits, being unable to speak, do not lie. This is thus an example of collective defeat. As an example of indeterministic defeat, E_9 attacks the set of arguments $\{E_1, E_8\}$ seen as a set: E_1 and E_8 cannot both be valid arguments if Joe loves all animals. However, both E_1 and E_8 can be valid seen as individual arguments, no matter how Joe feels about animals.

We claim that it is never necessary to specify a non-singleton set of arguments as attacked, as in $\{A_1, \dots, A_n\} \triangleright \{B_1, \dots, B_m\}$: If collective defeat is taken to heart, the attack can be reformulated as a series of attacks

$$\begin{aligned} \{A_1, \dots, A_n\} \triangleright B_1 \\ \vdots \\ \{A_1, \dots, A_n\} \triangleright B_m . \end{aligned}$$

It is easily seen that the above attacks would imply the attack, which is intended, as the validity of the A -arguments would ensure that none of the B -arguments are valid.

If instead indeterministic defeat is required, the attack can be reformulated as

$$\{A_1, \dots, A_n, B_2, \dots, B_m\} \triangleright B_1 ,$$

which ensures that in case the A -arguments are valid, then B_1 cannot be a valid argument if the remaining B -arguments are also true, thus preventing the entire set of B -arguments from being valid at once, if the A -arguments are true. In the example above, we would state that $\{E_8, E_1\}$ attacks E_9 . Notice that this “trick” is not dependent on the actual structure or language of the arguments, nor require the introduction of a new dummy argument, as was the case if only single arguments were allowed as attackers.

In conclusion, we have argued for the insufficiency of Dung’s treatment, when sets of arguments are taken into account, and that an attack relation that allows for sets of arguments attacking single arguments is sufficient to capture any kind of relation between sets of arguments.

3 Argumentation with Attacking Sets of Arguments

In this section we present our generalization of the framework of [11]. In an effort to ease comparison, we have labelled definitions, lemmas, and theorems with the same numbers as their counterparts in [11], even if this means that there are holes in the

numbering (e.g. there is no Lemma 2). Furthermore, we have omitted proofs where the original proofs of [11] suffice. As a result of the tight integration with [11] most definitions and results have been worded in a nearly identical manner, even if the proofs are different and the meaning of individual words are different. Those definitions and results that differ essentially from their counterparts in [11], or which is entirely new, have been marked with an asterix (*). The rest are identical to those in [11].

Throughout the presentation, it should be clear that the framework presented here reduces to that of [11] if only singleton sets are allowed as attackers.

Definition 1 (Argumentation System*). An argumentation system is a pair (A, \triangleright) , where A is a set of arguments, and $\triangleright \subseteq (\mathcal{P}(A) \setminus \{\emptyset\}) \times A$ is an attack relation.

We say that a set of arguments S attacks an argument A , if there is $S' \subseteq S$ such that $S' \triangleright A$. In that case we also say that A is attacked by S . If there is no set $S'' \subsetneq S'$ such that S'' attacks A , then we say that S' is a *minimal* attack on A . Obviously, if there exists a set that attacks an argument A , then there must also exist a minimal attack on A . If for two sets of arguments S_1 and S_2 , there is an argument A in S_2 that is attacked by S_1 , then we say that S_1 attacks S_2 , and that S_2 is attacked by S_1 .

Definition 2 (Conflict-free Sets*). A set of arguments S , is said to be conflict-free if it does not attack itself, i.e. there is no argument $A \in S$, such that S attacks A .

Let S_1 and S_2 be sets of arguments. If S_2 attacks an argument A , and S_1 attacks S_2 , then we say that S_1 is a *defense* of A from S_2 , and that S_1 *defends* A from S_2 . Obviously, if S_3 is a superset of S_1 , S_3 is also a defense of A from S_2 .

Definition 3 (Acceptable and Admissable Arguments*). An argument A is said to be acceptable with respect to a set of arguments S , if S defends A from all attacking sets of arguments in A .

A conflict-free set of arguments S is said to be admissable if each argument in S is acceptable with respect to S .

Intuitively, an argument A is acceptable with respect to some set S , if anyone believing in the validity of the arguments in S can defend A against all attacks. If a set of arguments is admissable, it means that anyone believing this set of arguments as valid is not contradicting himself and can defend his beliefs against all attacks.

Definition 4. An admissable set S is called a preferred extension if there is no admissable set $S' \subseteq A$, such that $S \subsetneq S'$.

Building on the intuition from before, taking on a preferred extension as your beliefs thus means that you would not be able to defend any more arguments without contradicting yourself.

Lemma 1 (Fundamental Lemma). Let S be an admissable set, and A and A' be arguments that each are acceptable with respect to S , then

1. $S' = S \cup \{A\}$ is admissable, and
2. A' is acceptable with respect to S' .

Proof. 1) As \mathcal{S} is admissible, and A is acceptable with respect to \mathcal{S} , it is obvious that \mathcal{S} , and therefore also \mathcal{S}' , defends each argument in \mathcal{S}' . Thus we only need to prove that \mathcal{S}' is conflict-free. Assume not. Then there is an argument $B \in \mathcal{S}'$ and an attack $\mathcal{S}'' \subseteq \mathcal{S}'$ on B . Since each argument in \mathcal{S}' is defended by \mathcal{S} it follows that \mathcal{S} attacks \mathcal{S}'' .

As \mathcal{S} attacks \mathcal{S}'' it follows that \mathcal{S} must attack at least one argument of \mathcal{S}'' . Let C be this argument. We consider two cases: First $C \equiv A$ and second $C \not\equiv A$. If $C \equiv A$ then it follows that \mathcal{S} attacks A . As A is acceptable with respect to \mathcal{S} , \mathcal{S} must then necessarily attack \mathcal{S} , which contradicts the assumption that \mathcal{S} is conflict-free. Assume then that $C \not\equiv A$. Then C must be part of \mathcal{S} , and consequently \mathcal{S} attacks \mathcal{S} yielding the same contradiction with the assumptions.

2) Obvious. □

Using the Fundamental Lemma the following important result, guaranteeing that an admissible set can be extended to a preferred extension, can be proven.

Theorem 1. *For any argumentation system the set of admissible sets forms a complete partial order with respect to set inclusion, and for each admissible set \mathcal{S} there exists a preferred extension \mathcal{S}' , such that $\mathcal{S} \subseteq \mathcal{S}'$.*

As the empty set is an admissible set, we have:

Corollary 2. *Every argumentation system has at least one preferred extension.*

A more aggressive semantics is the stable semantics:

Definition 5 (Stable Semantics). *A conflict free set \mathcal{S} is a stable extension if \mathcal{S} attacks all arguments in $\mathcal{A} \setminus \mathcal{S}$.*

Lemma 3. *\mathcal{S} is a stable extension iff $\mathcal{S} = \{A \mid A \text{ is not attacked by } \mathcal{S}\}$.*

Proof. “only if”: Obvious.

“if”: Assume not. Then \mathcal{S} is either not conflict-free, or there is an argument in $\mathcal{A} \setminus \mathcal{S}$ not attacked by \mathcal{S} . The latter possibility is precluded by the definition of \mathcal{S} , so there must be a set $\mathcal{S}' \subseteq \mathcal{S}$ and an argument $A \in \mathcal{S}$ such that \mathcal{S}' attacks A . But then \mathcal{S} also attacks A , which contradicts the definition of \mathcal{S} . □

The general connection between stable and preferred semantics is given by the following result:

Lemma 4. *Every stable extension is a preferred extension, but not vice versa.*

Both preferred and stable semantics are credulous in the sense that they represent beliefs that include as much as possible. Next, we consider semantics corresponding to more skeptical points of views. For this we need the notion of a characteristic function, and some general results on this:

Definition 6 (Characteristic Function). *The characteristic function of an argumentation system is the function $F : \mathcal{P}(\mathcal{A}) \rightarrow \mathcal{P}(\mathcal{A})$ defined as*

$$F(\mathcal{S}) = \{A \mid A \text{ is acceptable wrt. } \mathcal{S}\} .$$

Next, we state a couple of properties of the characteristic function F . The first result is not explicitly stated in [11], but included only as part of a proof. We make it explicit here as it is a property required of F by some proofs that have been left out.

Proposition 1 (*) *If S is a conflict-free set, then $F(S)$ is also conflict-free.*

Proof. Assume this is not the case, then there is $S' \subseteq F(S)$ and $A \in F(S)$ such that S' attacks A . Since A is acceptable wrt. S , S must attack at least one element B of S' . But since B is in $F(S)$ it must be acceptable wrt. S , and S must consequently attack itself. This contradicts the assumption that S is a conflict-free set. \square

Lemma 5. *A conflict-free set S is admissible iff $S \subseteq F(S)$.*

Proof. “only if”: All arguments of S are acceptable wrt. S , so $S \subseteq F(S)$.
 “if”: As $S \subseteq F(S)$ it follows that all arguments of S are acceptable wrt. S . \square

Lemma 6. *F is a monotonic function with respect to set inclusion.*

Proof. Follows since adding arguments to a set of arguments cannot cause the set to attack fewer arguments, and consequently cannot change the status of any of the arguments currently defended into being not defended. \square

Now, we can introduce the most skeptical semantics possible:

Definition 7 (Grounded Extension). *The grounded extension of an argumentation system, is the least fix-point of the corresponding characteristic function.*

A grounded extension is thus the set of arguments that are not challenged by any other arguments, along with the arguments defended by these arguments, those defended by those, and so on. [11] does not prove that the grounded extension of an argumentation system is well-defined, but we include a proof here.

Proposition 2 (*) *If G_1 and G_2 are both grounded extensions of an argumentation system, then $G_1 = G_2$.*

Proof. Assume not, and let $C = G_1 \cap G_2$. As G_1 and G_2 are different and also minimal, it follows that none of them can be the empty set, and hence that $F(\emptyset) \neq \emptyset$. As $F(\emptyset)$ consists of the arguments that are not attacked by any arguments at all, it follows that these are acceptable wrt. any set. In particular, $F(\emptyset)$ must be a subset of both G_1 and G_2 , so C is non-empty. Furthermore, as Lmm. 6 assures that F is monotonic, it follows that $F(C)$ must be a subset of both G_1 and G_2 . But then $F(C)$ must be equal to C , and is thus a fix point of F . As both G_1 and G_2 were supposed to be minimal and different, this yields the desired contradiction. \square

As a common class, encompassing all the semantics we have discussed so far, we introduce complete extensions:

Definition 8 (Complete Extension). *An admissible set S is called a complete extension, if all arguments that are acceptable with respect to S are in S .*

A couple of results tie the complete extension semantics to the other semantics we have discussed:

Lemma 7. *A conflict-free set S is a complete extension iff $S = F(S)$.*

Theorem 2. *Extensions are such that:*

1. *Each preferred extension is a complete extension, but not vice versa.*
2. *The grounded extension is the least complete extension with respect to set inclusion.*
3. *The complete extensions form a complete semi-lattice with respect to set inclusion.*

Next, we investigate classifying argumentation systems according to desirable properties of their corresponding semantics.

Definition 9 (Finitary System*). *An argumentation system is said to be finitary if for each argument A , there is at most a finite amount of minimal attacks on A , and each minimal attack is by a finite set of arguments.*

Lemma 8. *For any finitary system, F is ω -continuous.*

Proof. Let $S_1 \subseteq S_2 \subseteq \dots$ be an increasing series of sets of arguments, and $S = \cup_i S_i$. We need to show that $F(S) = \cup_i F(S_i)$. As adding arguments to a set cannot reduce the set of arguments attacked by this set, and therefore cannot reduce the set of arguments that are acceptable with respect to it, we have that $F(S_i) \subseteq F(S)$ for each i , and thus $F(S) \supseteq \cup_i F(S_i)$.

To see that $F(S) \subseteq \cup_i F(S_i)$, consider an argument $A \in F(S)$, and let T_1, \dots, T_n be the finitely many minimal attacks on A . As S attacks each attack on A , there must be an argument B_i in each T_i , which is attacked by S . Let $U_i \subseteq S$ be the minimal attack of B_i . As each minimal attack consists of a finite number of arguments, the set $U = U_1 \cup \dots \cup U_n$ is finite as well, and thus there must be a j , such that $U \subseteq S_j$. Consequently, A must be in $F(S_j)$ and therefore also in $\cup_i F(S_i)$. \square

Definition 10 (Well-founded System*). *An argumentation system is well-founded, if there exists no infinite sequence of sets S_1, S_2, \dots , such that S_i is a minimal attack on an argument in S_{i-1} for all i .*

Theorem 3. *Every well-founded argumentation system has exactly one complete extension, which is grounded, preferred, and stable.*

Proof. It suffices to prove that the grounded extension G is stable. Assume this is not the case, and let $S = \{A \mid A \notin G \text{ and } A \text{ is not attacked by } G\}$, which must be nonempty if the grounded extension is not stable. We prove that each argument A in S is attacked by a minimal set S' such that $S \cap S' \neq \emptyset$, and therefore that the system cannot be well-founded.

Since A is not in G it is not acceptable with respect to G . Therefore there must be a minimal attack T of A , not itself attacked by G . Since G does not attack A , at least one element of T must be outside of G . Let T' be $T \setminus G$, which is thus non-empty. As G does not attack T , it furthermore follows that T' must be a subset of S . Thus, T is the set S' we were looking for, and the proof is complete. \square

Definition 11 (Coherent and Relatively Grounded System). *An argumentation system is coherent if all its preferred extensions are stable. A system is relatively grounded if its grounded extension is the intersection of all its preferred extensions.*

Let A_1, A_2, \dots be a (possible finite) sequence of arguments, where each argument A_i is part of a minimal attack on A_{i-1} . Then the arguments $\{A_{2i}\}_{i \geq 1}$ are said to *indirectly attack* A_1 . The arguments $\{A_{2i-1}\}_{i \geq 1}$ are said to *indirectly defend* A_1 . If an argument A is both indirectly attacking and defending an argument B , then A is said to be *controversial with respect to* B , or simply *controversial*.

Definition 12 (Uncontroversial and Limited Controversial System). *An argumentation system is uncontroversial if none of its arguments are controversial. An argumentation system, for which there exists no infinite sequence of arguments A_1, A_2, \dots , such that for all i , A_i is controversial with respect to A_{i-1} , is said to be limited controversial.*

Obviously, a controversial argumentation system is also limited controversial.

Lemma 9. *In every limited controversial argumentation system there exists a nonempty complete extension.*

Proof. We construct the nonempty complete extension C . Since a nonempty grounded extension would suffice, we assume that it is empty. Since the system is limited controversial, every sequence of arguments, where A_i is controversial with respect to A_{i-1} , must have a last element, B . It follows that there is no argument that is controversial with respect to B . We define E_0 to be $\{B\}$, and E_i to be $E_{i-1} \cup D_i$, where D_i is a minimal set that defends E_{i-1} from $A \setminus E_{i-1}$, for all $i \geq 1$. As the grounded extension is empty, each argument is attacked by some other argument, and therefore each D_i is guaranteed to exist.

We then prove by induction that, for each $i \geq 0$, E_i is conflict-free and each argument in E_i indirectly defends B .

The hypothesis trivially holds true for $i = 0$. We assume it to be true for $i - 1$ and show that it also must be true for i : From the induction hypothesis we know that E_{i-1} consists of arguments that indirectly defend B . As each argument in D_i participates in attacking an argument, which participates in an attack on an argument in E_{i-1} , each of these must also indirectly defend B , and consequently this is true of all arguments in E_i . Assume then that E_i is not conflict-free. Then there is a set of arguments $S \subseteq E_i$, that attack an argument $B \in E_i$. But then the arguments in S are attacking an indirect defender of B , and thus are indirect attackers of B . This means that the arguments in S are controversial with respect to B , violating the assumptions of the lemma. Thus, the induction hypothesis is proved.

Next, let $E = \cup_i E_i$. We prove that this set is admissible, and then let C be the least complete extension containing E . We know such an extension exists as by Thm. 1 a preferred extension containing E must exist, and from Thm. 2 that extension must be a complete extension. To see that E is admissible, first let $C \in E$ be an argument. There must be some i , such that $C \in E_i$, and therefore a defense of C must be in D_i , and consequently in E_{i+1} . But then that defense is also in E , and hence C must be acceptable with respect to E . To see that E is conflict-free assume that it contains C and S , such that S attacks C . As each argument of S must be an element of some set

E_i , it follows that each of these indirectly defend B . But as C also indirectly defends B , each element of S must indirectly attack B also, and is thus controversial with respect to B . But this violates the assumption that no argument is controversial with respect to B , and there can therefore be no such S and C . \square

Lemma 10. *For any uncontroversial system, with an argument A that is neither a member of the grounded extension nor attacked by it,*

1. *there exists a complete extension containing A , and*
2. *there exists a complete extension that attacks A .*

Proof. 1) Similar to the proof in [11].

2) Proof by construction. Since A is not part of the grounded extension G , nor attacked by it, it is attacked by some minimal set of arguments S , such that $S \not\subseteq G$ and G does not attack S . As the system is uncontroversial, it is impossible for any member of S to participate in a minimal attack on S , so the set S is conflict-free. Following a process similar to the one in the proof of Lmm. 9, substituting S for $\{B\}$, we can build a series of conflict-free sets that consists of arguments that indirectly attack A . Extending the union of these sets to a complete extension provides the sought extension. \square

Theorem 4. *Every limited controversial system is coherent, and every uncontroversial system is also relatively grounded.*

Corollary 11. *Every limited controversial argumentation system possesses at least one stable extension.*

This ends our derivation of results mirroring those in [11]. [11] furthermore provides a series of results, showing how some formalisms are special cases of his framework. As Dung's framework itself is a special case of our framework, it follows that these frameworks are also special cases of our framework.

[11] ends with a procedure that turns any finitary argumentation system, as defined in [11], into a logic program, and thereby provides a tractable means for computing grounded extensions of such systems. As our framework is more general, it does not allow for Dung's procedure to be used directly. Instead we provide the following procedure for finitary systems: Given a finitary argumentation system (A, \triangleright) , we define a Prolog encoding of this system as the clauses

$$\{\mathbf{attacks}([S], A) \leftarrow | S \triangleright A \} ,$$

where $[S]$ is a Prolog list declaration containing the arguments in S .

Furthermore, a general interpreter for a Prolog encoding of a finitary argumentation system, is defined as:

$$\begin{aligned} &\{\mathbf{acceptable}(X) \leftarrow \neg \mathbf{defeated}(X); \\ &\quad \mathbf{defeated}(X) \leftarrow \mathbf{attacks}(Y, X), \mathbf{acc}(Y); \\ &\quad \mathbf{acc}(X|Y) \leftarrow \mathbf{acceptable}(X), \mathbf{acc}(Y); \\ &\quad \mathbf{acc}(X) \leftarrow \mathbf{acceptable}(X); \} . \end{aligned}$$

4 Related Work

While not explicitly analyzing the problem of Dung’s framework addressed here, nor trying to generalize it in a conservative manner, some previous works have allowed for sets of attacking arguments, although mainly as side effects. First and foremost, [29, Chapter 5] provided a framework, CumuLA, with a very general attack relation, which allows sets of arguments to attack other sets. However, the framework is focused on modelling the actual dialectic process of argumentation, rather than investigating the essentials of justified and acceptable arguments, and perhaps as a consequence of this, the semantics presented by Verheij is neither as clear as Dung’s nor does it allow for simple comparisons with other formalisms. Furthermore, there are some flaws in Verheij’s treatment, which effectively leave CumuLA with no well-founded semantics. Specifically, three requirements on allowed extensions turn out to prevent seemingly sensible systems from being analysed, and the semantics associated with an attack on sets of arguments is context dependent. For more on these problems see [22].

Later, Verheij has developed two additional frameworks that allow for sets of attacking arguments, namely Argue!, described in [30], and the formal logical framework of DefLog, described in [32] and implemented in [31]. Even though these frameworks builds on ideas from CumuLA, they avoid the problems associated with that framework. However, the two frameworks have other short-comings that make us prefer a conservative generalization of Dung’s framework: Argue! employs only a step-based procedural semantics, and thus lacks the analytical tools, theoretic results, and scope of [11]. DefLog, on the other hand, is well-investigated, but lacks a skeptical semantics, and allows sets of attacking arguments only as a rather contrived encoding. For instance, the attack $\{A, B\} \triangleright C$ would be encoded as

$$A \rightsquigarrow (B \rightsquigarrow \times(C)),$$

where \rightsquigarrow denotes primitive implication, and $\times(\cdot)$ denotes defeat of its argument. There are two two problems with this encoding, one technical and one aesthetic. The first is that systems involving infinite sets of attacking arguments cannot be analysed. The second that the symmetry of the set of attackers is broken. Consider for instance the case where A is “X weighs less than 80 kg”, B is “X is taller than 180 cms”, and C is “X is obese”; here encoding the fact that A and B together defeat C as “X weighs less than 80 kg” implies that “X is taller than 180 cms, so X is not obese” seems to us to be inelegant, and the larger the set of attackers, the larger the inelegance.

The power of encoding sets of attacking arguments wielded by DefLog is due to its expressive language, which is closed under both an implicative operator and an negative operator. Some other argumentation frameworks that are based on formal languages employing similar operators also have implicit methods for encoding attacks by sets of arguments. Most notable is the framework presented in [33], which allows for any sets of sentences to attack each other by encoding rules that from each of them lead to a contradiction. Undercutting attacks are, however, not expressible without assumptions on the underlying language. [8] and [15] present frameworks based on similar ideas. However, all of these fail to abstract from the structure of arguments and as a result do not clearly distinguish between arguments and their interactions, unlike the frameworks

of [11] and this paper. Moreover, the approach of encoding attacks in a logical language restrains sets of attackers to be finite.

5 Conclusions

In this paper we have started exploring formal abstract argumentation systems where synergy can arise between arguments. We believe that we have argued convincingly for the need for such systems, and have examined some of the semantics that can be associated with them. We have tried to do this in the most general fashion possible, by taking outset in the abstract frameworks of [11], and creating a new formalization that allows for sets of arguments to jointly attack other arguments. As we argued in Sect. 2 this degree of freedom ensures that all kinds of attacks between arguments can be modelled faithfully.

Acknowledgments: This work was partially supported by NSF IIS-0329037, and EU PF6-IST 002307 (ASPIC).

References

1. L. Amgoud. *Contribution a l'integration des préférences dans le raisonnement argumentatif*. PhD thesis, Université Paul Sabatier, Toulouse, July 1999.
2. L. Amgoud and C. Cayrol. On the acceptability of arguments in preference-based argumentation framework. In *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence*, pages 1–7, 1998.
3. 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.
4. L. Amgoud, N. Maudet, and S. Parsons. An argumentation-based semantics for agent communication languages. In *Proceedings of the Fifteenth European Conference on Artificial Intelligence*, 2002.
5. S. Benferhat, D. Dubois, and H. Prade. Argumentative inference in uncertain and inconsistent knowledge bases. In *Proceedings of the 9th Conference on Uncertainty in Artificial Intelligence*, pages 411–419, San Mateo, CA, 1993. Morgan Kaufmann.
6. L. Birnbaum. Argument molecules: a functional representation of argument structure. In *Proceedings of the 2nd National Conference on Artificial Intelligence*, pages 63–65, Los Altos, CA, 1982. William Kaufmann.
7. L. Birnbaum, M. Flowers, and R. McGuire. Towards an AI model of argumentation. In *Proceedings of the 1st National Conference on Artificial Intelligence*, pages 313–315, Los Altos, CA, 1980. William Kaufmann.
8. A. Bondarenko, P. M. Dung, R. A. Kowalski, and F. Toni. An abstract argumentation-theoretic approach to default reasoning. *Artificial Intelligence*, 93:63–101, 1997.
9. D. V. Carbogim, D. Robertson, and J. Lee. Argument-based applications to knowledge engineering. *The Knowledge Engineering Review*, 15(2), 2000.
10. C. Cayrol. On the relation between argumentation and non-monotonic coherence-based entailment. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pages 1443–1448, San Mateo, CA, 1995. Morgan Kaufmann.

11. P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–358, 1995.
12. M. Elvang-Gøransson and A. Hunter. Argumentative logics: reasoning with classically inconsistent information. *Data and Knowledge Engineering*, 16:125–145, 1995.
13. M. Elvang-Gøransson, P. Krause, and J. Fox. Dialectic reasoning with inconsistent information. In *Proceedings of the 9th Conference on Uncertainty in Artificial Intelligence*, pages 114–121, San Mateo, CA, 1993. Morgan Kaufmann.
14. M. Flowers, R. McGuire, and L. Birnbaum. Adversary arguments and the logic of personal attacks. In W. G. Lehnert and M. H. Ringle, editors, *Strategies for natural language processing*, pages 275–294. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1982.
15. A. J. Garcia and G. R. Simari. Defeasible logic programming an argumentative approach. *Theory and Practice of Logic Programming*, 4(1):95–138, 2004.
16. J. Kohlas. Symbolic evidence, arguments, supports and valuation networks. In M. Clarke, R. Kruse, and S. Moral, editors, *Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, pages 186–198. Springer Verlag, Berlin, Germany, 1993.
17. P. Krause, S. Ambler, M. Elvang-Gøransson, and J. Fox. A logic of argumentation for reasoning under uncertainty. *Computational Intelligence*, 11:113–131, 1995.
18. F. Lin. An argument-based approach to non-monotonic reasoning. *Computational Intelligence*, 9:254–267, 1993.
19. R. Loui. Defeat among arguments: a system of defeasible inference. *Computational Intelligence*, 3:100–106, 1987.
20. P. McBurney. *Rational Interaction*. PhD thesis, Department of Computer Science, University of Liverpool, 2002.
21. R. McGuire, L. Birnbaum, and M. Flowers. Opportunistic processing in arguments. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, pages 58–60, Menlo Park, CA, 1981. American Association for Artificial Intelligence.
22. S. H. Nielsen and S. Parsons. Note on the short-comings of CumulA. <http://www.cs.aau.dk/~holbech/cumulanote.ps>.
23. S. Parsons, C. Sierra, and N. R. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, 1998.
24. S. Parsons, M. Wooldridge, and L. Amgoud. Properties and complexity of formal inter-agent dialogues. *Journal of Logic and Computation*, 13(3):347–376, 2003.
25. John L. Pollock. *Cognitive Carpentry: A Blueprint for How to Build a Person*. MIT Press, 1995.
26. H. Prakken. Relating protocols for dynamic dispute with logics for defeasible argumentation. *Synthese*, 127:187–219, 2001.
27. H. Prakken. A study of accrual of arguments, with applications to evidential reasoning. In *Proceedings of the Tenth International Conference on Artificial Intelligence and Law*, pages 85–94. ACM, 2005.
28. H. Prakken and G. Vreeswijk. Logics for defeasible argumentation. In D. Gabbay, editor, *Handbook of Philosophical Logic*. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2000.
29. B. Verheij. *Rules, Reasons, Arguments. Formal studies of argumentation and defeat*. PhD thesis, Universiteit Maastricht, 1996.
30. B. Verheij. Argue! - an implemented system for computer-mediated defeasible argumentation. In H. La Poutre and H.J. van den Herik, editors, *Proceedings of the Tenth Netherlands/Belgium Conference on Artificial Intelligence*, pages 57–66. CWI, Amsterdam, 1998.
31. B. Verheij. Artificial argument assistants for defeasible argumentation. *Artificial Intelligence*, 150(1–2):291–324, 2003.

32. B. Verheij. Deflog: on the logical interpretation of prima facie justified assumptions. *Journal of Logic and Computation*, 13(3):319–346, 2003.
33. G. A. W. Vreeswijk. Abstract argumentation systems. *Artificial Intelligence*, 90(1):225–279, 1997.