

Abstract Solvers for Dung’s Argumentation Frameworks

Remi Brochenin¹, Thomas Linsbichler², Marco Maratea¹, Johannes Peter Wallner³, and Stefan Woltran²

¹ University of Genova, Italy

² TU Wien, Austria

³ HIIT, Department of Computer Science, University of Helsinki, Finland

Abstract. Abstract solvers are a quite recent method to uniformly describe algorithms in a rigorous formal way and have proven successful in declarative paradigms such as Propositional Satisfiability and Answer-Set Programming. In this paper, we apply this machinery for the first time to a dedicated AI formalism, namely Dung’s abstract argumentation frameworks. We provide descriptions of several advanced algorithms for the preferred semantics in terms of abstract solvers and, moreover, show how slight adaptations thereof directly lead to new algorithms.

1 Introduction

Dung’s concept of abstract argumentation [12] is nowadays a core formalism in AI [2, 21]. The problem of solving certain reasoning tasks on such frameworks is the centerpiece of many advanced higher-level argumentation systems. The problems to be solved are however intractable and might even be hard for the second level of the polynomial hierarchy [13, 15]. Thus, efficient and advanced algorithms have to be developed in order to deal with real-world size data with reasonable performance. The argumentation community is currently facing this challenge [7] and a first solver competition⁴ is organized this year. Thus, a number of new algorithms and systems will be developed in the near future. Being able to precisely analyze and compare already developed and new algorithms is a fundamental step in order to understand the ideas behind such high-performance systems, and to build a new generation of more efficient algorithms and solvers.

Usually, algorithms are presented by means of pseudo-code descriptions, but other communities have experienced that analyzing such algorithms on this basis may not be fruitful. More formal descriptions, which allow, e.g. for a uniform representation, have thus been developed: a recent and successful approach in this direction is the concept of *abstract solvers* [19]. Hereby, one characterizes the states of computation as nodes of a graph, the techniques as arcs between nodes, and the whole solving process as a path in the graph. This concept not only proved successful for SAT [19], but also has been applied for several variants of Answer-Set Programming [4, 16, 17].

⁴ <http://argumentationcompetition.org>

In this paper, we make a first step to investigate the appropriateness of abstract solvers for dedicated AI formalisms and focus on certain problems in Dung’s argumentation frameworks. In order to understand whether abstract solvers are powerful enough, we consider quite advanced algorithms – ranging from dedicated [20] to reduction-based [5, 14] approaches (see [8] for a recent survey) – for solving problems that are hard for the second level of the polynomial hierarchy. We show that abstract solvers allow for convenient algorithm design resulting in a clear and mathematically precise description, and how formal properties of the algorithms are easily specified by means of related graph properties. We also illustrate how abstract solvers simplify the *combination* of techniques implemented in different solvers in order to define new solving procedures. Consequently, our findings not only prove that abstract solvers are a valuable tool for specifying and analysing argumentation algorithms, but also indicate the broad range the novel concept of abstract solvers can be applied to. To sum up, our main contributions are as follows:

- We provide a full formal description of recent algorithms [5, 14, 20] for reasoning tasks under the preferred semantics in terms of abstract solvers, thus enabling a comparison of these approaches at a formal level.
- We exemplify one proof illustrating how formal correctness of algorithms can be shown with the help of descriptions in terms of abstract solvers.
- We outline how our reformulations can be used to gain more insight into the algorithms and how novel combinations of “levels” of abstract solvers might pave the way for new solutions.

The paper is structured as follows. Section 2 introduces the required preliminaries about abstract argumentation frameworks and abstract solvers. Then, Section 3 shows how our target algorithms are reformulated in terms of abstract solvers and introduces a new solving algorithm obtained from combining the target algorithms. The paper ends in Section 4 with final remarks and possible topics for future research.

2 Preliminaries

In this section we first review (abstract) argumentation frameworks [12] and their semantics (see [1] for an overview), and then introduce abstract transition systems [19] on the concrete instance describing the DPLL-procedure [9].

Abstract Argumentation Frameworks. An *argumentation framework (AF)* is a pair $F = (A, R)$ where A is a finite set of arguments and $R \subseteq A \times A$ is the *attack relation*. Semantics for argumentation frameworks assign to each AF $F = (A, R)$ a set $\sigma(F) \subseteq 2^A$ of *extensions*. We consider here for σ the functions *adm*, *com*, and *prf*, which stand for admissible, complete, and preferred semantics. Towards the definitions of the semantics we need some formal concepts. For an AF $F = (A, R)$, an argument $a \in A$ is *defended (in F)* by a set $S \subseteq A$ if for each $b \in A$ such that $(b, a) \in R$, there is a $c \in S$, such that $(c, b) \in R$ holds.

Definition 1. Let $F = (A, R)$ be an AF. A set $S \subseteq A$ is conflict-free (in F), denoted $S \in cf(F)$, if there are no $a, b \in S$ such that $(a, b) \in R$. For $S \in cf(F)$, it holds that

- $S \in adm(F)$ if each $a \in S$ is defended by S ;
- $S \in com(F)$ if $S \in adm(F)$ and for each $a \in A$ defended by S , $a \in S$ holds; and
- $S \in prf(F)$ if $S \in adm(F)$ (resp. $S \in com(F)$) and there is no $T \in adm(F)$ (resp. $T \in com(F)$) with $T \supset S$.

Given an AF $F = (A, R)$, an argument $a \in A$, and a semantics σ , the problem of skeptical acceptance asks whether it is the case that a is contained in all σ -extensions of F . While skeptical acceptance is trivial for adm and decidable in polynomial time for com , it is Π_2^P -complete for prf , see [10, 12, 13]. The class $\Pi_2^P = \text{coNP}^{\text{NP}}$ denotes the class of problems P , such that the complementary problem \bar{P} can be decided by a nondeterministic polynomial time algorithm that has (unrestricted) access to an NP-oracle.

Abstract Solvers. Most SAT solvers are based on the Davis-Putnam-Logemann-Loveland (DPLL) procedure [9]. We give an abstract transition system for DPLL following the work of Nieuwenhuis et al. in [19]. We start with basic notation for Boolean logic.

For a Conjunctive Normal Form (CNF) formula φ (resp. a set of literals M), we denote the set of atoms occurring in φ (resp. in M) by $atoms(\varphi)$ (resp. $atoms(M)$). We identify a consistent set E of literals (i.e. a set that does not contain complementary literals, for example a and $\neg a$) with an assignment to $atoms(E)$ as follows: if $a \in E$ then a maps to *true*, while if $\neg a \in E$ then a maps to *false*. For sets X and Y of atoms such that $X \subseteq Y$, we identify X with an assignment over Y as follows: if $a \in X$ then a maps to *true*, while if $a \in Y \setminus X$ then a maps to *false*. By $\text{Sat}(\varphi)$ we refer to the set of satisfying assignments of φ .

We now introduce the abstract procedure for deciding whether a CNF formula is satisfiable. A *decision literal* is a literal annotated by d , as in l^d . An *annotated literal* is a literal, a decision literal or the false constant \perp . For a set X of atoms, a *record* relative to X is a string E composed of annotated literals over X without repetitions. For instance, \emptyset , $\neg a^d$ and $a \neg a^d$ are records relative to the set $\{a\}$. We say that a record E is *inconsistent* if it contains \perp or both a literal l and its complement \bar{l} , and consistent otherwise. We sometimes identify a record with the set containing all its elements without annotations. For example, we identify the consistent record $b^d \neg a$ with the consistent set $\{\neg a, b\}$ of literals, and so with the assignment which maps a to *false* and b to *true*.

Each CNF formula φ determines its DPLL *graph* DP_φ . The set of nodes of DP_φ consists of the records relative to the set of atoms occurring in φ and the distinguished states *Accept* and *Reject*. A node in the graph is *terminal* if no edge originates from it; in practice, the terminal nodes are *Accept* and *Reject*. The edges of the graph DP_φ are specified by the transition rules presented in

Oracle rules			
<i>Backtrack</i>	$E l^d E'' \Rightarrow E \bar{l}$	if	$\left\{ \begin{array}{l} E l^d E'' \text{ is inconsistent and} \\ E'' \text{ contains no decision literal} \end{array} \right.$
<i>UnitPropagate</i>	$E \Rightarrow E l$	if	$\left\{ \begin{array}{l} l \text{ does not occur in } E \text{ and} \\ \text{all the literals of } \bar{C} \text{ occur in } E \text{ and} \\ C \vee l \text{ is a clause in } \varphi \end{array} \right.$
<i>Decide</i>	$E \Rightarrow E l^d$	if	$\left\{ \begin{array}{l} E \text{ is consistent and} \\ \text{neither } l \text{ nor } \bar{l} \text{ occur in } E \end{array} \right.$
Failing rule			
<i>Fail</i>	$E \Rightarrow \textit{Reject}$	if	$\{ E \text{ is inconsistent and decision-free}$
Succeeding rule			
<i>Succeed</i>	$E \Rightarrow \textit{Accept}$	if	$\{ \text{no other rule applies}$

Fig. 1. The transition rules of DP_φ .

Figure 1. In solvers, generally the oracle rules are chosen with the preference order following the order in which they are stated in Figure 1, while the failing rule has a higher priority than all the oracle rules.

Intuitively, every state of the DPLL graph represents some hypothetical state of the DPLL computation whereas a path in the graph is a description of a process of search for a satisfying assignment of a given CNF formula. The rule *Decide* asserts that we make an arbitrary decision to add a literal or, in other words, to assign a value to an atom. Since this decision is arbitrary, we are allowed to backtrack at a later point. The rule *UnitPropagate* asserts that we can add a literal that is a logical consequence of our previous decisions and the given formula. The rule *Backtrack* asserts that the present state of computation is failing but can be fixed: at some point in the past we added a decision literal whose value we can now reverse. The rule *Fail* asserts that the current state of computation has failed and cannot be fixed. The rule *Succeed* asserts that the current state of computation corresponds to a successful outcome.

To decide the satisfiability of a CNF formula it is enough to find a path in DP_φ leading from node \emptyset to a terminal node. If it is *Accept*, then the formula is satisfiable, and if it is *Reject*, then it is unsatisfiable. Since there is no infinite path, a terminal node is always reached.

3 Algorithms for Preferred Semantics

In this section we abstract two SAT-based algorithms for preferred semantics, namely PrefSat [5] (implemented in the tool ARGSEMSAT [6]) for extension enumeration, and an algorithm for deciding skeptical acceptance of CEGARTIX [14]. Moreover, we abstract the dedicated approach for enumeration of [20]. In Section 3.4 we show how our graph representations can be used to develop novel algorithms, by combining parts of PrefSat and parts of the dedicated algorithm.

We will present these algorithms in a uniform way, abstracting from some minor tool-specific details. Moreover, even if abstract solvers are mainly conceived as a modeling formalism, in our solutions a certain level of systematicity can be outlined, that helps in the design of such abstract solvers. In fact, common to all algorithms is a conceptual two-level architecture of computation, similar to Answer Set Programming solvers for disjunctive logic programs [4]. The lower level corresponds to a DPLL-like search subprocedure, while the higher level part takes care of the control flow and drives the overall algorithm. For PrefSat and CEGARTIX, the subprocedures actually are delegated to a SAT solver, while the dedicated approach carries out a tailored search procedure.

Each algorithm uses its own data structures, and, by slight abuse of notation, for a given AF $F = (A, R)$ we denote their used variables in our graph representation by $atoms(F)$. For this set it holds that $A \subseteq atoms(F)$, i.e. the status of the arguments can be identified from this set of atoms. The states of our graph representations of all algorithms are either

1. an annotated triple $(\epsilon, E', E)_i$ where $i \in \{out, base, max\}$, $\epsilon \subseteq 2^A$ is a set of sets of arguments, and both E' and E are records over $atoms(F)$; or
2. $Ok(\epsilon)$ for $\epsilon \subseteq 2^A$; or
3. a distinguished state *Accept* or *Reject*.

The intended meaning of a state $(\epsilon, E', E)_i$ is that ϵ is the set of already found preferred extensions of F (visited part of the search space), E' is a record representing the current candidate extension (which is admissible or complete in F), and E is a record that may be currently modified by a subprocedure. Note that both E and E' are records, since they will be modified by subprocedures, while found preferred extensions will be translated to a set of arguments before being stored in ϵ . The annotation i denotes the current (sub)procedure we are in. Both *base* and *max* correspond to different lower level computations, typically SAT calls, while *out* is used solely for (simple) checks outside such subprocedures. Transition rules reflecting the higher level of computation shift these annotations, e.g. from a terminated subprocedure *base* to subprocedure *max*, and transition rules mirroring rules “inside” a SAT solver do not modify i .

The remaining states denote terminated computation: $Ok(\epsilon)$ contains all solutions, while *Accept* or *Reject* denote an answer to a decision problem.

The SAT-based algorithms construct formulas by an oracle function f s.t. $A \subseteq atoms(f(\epsilon, E, F, \alpha)) \subseteq atoms(F)$ for all possible arguments of f , in particular for $\alpha \in A$. The formulas $f(\epsilon, E, F, \alpha)$ are adapted from [3]. The argument α is relevant only for CEGARTIX to decide skeptical acceptance of α . Finally, we use $e(M) = M \cap A$ to project the arguments from a record M .

3.1 SAT-based Algorithm for Enumeration

PrefSat (Algorithm 1 of [5]) is a SAT-based algorithm for finding all preferred extensions of a given AF F . The algorithm maintains a list of visited preferred extensions. It first searches for a complete extension not contained in previously

<i>i</i> -oracle rules ($i \in \{base, max\}$)	
$Backtrack_i (\epsilon, E', El^d E'')_i \Rightarrow (\epsilon, E', E\bar{l})_i$	if $\left\{ \begin{array}{l} El^d E'' \text{ is inconsistent and} \\ E'' \text{ contains no decision literal} \end{array} \right.$
$UnitPropagate_i (\epsilon, E', E)_i \Rightarrow (\epsilon, E', El)_i$	if $\left\{ \begin{array}{l} l \text{ does not occur in } E \text{ and} \\ \text{all the literals of } \bar{C} \text{ occur in } E \text{ and} \\ C \vee l \text{ is a clause in } f_i^{com}(\epsilon, E', F, \alpha) \end{array} \right.$
$Decide_i (\epsilon, E', E)_i \Rightarrow (\epsilon, E', El^d)_i$	if $\left\{ \begin{array}{l} E \text{ is consistent and} \\ \text{neither } l \text{ nor } \bar{l} \text{ occur in } E \end{array} \right.$
Succeeding rules	
$Succeed_{base} (\epsilon, E', E)_{base} \Rightarrow (\epsilon, E, \emptyset)_{max}$	if $\left\{ \begin{array}{l} \text{no other rule applies} \end{array} \right.$
$Succeed_{max} (\epsilon, E', E)_{max} \Rightarrow (\epsilon, E, \emptyset)_{max}$	if $\left\{ \begin{array}{l} \text{no other rule applies} \end{array} \right.$
Failing rules	
$Fail_{base} (\epsilon, E', E)_{base} \Rightarrow Ok(\epsilon)$	if $\left\{ \begin{array}{l} E \text{ is inconsistent and decision-free} \end{array} \right.$
$Fail_{max} (\epsilon, E', E)_{max} \Rightarrow (\epsilon \cup \{e(E')\}, \emptyset, \emptyset)_{base}$	if $\left\{ \begin{array}{l} E \text{ is inconsistent and decision-free} \end{array} \right.$

Fig. 2. The rules of $ENUM_{\bar{f}}^F$.

found preferred extensions. If such an extension is found, it is iteratively extended until we reach a subset-maximal complete extension, i.e. a preferred extension. This preferred extension is stored, and we repeat the process.

In PrefSat we have two subprocedures that are delegated to a SAT solver. The first has to generate a complete extension not contained in one of the enumerated preferred extensions, and the second searches for a complete extension that is a strict superset of a given one.

We now represents PrefSat via abstract solver. The graph $ENUM_{\bar{f}}^F$ for an AF $F = (A, R)$ and a vector of oracle functions \bar{f} is defined by the states over $atoms(F)$ and the transition rules presented in Figure 2. Its initial state is $(\emptyset, \emptyset, \emptyset)_{base}$. We assume the functions f_{base}^{com} and f_{max}^{com} that generate CNF formulas for $\epsilon \subseteq 2^A$, a record E , and an argument $\alpha \in A$ such that:

1. $\{e(M) \mid M \in \text{Sat}(f_{base}^{com}(\epsilon, E, F, \alpha))\} = \{E'' \in \text{com}(F) \mid \neg \exists E' \in \epsilon : E'' \subseteq E'\}$;
2. $\{e(M) \mid M \in \text{Sat}(f_{max}^{com}(\epsilon, E, F, \alpha))\} = \{E'' \in \text{com}(F) \mid e(E) \subset E''\}$.

We remark that α is not relevant for enumeration of extensions and only used for acceptance later on. In a state $(\epsilon, E', E)_i$, the set ϵ represents preferred extensions found as of now, E' is a record for the complete extension found in the previous subprocedure, and E is a record for the complete extension that the current oracle is trying to build. The annotation $i \in \{base, max\}$ corresponds to different kinds of SAT calls.

If the conditions of a rule with annotation i check for consistency, we implicitly refer to the formula generated by f_i^{com} . That is, if a $Fail_i$ rule is applied to the state $(\epsilon, E', E)_i$ for $i \in \{base, max\}$, the formula $f_i^{com}(\epsilon, E', F, \alpha)$ is unsatisfi-




Fig. 3. AF F with $\text{prf}(F) = \{\{a, c\}, \{a, d\}\}$.

able. Conversely, when a Succeed_i rule is applied, the formula $f_i^{\text{com}}(\epsilon, E', F, \alpha)$ is satisfied by E . Notice that Fail_i and Succeed_i might shift i to reflect a change of type of SAT calls. When $i = \text{base}$, the oracle searches for a complete extension that has not been found before. In case of failure all the preferred extensions have been found. In case of success, it is necessary to search whether there are strictly larger complete extensions than the one found. This is handled by sub-procedure max . In case of success, $\text{Succeed}_{\text{max}}$ is applied and the procedure is repeated, since the current complete extension might still not be maximal. Failure by Fail_{max} means we have found a preferred extension.

Example 1. Consider the AF F depicted in Figure 3, where nodes of the graph represent arguments and edges represent attacks. F has two preferred extensions, namely $\{a, c\}$ and $\{a, d\}$. Figure 4 shows a possible path in the graph ENUM_F^F . As expected, the computation terminates in the state $\text{Ok}(\{\{a, d\}, \{a, c\}\})$. Note that we abbreviate the parts of the path where we are “inside” the SAT-solver. Also, we only show literals over A , and do not state the extra literals that may have been assigned during the call to the SAT-solver. By unsat we represent an inconsistent and decision-free record.

Initial state : $(\emptyset, \emptyset, \emptyset)_{\text{base}}$
 base-oracle : $(\emptyset, \emptyset, E_1 \supseteq \{a, \neg b, \neg c, \neg d\})_{\text{base}}$
 Succeed_{base} : $(\emptyset, E_1, \emptyset)_{\text{max}}$
 max-oracle : $(\emptyset, E_1, E_2 \supseteq \{a, \neg b, \neg c, d\})_{\text{max}}$
 Succeed_{max} : $(\emptyset, E_2, \emptyset)_{\text{max}}$
 max-oracle : $(\emptyset, E_2, \text{unsat})_{\text{max}}$
 Fail_{max} : $(\{\{a, d\}\}, \emptyset, \emptyset)_{\text{base}}$
 base-oracle : $(\{\{a, d\}\}, \emptyset, E_3 \supseteq \{a, \neg b, c, \neg d\})_{\text{base}}$
 Succeed_{base} : $(\{\{a, d\}\}, E_3, \emptyset)_{\text{max}}$
 max-oracle : $(\{\{a, d\}\}, E_3, \text{unsat})_{\text{max}}$
 Fail_{max} : $(\{\{a, d\}, \{a, c\}\}, \emptyset, \emptyset)_{\text{base}}$
 base-oracle : $(\{\{a, d\}, \{a, c\}\}, \emptyset, \text{unsat})_{\text{base}}$
 Fail_{base} : $\text{Ok}(\{\{a, d\}, \{a, c\}\})$

Fig. 4. Path in ENUM_F^F where F is the AF from Figure 3.

It remains to show that ENUM_F^F correctly describes PrefSat by showing that we reach a terminal state containing all preferred extensions of F . We begin with

a lemma stating that we only add preferred extensions to ϵ which have not been found at this point.

Lemma 1. *For any AF $F = (A, R)$, if the rule $Fail_{max}$ is applied from state $(\epsilon, E', E)_{max}$ in the graph $ENUM_F^F$ then $e(E') \in prf(F)$ and $e(E') \notin \epsilon$.*

Proof. Let $S_1 = (\epsilon_1, E'_1, E_1)_{max}$ be the state from which $Fail_{max}$ is applied. This means that f_{max}^{com} is unsatisfiable, hence, by the definition of formula f_{max}^{com} , there is no $C \in com(F)$ with $C \supset e(E'_1)$. To get $e(E'_1) \in prf(F)$ it remains to show that $e(E'_1) \in com(F)$. Observe that $Succeed_{base}$ is applied at least once, since every framework has a complete extension. Moreover, an update of the value of E'_1 is only done by an application of $Succeed_{base}$ or $Succeed_{max}$. In both cases $e(E'_1)$ corresponds to a complete extension of F , since E'_1 is a satisfying assignment of the formula f_{base}^{com} or f_{max}^{com} , respectively. Therefore E'_1 is a complete extension of F .

Since the initial state is $(\emptyset, \emptyset, \emptyset)_{base}$, an application of $Succeed_{base}$ must precede $Fail_{max}$. From this application of $Succeed_{base}$ it follows that there is a record E' such that $\neg \exists C \in \epsilon : e(E') \subseteq C$ follows. Moreover every application of $Succeed_{max}$ updates E' by a proper superset of itself. Therefore $e(E'_1) \supseteq e(E')$ and also $\neg \exists C \in \epsilon : e(E'_1) \subseteq C$, in particular $e(E'_1) \notin \epsilon$. \square

Now we are ready to show correctness of $ENUM_F^F$.

Theorem 1. *For any AF F , the graph $ENUM_{(f_{base}^{com}, f_{max}^{com})}^F$ is finite, acyclic and the only terminal state reachable from the initial state is $Ok(\epsilon)$ where $\epsilon = prf(F)$.*

Proof. In order to show that $ENUM_F^F$ is finite, consider some state $(\epsilon, E', E)_i$ of $ENUM_F^F$. Since both E and E' are records over $atoms(F)$, and F is finite by definition, the number of possible records E and E' is finite. Similarly, there is only a finite number of sets of sets of arguments ϵ . Finally, $ENUM_F^F$ only contains states with $i \in \{base, max\}$. Thus the number of states is finite in the graph $ENUM_F^F$.

In order to show that it is acyclic consider two states $S_1 = (\epsilon_1, E'_1, E_1)_{i_1}$ and $S_2 = (\epsilon_2, E'_2, E_2)_{i_2}$. For a record E let $s(E) = |L^0|, |L^1|, \dots, |L^p|$, where E is of the form $L^0 l^1 L^1 \dots l^p L^p$ with l^1, \dots, l^p being the decision literals of E . We define the strict partial order on states such that $S_1 < S_2$ iff

- (i) $\epsilon_1 \subset \epsilon_2$, or
- (ii) $\epsilon_1 = \epsilon_2$ and $i_1 <_i i_2$, or
- (iii) $\epsilon_1 = \epsilon_2$ and $i_1 = i_2$ and $e(E'_1) \subset e(E'_2)$, or
- (iv) $\epsilon_1 = \epsilon_2$ and $i_1 = i_2$ and $e(E'_1) = e(E'_2)$ and $E_1 <_l E_2$,

where $base <_i max$ and $E_1 <_l E_2$ iff $s(E_1)$ is lexicographically smaller than $s(E_2)$. We show that each transition rule is increasing with respect to $<$: First of all, the i -oracle rules (i.e. $Backtrack_i$, $UnitPropagate_i$, and $Decide_i$) fulfill $S_1 < S_2$ because of (iv). For all of these rules $\epsilon_1 = \epsilon_2$, $E'_1 = E'_2$ and $i_1 = i_2$, but $s(E_1)$ is lexicographically smaller than $s(E_2)$, therefore $E_1 <_l E_2$. Moreover, $Fail_{max}$

fulfills $S_1 < S_2$ due to (i) since $e(E'_1) \notin \epsilon_1$ by Lemma 1. $Succeed_{base}$ guarantees $S_1 < S_2$ because of (ii). Finally, $Succeed_{max}$ fulfills $S_1 < S_2$ due to (iii), since the max -oracle rules work on the formula f_{max}^{com} and the extension associated with a satisfying assignment $E_1 = E'_2$ thereof must be a proper superset of $e(E'_1)$. We have shown that each transition rule is increasing with respect to $<$. Therefore, for any two states S_1 and S_n such that S_n is reachable from S_1 in $ENUM_{\bar{f}}^F$ it holds that $S_1 < S_n$, showing that the graph is acyclic.

The only terminal state reachable from the initial state is $Ok(\epsilon)$ (via rule $Fail_{base}$) since all states $S = (\epsilon, E, E')_i$ of $ENUM_{\bar{f}}^F$ have $i \in \{base, max\}$ and for each $i \in \{base, max\}$ there is a rule $Succeed_i$ with the condition “no other rule applies”. It remains to show that, when state $Ok(\epsilon)$ is reached, ϵ coincides with $prf(F)$. Since elements are only added to ϵ by application of the rule $Fail_{max}$ we know from Lemma 1 that for each $P \in \epsilon$ it holds that $P \in prf(F)$. To reach $Ok(\epsilon)$, the rule $Fail_{base}$ must have been applied from state a $(\epsilon, E', E)_{base}$. This means, by the definition of f_{base}^{com} , that for each complete extension C of F there is some $P \in \epsilon$ such that $C \subseteq P$. Hence $\epsilon = prf(F)$. \square

3.2 SAT-based Algorithm for Acceptance

CEGARTIX [14] is a SAT-based tool for deciding several acceptance questions for AFs. We focus here on Algorithm 1 of [14] for deciding skeptical acceptance under preferred semantics of an argument $\alpha \in A$. Similarly as PrefSat, CEGARTIX traverses the search space of a certain semantics, generates candidate extensions not contained in already visited preferred extensions, and maximizes the candidate until a preferred extension is found. The main differences to PrefSat are (1) the parametrized usage of the base semantics σ (the search space), which can be either admissible or complete semantics, and (2) the incorporation of the queried argument α . To prune the search space, it is required that α is contained in the candidate σ -extension before maximization. Again, we have two kinds of SAT-calls.

The graph $SKEPT-PRF_{\bar{f}}^{F,\alpha}$ for an AF F , an argument α and a vector of oracle functions \bar{f} is defined by the states over $atoms(F)$ and the rules in Figure 2 replacing the $Fail_i$ rules and adding the *out* rules as depicted in Figure 5. The initial state is $(\emptyset, \emptyset, \emptyset)_{base}$. For $\sigma \in \{adm, com\}$ we assume the functions f_{base}^σ and f_{max}^σ such that:

1. $\{e(M) \mid M \in \text{Sat}(f_{base}^\sigma(\epsilon, E, F, \alpha))\} = \{E'' \in \sigma(F) \mid \alpha \notin E'' \wedge \neg \exists E' \in \epsilon : E'' \subseteq E'\}$;
2. $\{e(M) \mid M \in \text{Sat}(f_{max}^\sigma(\epsilon, E, F, \alpha))\} = \{E'' \in \sigma(F) \mid e(E) \subset E''\}$.

The graph $SKEPT-PRF_{\bar{f}}^{F,\alpha}$ is nearly identical to $ENUM_{\bar{f}}^F$. It differs only in case of failure in subprocedure *base* or *max*. When all the preferred extensions have been enumerated in subprocedure *base*, we can report a positive outcome with *Accept*, since we have ensured that α belongs to all of them. In subprocedure *max*, when a preferred extension has been found, it is here necessary to check whether α belongs to it. The *out* rules correspond to an if-then-else construct:

Failing rules

$$\begin{array}{ll}
Fail_{base} & (\epsilon, E', E)_{base} \Rightarrow Accept & \text{if } \{ E \text{ is inconsistent and decision-free} \\
Fail_{max} & (\epsilon, E', E)_{max} \Rightarrow (\epsilon, E', \emptyset)_{out} & \text{if } \{ E \text{ is inconsistent and decision-free} \\
Fail_{out} & (\epsilon, E', E)_{out} \Rightarrow (\epsilon \cup \{e(E')\}, \emptyset, \emptyset)_{base} & \text{if } \{ \alpha \in e(E')
\end{array}$$

Succeeding rules

$$Succeed_{out} (\epsilon, E', E)_{out} \Rightarrow Reject \quad \text{if } \{ \alpha \notin e(E')$$

Fig. 5. Changed transition rules for $SKEPT-PRF_{\bar{f}}^{F,\alpha}$.

if the condition $\alpha \notin E'$ holds then we follow the $Succeed_{out}$ rule else follow the $Fail_{out}$ rule. In other words, if α is not in the extension then the procedure can terminate with a negative answer; else proceed as in the previous graph: add the preferred extension to ϵ and search for a new one by going back to *base*.

Example 2. Again consider the AF F from Figure 3 and note that skeptical acceptance of the argument c is rejected since c is not contained in the preferred extension $\{a, d\}$ of F . Accordingly, the possible path of the graph $SKEPT-PRF_{\bar{f}}^{F,c}$ which is depicted in Figure 6 (with base semantics *adm*) terminates in the *Reject*-state.

Initial state : $(\emptyset, \emptyset, \emptyset)_{base}$
base-oracle : $(\emptyset, \emptyset, E_1 \supseteq \{a, \neg b, \neg c, \neg d\})_{base}$
Succeed_{base} : $(\emptyset, E_1, \emptyset)_{max}$
max-oracle : $(\emptyset, E_1, E_2 \supseteq \{a, \neg b, c, \neg d\})_{max}$
Succeed_{max} : $(\emptyset, E_2, \emptyset)_{max}$
max-oracle : $(\emptyset, E_2, unsat)_{max}$
Fail_{max} : $(\emptyset, E_2, \emptyset)_{out}$
Fail_{out} : $(\{\{a, c\}\}, \emptyset, \emptyset)_{base}$
base-oracle : $(\{\{a, c\}\}, \emptyset, E_3 \supseteq \{a, \neg b, \neg c, d\})_{base}$
Succeed_{base} : $(\{\{a, c\}\}, E_3, \emptyset)_{max}$
max-oracle : $(\{\{a, c\}\}, E_3, unsat)_{max}$
Fail_{max} : $(\{\{a, c\}\}, E_3, \emptyset)_{out}$
Succeed_{out} : *Reject*

Fig. 6. Reject-path for argument c in $SKEPT-PRF_{\bar{f}}^{F,c}$.

On the other hand, argument a is skeptically accepted under preferred semantics in F as it belongs to all preferred extensions enumerated in $\{\{a, d\}, \{a, c\}\}$. For checking whether a is skeptically accepted in F , a possible path in the graph $SKEPT-PRF_{\bar{f}}^{F,a}$ (again with base semantics *adm*) is shown in Figure 7. As expected, the path terminates in the state *Accept*.

Initial state : $(\emptyset, \emptyset, \emptyset)_{base}$
 base-oracle : $(\emptyset, \emptyset, E_1 \supseteq \{\neg a, \neg b, \neg c, \neg d\})_{base}$
 Succeed_{base} : $(\emptyset, E_1, \emptyset)_{max}$
 max-oracle : $(\emptyset, E_1, E_2 \supseteq \{a, \neg b, \neg c, \neg d\})_{max}$
 Succeed_{max} : $(\emptyset, E_2, \emptyset)_{max}$
 max-oracle : $(\emptyset, E_2, E_3 \supseteq \{a, \neg b, \neg c, d\})_{max}$
 Succeed_{max} : $(\emptyset, E_3, \emptyset)_{max}$
 max-oracle : $(\emptyset, E_3, unsat)_{max}$
 Fail_{max} : $(\emptyset, E_3, \emptyset)_{out}$
 Fail_{out} : $(\{\{a, d\}\}, \emptyset, \emptyset)_{base}$
 base-oracle : $(\{\{a, d\}\}, \emptyset, unsat)_{base}$
 Fail_{base} : *Accept*

Fig. 7. Accept-path for argument a in $\text{SKEPT-PRF}_{\mathcal{F}}^{F, \alpha}$.

Theorem 2. *For any AF $F = (A, R)$, argument $\alpha \in A$, and $\sigma \in \{adm, com\}$, the graph $\text{SKEPT-PRF}_{(f_{base}^\sigma, f_{max}^\sigma)}^{F, \alpha}$ is finite, acyclic and any terminal state reachable from the initial state is either *Accept* or *Reject*; *Accept* is reachable iff α is skeptically accepted in F w.r.t. *prf*.*

3.3 Dedicated Approach for Enumeration

Algorithm 1 of [20] presents a direct approach for enumerating preferred extensions. One function is important for this algorithm, which is called IN-TRANS. It marks an argument $x \in A$ as belonging to the currently built extension, and marks all attackers $\{y \mid (y, x) \in R\}$ and all attacked arguments $\{y \mid (x, y) \in R\}$ as outside of this extension. Intuitively, IN-TRANS *decides* to accept x , and then *propagates* the immediate consequences to the neighboring nodes. It actually does an additional task. It labels the attacked arguments as “attacked”, and the attackers that are not yet labelled as attacked as “to be attacked”: this allows later to easily check the admissibility of the extension by just looking whether there is any argument “to be attacked”.

The algorithm is recursive, and stores the admissible extensions in a global variable. First, it checks whether all the arguments are marked as either belonging or outside the extension, and if so it returns after adding the extension to the global variable if the extension is actually admissible. Second, it applies the function IN-TRANS to some unmarked argument and calls itself recursively. Third, it reverts the effects of IN-TRANS, marks the argument it chose as outside of this extension, and calls itself recursively. This can be seen as a *backtrack*.

We have defined an equivalent representation of this algorithm that follows the framework of abstract solvers with binary logics as previously used in this article. Binary truth values are sufficient to represent the arguments marked, but we see the labels “attacked” and “to be attacked” as an optimization as they can be easily recovered at the end of the algorithm. Indeed, they correspond to the

Oracle rules

$$\begin{array}{ll}
\text{Backtrack}'_{max} & (\epsilon, \emptyset, Ea^d E'')_{max} \Rightarrow (\epsilon, \emptyset, E\neg a)_{max} \quad \text{if } \begin{cases} Ea^d E'' \text{ is inconsistent and} \\ E'' \text{ contains no decision literal} \end{cases} \\
\text{Propagate}'_{max} & (\epsilon, \emptyset, E)_{max} \Rightarrow (\epsilon, \emptyset, E\neg a)_{max} \quad \text{if } \begin{cases} E \text{ attacks } a \text{ or } a \text{ attacks } E \end{cases} \\
\text{Decide}'_{max} & (\epsilon, \emptyset, E)_{max} \Rightarrow (\epsilon, \emptyset, Ea^d)_{max} \quad \text{if } \begin{cases} E \text{ is consistent and} \\ \text{neither } a \text{ nor } \neg a \text{ occur in } E \text{ and} \\ \text{Propagate}'_{max} \text{ does not apply} \end{cases}
\end{array}$$

Succeeding and failing rules

$$\begin{array}{ll}
\text{Fail}_{max} & (\epsilon, \emptyset, E)_{max} \Rightarrow Ok(\epsilon) \quad \text{if } \begin{cases} E \text{ is incons. and decision-free} \\ \text{no other rule applies} \end{cases} \\
\text{Succeed}_{max} & (\epsilon, \emptyset, E)_{max} \Rightarrow (\epsilon, \emptyset, E)_{out} \quad \text{if } \begin{cases} \exists E' \in \epsilon : E \subseteq E' \text{ or} \\ \text{there is an argument } a \text{ s.t.} \\ E \text{ does not attack } a \text{ and} \\ a \text{ attacks } E \end{cases} \\
\text{Fail}_{out} & (\epsilon, \emptyset, E)_{out} \Rightarrow (\epsilon, \emptyset, E\perp)_{max} \quad \text{if } \begin{cases} \text{no other rule applies} \end{cases} \\
\text{Succeed}_{out} & (\epsilon, \emptyset, E)_{out} \Rightarrow (\epsilon \cup \{e(E)\}, \emptyset, E\perp)_{max} \quad \text{if } \begin{cases} \text{no other rule applies} \end{cases}
\end{array}$$

Fig. 8. The rules of the graph DIRECT^F .

condition “there is an argument a such that E does not attack a and a attacks E'' of the rule Fail_{out} .”

The graph DIRECT^F for an AF F is defined by the states over $\text{atoms}(F)$ and the transition rules presented in Figure 8. Its initial state is $(\emptyset, \emptyset, \emptyset)_{max}$. The structure of the graph is similar to that of ENUM_F^F . It differs from this graph in two ways. First, it has only one subprocedure. Second, the rules of the oracle differ from the previous oracle rules since they are not a call to a SAT solver; we primed them to emphasize the difference.

More precisely, among the oracle rules, propagation now only occurs so as to negatively add an atom if it attacks or is attacked by an atom of the extension being built. The Decide'_{max} rule only adds atoms positively, which is useful in Algorithm 2 of [20], but does not seem to be crucial here. When a record assigning all arguments is found, the rule Succeed_{max} is applied so as to allow the test of the outer rules to be carried on. If the record corresponds to a preferred extension, then it is stored by Succeed_{out} and the process of trying all possible records continues. In both Succeed_{out} and Fail_{out} , the use of one of the rules $\text{Backtrack}'_{max}$ or Fail_{max} is forced by making the record inconsistent. This way the process of browsing records is forced to continue.

Theorem 3. *For any AF F , the graph DIRECT^F is finite, acyclic and the only terminal state reachable from its initial state is $Ok(\epsilon)$ where $\epsilon = \text{prf}(F)$.*

3.4 Combining Algorithms

We can now define a new algorithm which is a combination of the PrefSat approach and the dedicated approach. In fact, it replaces the loop of SAT-calls for

$$\begin{array}{l}
\text{Succeeding and failing rules} \\
\textit{Succeed}_{base} (\epsilon, \emptyset, E)_{base} \Rightarrow (\epsilon, \emptyset, e(E))_{max} \quad \text{if } \left\{ \begin{array}{l} \text{no other rule applies} \\ \text{no other rule applies} \end{array} \right. \\
\textit{Succeed}_{out} (\epsilon, \emptyset, E)_{out} \Rightarrow (\epsilon \cup \{e(E)\}, \emptyset, \emptyset)_{base} \quad \text{if } \left\{ \begin{array}{l} \text{no other rule applies} \\ \text{no other rule applies} \end{array} \right.
\end{array}$$

Fig. 9. The rules of the graph MIX-PRF_F^F .

maximizing a complete extension of PrefSat by a part of the dedicated algorithm of [20]. In particular, instead of having subsequent oracle calls for maximization, we utilize the dedicated algorithm with a different initialization and stop already when the first preferred extension has been found. The graph MIX-PRF_F^F representing this algorithm consists of the oracle rules and the rules $\textit{Succeed}_{max}$ and \textit{Fail}_{out} of DIRECT^F , the *base*-oracle rules and the rule \textit{Fail}_{base} of ENUM_F^F and the rules in Figure 9. The initial state is $(\emptyset, \emptyset, \emptyset)_{base}$.

As in ENUM_F^F , whenever a $\textit{Succeed}_{base}$ rule is applied, a complete extension has been generated and it has to be validated or extended by the subprocedure identified with *max*. When $\textit{Succeed}_{max}$ is applied, a preferred extension has been found and the search for another complete extension can be started. Whenever an extension has been found by procedure *base*, there is a preferred extension that is a superset of the found extension. Hence, there is no need for a rule \textit{Fail}_{max} , since subprocedure *max* will always succeed.

Theorem 4. *For any AF F , the graph MIX-PRF_F^F is finite, acyclic and the only terminal state reachable from its initial state is $Ok(\epsilon)$ where $\epsilon = \text{prf}(F)$.*

4 Discussion and Conclusions

In this paper we have shown the applicability and the advantages of using a rigorous formal way for describing certain algorithms for solving decision problems for AFs through graph-based abstract solvers instead of pseudo-code-based descriptions. Both SAT-based and dedicated approaches for solving hard problems have been analyzed and compared. Moreover, by a combination of these approaches we have obtained a novel algorithm for enumeration of preferred extensions.

Our work shows the potential of abstract transition systems to describe, compare and combine algorithms also in the research field of abstract argumentation, as already happened in, e.g. SAT, SMT and ASP. In particular, the last feature, which allows the design of new solving procedures by combining reasoning modules from different algorithms, seems to be particularly appealing. However, we do not claim about the efficiency of a new tool built on this basis, given that it usually requires many iterations of theoretical analysis, practical engineering, and domain-specific optimizations to develop efficient systems.

We have focused on the well-studied preferred semantics, and we have presented core algorithms. However, the machinery can be easily employed to de-

scribing algorithms for solving other reasoning tasks, such as credulous acceptance, or different semantics, e.g. semi-stable and stage semantics, as employed in CEGARTIX [14]. Moreover, specific optimization techniques can be taken into account by means of modular addition of transition rules to the graph describing the core parts of the algorithms. As future work we plan to make these points more concrete.

Concerning further future work we envisage to formally describe further algorithms for reasoning tasks within abstract argumentation (e.g. [11, 18], see [8] for a comprehensive survey). In particular, the results of the upcoming competition will suggest promising candidates for the application of the newly gained technique of algorithm combination via abstract solvers.

Acknowledgements

This work has been funded by the Austrian Science Fund (FWF) through project I1102, and by Academy of Finland through grants 251170 COIN and 284591.

References

1. Baroni, P., Caminada, M.W.A., Giacomin, M.: An introduction to argumentation semantics. *The Knowledge Engineering Review* 26(4), 365–410 (2011)
2. Bench-Capon, T.J.M., Dunne, P.E.: Argumentation in artificial intelligence. *Artificial Intelligence* 171(10-15), 619–641 (2007)
3. Besnard, P., Doutre, S.: Checking the Acceptability of a Set of Arguments. In: Delgrande, J.P., Schaub, T. (eds.) *Proceedings of the Tenth International Workshop on Non-Monotonic Reasoning, NMR 2004*. pp. 59–64 (2004)
4. Brochenin, R., Lierler, Y., Maratea, M.: Abstract disjunctive answer set solvers. In: Schaub, T., Friedrich, G., O’Sullivan, B. (eds.) *Proceedings of the 21st European Conference on Artificial Intelligence, ECAI 2014. Frontiers in Artificial Intelligence and Applications*, vol. 263, pp. 165–170. IOS Press (2014)
5. Cerutti, F., Dunne, P.E., Giacomin, M., Vallati, M.: Computing preferred extensions in abstract argumentation: A SAT-based approach. In: Black, E., Modgil, S., Oren, N. (eds.) *Proceedings of the Second International Workshop on Theory and Applications of Formal Argumentation, TAFA 2013. Lecture Notes in Computer Science*, vol. 8306, pp. 176–193. Springer (2014), http://dx.doi.org/10.1007/978-3-642-54373-9_12
6. Cerutti, F., Giacomin, M., Vallati, M.: ArgSemSAT: Solving argumentation problems using SAT. In: Parsons, S., Oren, N., Reed, C., Cerutti, F. (eds.) *Proceedings of the Fifth International Conference on Computational Models of Argument, COMMA 2014. Frontiers in Artificial Intelligence and Applications*, vol. 266, pp. 455–456. IOS Press (2014)
7. Cerutti, F., Oren, N., Strass, H., Thimm, M., Vallati, M.: A benchmark framework for a computational argumentation competition. In: Parsons, S., Oren, N., Reed, C., Cerutti, F. (eds.) *Proceedings of the Fifth International Conference on Computational Models of Argument, COMMA 2014. Frontiers in Artificial Intelligence and Applications*, vol. 266, pp. 459–460. IOS Press (2014)

8. Charwat, G., Dvořák, W., Gaggl, S.A., Wallner, J.P., Woltran, S.: Methods for Solving Reasoning Problems in Abstract Argumentation - A Survey. *Artificial Intelligence* 220, 28–63 (2015)
9. Davis, M., Logemann, G., Loveland, D.: A machine program for theorem proving. *Communications of the ACM* 5(7), 394–397 (1962)
10. Dimopoulos, Y., Torres, A.: Graph theoretical structures in logic programs and default theories. *Theoretical Computer Science* 170(1-2), 209–244 (1996)
11. Doutre, S., Mengin, J.: Preferred extensions of argumentation frameworks: Query answering and computation. In: Goré, R., Leitsch, A., Nipkow, T. (eds.) *Proceedings of the First International Joint Conference on Automated Reasoning, IJCAR 2001*. *Lecture Notes in Computer Science*, vol. 2083, pp. 272–288. Springer (2001)
12. Dung, P.M.: 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)
13. Dunne, P.E., Bench-Capon, T.J.M.: Coherence in finite argument systems. *Artificial Intelligence* 141(1/2), 187–203 (2002)
14. Dvořák, W., Jarvisalo, M., Wallner, J.P., Woltran, S.: Complexity-sensitive decision procedures for abstract argumentation. *Artificial Intelligence* 206, 53–78 (2014)
15. Dvořák, W., Woltran, S.: Complexity of semi-stable and stage semantics in argumentation frameworks. *Information Processing Letters* 110(11), 425–430 (2010)
16. Lierler, Y.: Abstract answer set solvers with backjumping and learning. *Theory and Practice of Logic Programming* 11(2-3), 135–169 (2011)
17. Lierler, Y.: Relating constraint answer set programming languages and algorithms. *Artificial Intelligence* 207, 1–22 (2014)
18. Modgil, S., Caminada, M.W.A.: Proof theories and algorithms for abstract argumentation frameworks. In: Rahwan, I., Simari, G.R. (eds.) *Argumentation in Artificial Intelligence*, pp. 105–129. Springer (2009)
19. Nieuwenhuis, R., Oliveras, A., Tinelli, C.: Solving SAT and SAT modulo theories: From an abstract davis–putnam–logemann–loveland procedure to DPLL(T). *Journal of the ACM* 53(6), 937–977 (2006)
20. Nofal, S., Atkinson, K., Dunne, P.E.: Algorithms for decision problems in argument systems under preferred semantics. *Artificial Intelligence* 207, 23–51 (2014)
21. Rahwan, I., Simari, G.R. (eds.): *Argumentation in Artificial Intelligence*. Springer (2009)