Computing Stable Conclusions under the Weakest-Link Principle in the ASPIC+ Argumentation Formalism

Tuomo Lehtonen[1] , **Johannes P. Wallner**[2] , **Matti Järvisalo**[1]

[1]University of Helsinki, Helsinki, Finland
[2]Graz University of Technology, Graz, Austria
{tuomo.lehtonen, matti.jarvisalo}@helsinki.fi, wallner@ist.tugraz.at

Abstract

Rephrasing argumentation semantics in terms of subsets of defeasible elements allows for gaining new insights for reasoning about acceptance in established fragments of the central structured argumentation formalism of ASPIC$^+$. We provide a non-trivial generalization of these recent results, capturing preferences in ASPIC$^+$. In particular, considering preferences under the weakest-link principle, we show that the stable semantics can be phrased in terms of subsets of defeasible elements. We employ the rephrasing for establishing both complexity results and practical algorithms for reasoning about acceptance in this variant of ASPIC$^+$. Justified by completeness for the second level of the polynomial hierarchy, we develop an iterative answer set solving based approach to reasoning about acceptance under the so-called elitist lifting in ASPIC$^+$ frameworks. Our implementation of the approach scales well in practice.

1 Introduction

Argumentation is an established area of knowledge representation and reasoning research, with the fundamental aim of drawing conclusions from internally inconsistent or incomplete knowledge bases (Baroni et al. 2018; Gabbay et al. 2021; Atkinson et al. 2017). Arguments most often have an intrinsic structure made explicit through derivations from more basic structures. Computational models for structured argumentation (Bondarenko et al. 1997; Cyras et al. 2018; García and Simari 2014; García and Simari 2018; Besnard and Hunter 2008; Besnard and Hunter 2018)—ASPIC$^+$ (Modgil and Prakken 2018; Modgil and Prakken 2013) among the most prominent formalisms—enable making the internal structure of arguments explicit. In its general form, ASPIC$^+$ allows for arguments that combine strict inference rules—capturing deductively valid inferences—and defeasible inference rules—capturing presumptive inference—as well as accounting for preferential information. On one hand, this generality enables application in various settings, including legal reasoning (Prakken et al. 2015; Prakken 2012), ontology-based data access (Yun and Croitoru 2016), intelligence analysis (Toniolo et al. 2015), and information extraction for online crime reports (Schraagen et al. 2018). On the other hand, developing insights on the complexity and algorithmic aspects of reasoning in ASPIC$^+$ is still called for and remains a challenge.

Recently, an answer set programming (ASP) (Niemelä 1999; Brewka et al. 2015) approach to reasoning in ASPIC$^+$ without preferences was proposed (Lehtonen, Wallner, and Järvisalo 2020). Based on phrasing argumentation semantics in terms of subsets of defeasible elements in ASPIC$^+$, the approach avoids a potentially exponential translation to Dung's abstract argumentation frameworks (Dung 1995) employed earlier for realizing reasoning in ASPIC$^+$ (Snaith and Reed 2012; Thimm 2017) altogether, and instead employs an ASP solver on a direct ASP encoding on the level of ASPIC$^+$ for the reasoning task at hand.

In this work we study possibilities of generalizing these recent insights to preferential reasoning in ASPIC$^+$. In particular, we consider an instantiation of ASPIC$^+$ composed of atomic sentences, including axioms and ordinary premises, and allowing asymmetric negation. Towards a computational approach supporting preferences, we formally extend the foundations of rephrasing semantics developed earlier (Lehtonen, Wallner, and Järvisalo 2020) to cover preferences under the central weakest-link principle and elitist lifting (Modgil and Prakken 2018), focusing on the stable semantics as the first goal.

We provide both new complexity results as well as algorithms for credulous and skeptical reasoning, together with an experimental evaluation of an implementation of the algorithms. The non-trivial extension we develop of the earlier rephrasing to incorporate preferential reasoning is central for these contributions. In terms of complexity results, we show Σ_2^P and Π_2^P completeness for credulous and skeptical reasoning, respectively, for ASPIC$^+$ frameworks satisfying prominent rationality criteria (Modgil and Prakken 2018). Furthermore, we establish that reasoning has milder complexity (NP- and coNP-complete) for the case without defeasible rules, thereby being of the same complexity as the case without preferences (Lehtonen, Wallner, and Järvisalo 2020). Our rephrasing is vital here, since an explicated abstract framework yields structures not bounded polynomially, and these complexity classes restrict algorithms (among other aspects) to polynomial space consumption. The complexity results clearly indicate that inclusion of preferences increase complexity of reasoning in ASPIC$^+$. Moreover, our rephrasing highlights that the jump in complexity is due to a particular type of attack ("contradictory rebut") when combined with preferential reasoning.

In addition to the complexity results, in light of completeness of acceptance for the second level of the polynomial hierarchy, we employ the rephrasing in developing ASP-based counterexample-guided abstraction refinement algorithms for credulous and skeptical reasoning in ASPIC$^+$ with preferences, making use of incremental ASP solving techniques, and show that an implementation of the approach shows promising scalability up to hundreds of sentences.

After recalling necessary preliminaries of ASPIC$^+$ in Section 2, we overview our complexity results in Section 3. The extended rephrasing with preferences is presented in two parts: in Section 4 defeats on assumption are presented and Section 5 shows the rephrasing of stable semantics. In Section 6 the novel algorithms are presented, and experiments are presented in Section 7. We conclude after discussing related works in Section 8. An extended version of the paper with formal proofs is available via the authors' webpages.

2 ASPIC$^+$ Framework

We recall background on ASPIC$^+$ (Modgil and Prakken 2018; Modgil and Prakken 2013; Prakken 2010). We assume a set (language) $\mathcal{L}$ composed of atoms x. We start with contrariness.

Definition 1. *Let a contrary function be* $^- : \mathcal{L} \to 2^{\mathcal{L}}$. *We say that* $a \in \mathcal{L}$ *is a contrary of* $b \in \mathcal{L}$ *if* $a \in \overline{b}$ *and* $b \notin \overline{a}$. *We say that* a *is a contradictory of* b *if* $a \in \overline{b}$ *and* $b \in \overline{a}$.

That is, contraries represent an asymmetric relation, while contradictories are symmetric. When a and b are contradictory to each other we sometimes write $-a = b$ (or $a = -b$). An atom may be a contradictory to several atoms.

A central part of an ASPIC$^+$ framework is a knowledge base $\mathcal{K} \subseteq \mathcal{L}$ comprised of a defeasible part (ordinary premises $\mathcal{K}_p$) and a non-defeasible part (axioms $\mathcal{K}_n$).

Definition 2. *A knowledge base is a set* $\mathcal{K}_n \cup \mathcal{K}_p = \mathcal{K} \subseteq \mathcal{L}$, *with disjoint sets* $\mathcal{K}_n$ *(axioms) and* $\mathcal{K}_p$ *(ordinary premises).*

Another part of ASPIC$^+$ is a set of rules over $\mathcal{L}$, denoted by $\mathcal{R}$. This set is composed of defeasible rules $a_1, \ldots, a_n \Rightarrow b$ and strict rules $a_1, \ldots, a_n \to b$. We denote the set of defeasible rules by $\mathcal{R}_d$ and the set of strict rules by $\mathcal{R}_s$. When we do not distinguish between strict or defeasible rules, we write $a_1, \ldots, a_n \rightsquigarrow b$. A partial function $n : \mathcal{R}_d \to \mathcal{L}$ gives names to defeasible rules. For a rule $r = a_1, \ldots, a_n \rightsquigarrow b$, we denote its head by $head(r) = b$ and its body by $body(r) = \{a_1, \ldots, a_n\}$.

For preferences, we consider preorders (i.e., reflexive and transitive binary relations) $\leq = \leq_p \cup \leq_d$, composed of a preorder on ordinary premises $\leq_p$ and a preorder on defeasible rules $\leq_d$.

Definition 3. *An argumentation theory (AT) is a tuple* $(\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$, *with a knowledge base* $\mathcal{K} \subseteq \mathcal{L}$, *rules* $\mathcal{R} = \mathcal{R}_d \cup \mathcal{R}_s$ *over* $\mathcal{L}$, *a contrary function* $^- : \mathcal{L} \to 2^{\mathcal{L}}$, *a partial function* $n : \mathcal{R}_d \to \mathcal{L}$, *and a preorder* $\leq = \leq_p \cup \leq_d$.

Each part of an AT is assumed to be finite. Arguments are constructed from parts of an AT. An argument represents a "derivation tree" starting from elements in the knowledge base and uses rules to derive a conclusion.

Definition 4. *Given an AT* $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$, *the set of arguments in* T *is inductively defined as follows.*

- *If* $x \in \mathcal{K}$, *then* $A = x$ *is an argument with* $\texttt{Conc}(A) = x$.
- *If* $A_1, \ldots, A_n$ *are arguments,* $x_i = \texttt{Conc}(A_i)$ *for* $1 \leq i \leq n$, *and* $(x_1, \ldots, x_n \rightsquigarrow x) \in \mathcal{R}$, *then* $A = A_1, \ldots, A_n \rightsquigarrow x$ *is an argument with* $\texttt{Conc}(A) = x$.

There are no other arguments.

We make use of the following shorthands.

Definition 5. *Let* $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$ *be an AT, and* A *an argument in* T.

- *If* $A = x \in \mathcal{K}$ *then* $\texttt{Sub}(A) = \{A\}$ *and* $\texttt{Rules}(A) = \emptyset$.
- *If* $A = A_1, \ldots, A_n \rightsquigarrow x$, *then*
  $\texttt{Sub}(A) = \{A\} \cup \bigcup_{i=1}^n \texttt{Sub}(A_i)$,
  $\texttt{TopRule}(A) = (\texttt{Conc}(A_1), \ldots, \texttt{Conc}(A_n) \rightsquigarrow x)$, *and*
  $\texttt{Rules}(A) = \{\texttt{TopRule}(A)\} \cup \bigcup_{i=1}^n \texttt{Rules}(A_i)$.

Further, $\texttt{Prem}(A) = \texttt{Sub}(A) \cap \mathcal{K}$, $\texttt{Prem}_\texttt{d}(A) = \texttt{Prem}(A) \cap \mathcal{K}_p$, *and* $\texttt{DefRules}(A) = \texttt{Rules}(A) \cap \mathcal{R}_d$.

That is, we define shorthands for the subarguments (Sub) of an argument, the rules and defeasible rules in the argument (Rules and DefRules), the topmost rule (TopRule), the premises of the argument within the knowledge base (Prem), and the ordinary premises ($\texttt{Prem}_\texttt{d}$). Further, $\texttt{defPart}(A) = \texttt{Prem}_\texttt{d}(A) \cup \texttt{DefRules}(A)$. If $A \in \mathcal{K}$, then $\texttt{TopRule}(A)$ is undefined. We extend the shorthands for a set of arguments $\mathcal{A}$ as $\texttt{Conc}(\mathcal{A}) = \{\texttt{Conc}(A) \mid A \in \mathcal{A}\}$ and $\texttt{TopRule}(\mathcal{A}) = \{\texttt{TopRule}(A) \mid A \in \mathcal{A}\}$. For each shorthand $f \in \{\texttt{Sub}, \texttt{Rules}, \texttt{DefRules}, \texttt{Prem}, \texttt{Prem}_\texttt{d}\}$ returning a set, we define $f(\mathcal{A}) = \bigcup_{A \in \mathcal{A}} f(A)$. An argument A is an immediate subargument of $B = A_1, \ldots, A_n \rightsquigarrow x$ if $A \in \{A_1, \ldots, A_n\}$. We allow only finite structures as arguments (i.e., arguments which are "trees" of finite size), and consider as arguments those arguments A for which $\texttt{Sub}(A)$ is finite (disallowing infinite chaining of rules, e.g., via $x \rightsquigarrow x$), as also defined by Modgil and Prakken (2018).

Conflicts among arguments are defined via attacks between arguments.

Definition 6. *Given an AT* $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$ *and two arguments* A *and* B *in* T, *argument* A *attacks argument* B *iff* A *undercuts, rebuts, or undermines* B, *where*

- A *undercuts* B *(on* B'*) iff* $\texttt{Conc}(A) \in \overline{n(r)}$ *for some* $B' \in \texttt{Sub}(B)$ *such that* $\texttt{TopRule}(B') = r$ *is defeasible;*
- A *rebuts* B *(on* B'*) iff* $\texttt{Conc}(A) \in \overline{x}$ *for some* $B' = B_1, \ldots, B_n \Rightarrow x \in \texttt{Sub}(B)$*; and*
- A *undermines* B *(on* x*) iff* $\texttt{Conc}(A) \in \overline{x}$ *and* $x \in \texttt{Prem}_\texttt{d}(B)$.

That is, an argument attacks another argument on the defeasible parts of the latter. Ordinary premises can be attacked by undermining, and defeasible rules can be attacked by rebutting the conclusion or undercutting the rule itself. For rebuts and undermining one distinguishes further if $\texttt{Conc}(A)$ and x are contraries or contradictories: in the former case we say that A contrary undermines (rebuts) B and in the latter that A contradictory undermines (rebuts) B.

The preorders $\leq_p$ and $\leq_d$ on the defeasible parts are lifted to strict partial orders (i.e., to irreflexive, asymmetric, and transitive binary relations) $\triangleleft_p$ and $\triangleleft_d$, via lifting operators. We focus on the elitist lifting (Modgil and Prakken 2018). As usual, the strict part of $\leq_p$ ($\leq_d$) is defined by $x <_p y$ iff $x \leq_p y$ and $y \not\leq_p x$ (same for $\leq_d$).

Definition 7. *Let $\leq$ be a preorder on a set X. Define $\triangleleft^{Eli}$ for two non-empty Y, $Z \subseteq X$ by $Y \triangleleft^{Eli} Z$ iff $\exists a \in Y$ s.t. $\forall b \in Z$ we have $a < b$. Moreover, $\emptyset \not\triangleleft^{Eli} Y$ and $Z \triangleleft^{Eli} \emptyset$ for each non-empty Z.*

That is, Z is preferred to Y if there is at least one element in Y that is strictly less preferred to each element in Z. The empty set is a special case, being strictly preferred to each non-empty set, and cannot be less preferred than any set.

For two preorders $\leq_p$ and $\leq_d$ and their liftings $\triangleleft^{Eli} = \triangleleft_p^{Eli} \cup \triangleleft_d^{Eli}$, one defines a strict partial order $\prec$ on arguments, denoting the preference (ranking) on arguments using certain principles. We omit the superscript *Eli* in the remainder of the paper. We focus here on the weakest-link principle, by which one considers all defeasible elements of arguments in the comparison. An argument B is strictly preferred to A ($A \prec B$) whenever

- if $\texttt{DefRules}(A) = \texttt{DefRules}(B) = \emptyset$, then $\texttt{Prem}_\texttt{d}(A) \triangleleft \texttt{Prem}_\texttt{d}(B)$;
- if $\texttt{Prem}_\texttt{d}(A) = \texttt{Prem}_\texttt{d}(B) = \emptyset$, then $\texttt{DefRules}(A) \triangleleft \texttt{DefRules}(B)$;
- else $\texttt{Prem}_\texttt{d}(A) \triangleleft \texttt{Prem}_\texttt{d}(B)$ and $\texttt{DefRules}(A) \triangleleft \texttt{DefRules}(B)$.

Defeats between arguments, i.e., successful attacks, are defined based on argument preferences as follows.

Definition 8. *Given an AT $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$ and two arguments A and B in T, argument A defeats argument B iff A successfully undercuts, rebuts, or undermines B, where*

- *A successfully undercuts B if A undercuts B;*
- *A successfully rebuts B (on B') iff A contrary rebuts B, or A rebuts B on B' and $A \not\prec B'$; and*
- *A successfully undermines B (on x) iff A contrary undermines B, or A undermines B on x and $A \not\prec x$.*

In other words, an undercut always succeeds, as do contrary rebuts and undermining attacks. For intuition, an undercut attacks the very applicability of a rule, and such an attack takes precedence over preferences between arguments. Similarly, as contraries signal an asymmetric incompatibility of two atoms, the presence of a contrary of an atom x results in x being attacked regardless of preferences. Modgil and Prakken (2013; 2018) provide a discussion on this. For the contradictory variants of rebut and undermining (i.e., when conclusion of A and the ordinary premise or conclusion of a defeasible rule in B are contradictories of each other), the preference order $\prec$ on arguments decides whether the attack succeeds: if A is strictly less preferred to the attacked subargument, the attack fails, otherwise it is a defeat.

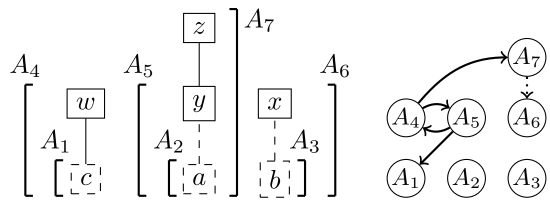

Figure 1: Example AT and corresponding AF

Example 1. *Let $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$ be an AT with $\mathcal{L} = \{a, b, c, w, x, y, z\}$, $\mathcal{K}_p = \{a, b, c\}$, $\mathcal{K}_n = \emptyset$, $\overline{x} = \{z\}$, $\overline{z} = \{x\}$, $\overline{c} = \{y\}$, $\overline{n(r_1)} = \{w\}$, $\mathcal{R}_d = \{(r_1 : a \Rightarrow y), (r_2 : b \Rightarrow x)\}$, and $\mathcal{R}_s = \{(y \to z), (c \to w)\}$. We write names of defeasible rules by r_i before the rule. Moreover, let $a \leq_p b$ and $r_1 \leq_d r_2$. The AT is shown in Figure 1(left) with the arguments it gives rise to. Defeasible premises or inference is marked with dashed lines. The arguments and their defeat relation are shown on the right side of this figure. It holds that A_4 undercuts both A_5 and A_7 (on A_5), A_5 contrary undermines both A_1 and A_4 (on A_1), and A_7 contradictory rebuts A_6. Since x and z are contradictory to each other, but A_7 concludes z via a strict rule, it is the case that A_6 does not contradictory rebut A_7. Except for the last attack, all are successful and thus defeats, since $A_7 \prec A_6$ ($\texttt{Prem}_\texttt{d}(A_7) = \{a\} \triangleleft_p \{b\} = \texttt{Prem}_\texttt{d}(A_6)$ due to $a <_p b$ and $\texttt{DefRules}(A_7) = \{r_1\} \triangleleft_d \{r_2\} = \texttt{DefRules}(A_6)$ due to $r_1 <_d r_2$). The unsuccessful contradictory rebut is denoted as a dotted arrow.*

In ASPIC$^+$, there are important conditions which ensure satisfaction of rationality postulates, which might fail in the very general case (Modgil and Prakken 2018; Caminada 2018). It turns out that these conditions are useful for our computational approach, as well. These conditions constrain the set of strict rules in an AT, and we make use of three such conditions here, borrowing from Modgil and Prakken (2018).

A set of strict rules $\mathcal{R}_s$ is said to be closed under transposition if for each $a_1, \ldots, a_n \to b \in \mathcal{R}_s$ it holds that for each i, $1 \leq i \leq n$, we have $a_1, \ldots, a_{i-1}, b', a_{i+1}, \ldots, a_n \to a'_i \in \mathcal{R}_s$ for all contradictories $-b' = b$ and $-a'_i = a_i$, and at least one contradictory of each must exist. An AT T is strict-consistent if there are no arguments A, B with $\texttt{defPart}(A) = \texttt{defPart}(B) = \emptyset$ s.t. $\texttt{Conc}(A)$ is a contrary of $\texttt{Conc}(B)$, or $\texttt{Conc}(A)$ and $\texttt{Conc}(B)$ are contradictory to each other.

Definition 9. *We say an AT T is well-formed if*

- *$\mathcal{R}_s$ is closed under transposition,*
- *T is strict-consistent, and*
- *if x is a contrary of y then $y \notin \mathcal{K}_n$ and $y \neq head(r)$ for all $r \in \mathcal{R}_s$.*

For intuition, these conditions aim to avoid certain inconsistencies that may arise, with strict-consistency likely the most immediate: if an AT is not strict-consistent then there are a arguments composed only of axioms and strict rules concluding contrary or contradictory atoms. For details we refer the reader to the work of Modgil and Prakken (2018). We remark that we are using a subset of the conditions of "well-defined" ATs as defined by Modgil and Prakken (2018).

Semantics of ATs are defined via a translation to (abstract) argumentation frameworks (AFs) (Dung 1995). An AF is a pair $F = (\mathcal{A}, \mathcal{D})$ of a set of (abstract) arguments $\mathcal{A}$ and defeats $\mathcal{D} \subseteq \mathcal{A} \times \mathcal{A}$ between arguments. If $(A, B) \in \mathcal{D}$ we say that A defeats B. Similarly, $\mathcal{S} \subseteq \mathcal{A}$ defeats $B \in \mathcal{A}$ if there is an $A \in \mathcal{S}$ with $(A, B) \in \mathcal{D}$. We say that $\mathcal{S}$ defends an argument A if for each $B \in \mathcal{A}$ such that $(B, A) \in \mathcal{D}$, there is a $C \in \mathcal{S}$ such that $(C, B) \in \mathcal{D}$. We consider the AF semantics of conflict-free and admissible sets, and complete and stable extensions, with the corresponding functions $\sigma \in \{cf, adm, com, stb\}$. A semantics $\sigma(F) \subseteq 2^{\mathcal{A}}$ returns a set of extensions. An extension under a semantics σ is a σ-extension for short.

Definition 10. *Given an AF $F = (\mathcal{A}, \mathcal{D})$, a set $\mathcal{E} \subseteq \mathcal{A}$ is conflict-free (in F) if there are no A, B in $\mathcal{E}$ such that $(A, B) \in \mathcal{D}$. The set of all conflict-free sets of F is denoted by $cf(F)$. For an $\mathcal{E} \in cf(F)$, we have*

- *$\mathcal{E} \in adm(F)$ iff each $A \in \mathcal{E}$ is defended by $\mathcal{E}$;*
- *$\mathcal{E} \in com(F)$ iff $\mathcal{E} \in adm(F)$ and each A defended by $\mathcal{E}$ is in $\mathcal{E}$;*
- *$\mathcal{E} \in stb(F)$ iff $\mathcal{E}$ defeats each argument in $\mathcal{A} \setminus \mathcal{E}$.*

ATs can be translated to AFs as follows.

Definition 11. *Let $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$ be an AT. An AF $F = (\mathcal{A}, \mathcal{D})$ corresponds to T if $\mathcal{A}$ is the set of all arguments in T and $\mathcal{D}$ the defeat relation based on T.*

Reasoning on ATs consists of checking whether a queried atom is warranted, by asking (credulously) whether there is a σ-extension having an argument concluding the atom, or (skeptically) whether all σ-extensions have such an argument.

Definition 12. *Given an AT $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$ and an AF F corresponding to T, we say that $x \in \mathcal{L}$ is*

- *skeptically justified in T under semantics σ if in each $\mathcal{E} \in \sigma(F)$ there is an $A \in \mathcal{E}$ with $\texttt{Conc}(A) = x$;*
- *credulously justified in T under semantics σ if there is an $\mathcal{E} \in \sigma(F)$ with an $A \in \mathcal{E}$ s.t. $\texttt{Conc}(A) = x$.*

Example 2. *Continuing Example 1, the AF corresponding to the AT is shown in Figure 1 (right). There are two stable extensions: $\mathcal{E}_1 = \{A_1, A_2, A_3, A_4, A_6\}$ and $\mathcal{E}_2 = \{A_2, A_3, A_5, A_6, A_7\}$. Since $\texttt{Conc}(A_4) = w$, $A_4 \in \mathcal{E}_1$, and there is no argument in $\mathcal{E}_2$ with conclusion w, it holds that w is credulously but not skeptically justified under stable semantics.*

A useful property is that each complete extension of an AF corresponding to an AT is closed under subarguments, i.e., if $\mathcal{E}$ is complete and $A \in \mathcal{E}$, then all $\texttt{Sub}(A)$ are in $\mathcal{E}$, as well. Note that stable extensions are also complete extensions.

Proposition 1 (Modgil and Prakken 2013)**.** *Let $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$ be an AT, $F = (\mathcal{A}, \mathcal{D})$ the AF corresponding to T, and $\mathcal{E} \in com(F)$. It holds that $\mathcal{E}$ is closed under subarguments.*

Building on earlier work (Lehtonen, Wallner, and Järvisalo 2020), we utilize the concept of so-called assumptions (P, D) for a given AT $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$, which represent parts of the defeasible elements: $P \subseteq \mathcal{K}_p$ and $D \subseteq \mathcal{R}_d$. We define $(P, D) \sqsubseteq (P', D')$ if $P \subseteq P'$ and $D \subseteq D'$.

Given a set of rules $\mathcal{R}$ and a set of atoms $L \subseteq \mathcal{L}$, we say that $x \in \mathcal{L}$ is derivable from L via $\mathcal{R}$, denoted by $L \vdash_{\mathcal{R}} x$, if (i) $x \in L$ or (ii) there is a sequence of rules $(r_1, \ldots, r_n)$ from $\mathcal{R}$ s.t. $head(r_n) = x$ and for each rule r_i it holds that each atom in the body of r_i is derived from rules earlier in the sequence or is in L, i.e., $body(r_i) \subseteq L \cup \bigcup_{j<i} head(r_j)$.

We extend the derivability notion to assumptions in a straightforward way: given an AT $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$ and an assumption (P, D) in T, we say that from (P, D) one can derive (in T) an atom $x \in \mathcal{L}$, denoted by $(P, D) \vdash_T x$, if $P \cup \mathcal{K}_n \vdash_{D \cup \mathcal{R}_s} x$, i.e., x is derivable from the defeasible elements in the assumption and all non-defeasible parts in the AT. The deductive closure of an assumption in T is then defined as $Th_T(P, D) = \{x \in \mathcal{L} \mid (P, D) \vdash_T x\}$. We say that a rule r is applicable by an assumption (P, D) if $body(r) \subseteq Th_T(P, D)$, i.e., all elements of the body of r can be derived using the assumption. For an assumption (P, D) in T and an argument A in T we say that A is based on (P, D) if A uses only defeasible elements from this assumption, i.e., A is based on (P, D) in T if $\texttt{DefRules}(A) \subseteq D$ and $\texttt{Prem}_{\texttt{d}}(A) \subseteq P$.

Assumptions and arguments have a direct connection.

Proposition 2 (Lehtonen, Wallner, and Järvisalo 2020)**.** *Let $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$ be an AT and (P, D) an assumption in T. There is an argument A based on (P, D) in T with $\texttt{Conc}(A) = x$ iff $(P, D) \vdash_T x$.*

3 Complexity of Reasoning

In this section, we provide an overview of our complexity results for preferential reasoning under stable semantics in ASPIC$^+$. The rephrasing of stable semantics essential for the formal proofs of the complexity results (as well as for the algorithms presented in Section 6) is postponed until Sections 4–5.

While an AF corresponding to a given AT is, in general, not bounded polynomially in size w.r.t. the given AT, we showed (Lehtonen, Wallner, and Järvisalo 2020) that one can define criteria on assumptions (P, D) so that there is a direct correspondence between assumptions and extensions of arguments, without explicating the extension (and without constructing the corresponding AF), with assumptions being bounded polynomially regarding the input AT. Such an approach is essential to show complexity results for classes in the polynomial hierarchy (which are restricted to polynomial space). Concretely, assumptions, and criteria on assumptions, allow for a design of algorithms that operate in polynomial space, and fall into a class of the polynomial hierarchy, in the form of a non-deterministic construction of a (P, D) assumption and subsequent verification of the conditions imposed by the semantics.

Preferential reasoning is important to ASPIC$^+$, but obtaining criteria on assumptions reflecting preferences is, as we will show, non-immediate, and, moreover, increases computational complexity. Recall that preferences in ASPIC$^+$ lead to a modified attack structure (the defeat relation), which means that preferences change conflicts in the corresponding AF, an object we have to avoid explicitly constructing in order to obtain tight complexity results.

We show that preferential reasoning, under stable semantics, the weakest-link principle, and elitist lifting in ASPIC$^+$, remains on the same level of complexity in case no defeasible rules are present as in the case no preferences are present. This result, in fact, aligns reasoning under stable semantics in ASPIC$^+$ and assumption-based argumentation (ABA) with certain kinds of preferences (called ABA$^+$) (Čyras and Toni 2016) from the view of complexity (Lehtonen, Wallner, and Järvisalo 2021).

Theorem 3. *For ATs without defeasible rules and stable semantics, credulous justification is NP-complete and skeptical justification is coNP-complete.*

While hardness for these cases follows in a rather direct fashion from existing results (e.g., see hardness on AFs (Dvořák and Dunne 2018) or ASPIC$^+$ without preferences (Lehtonen, Wallner, and Järvisalo 2020)), membership results are based on our rephrasing that characterizes stable semantics in terms of conditions on assumption sets, formally presented in the next subsections.

However, in contrast, when including defeasible rules we show that reasoning jumps to the second level of the polynomial hierarchy, under prominent instantiations of ASPIC$^+$ satisfying rationality postulates, concretely what we refer to as well-formed ATs. These results clearly set ASPIC$^+$ and ABA$^+$ apart, in terms of computational complexity of reasoning. Hardness holds even in the case of no strict rules.

Theorem 4. *For well-formed ATs and stable semantics, credulous justification is Σ_2^P-complete and skeptical justification is Π_2^P-complete. Hardness holds even when there are no strict rules.*

For intuition on why well-formedness, i.e., conditions usually used for showing rationality postulates, are useful for showing complexity results, these conditions also lead to properties that allow to align stable extensions and stable assumptions in a more direct way. In the following section we present a rephrasing of defeats to the viewpoint of assumptions, where we clearly distinguish between two types of defeats with one being the underlying reason for the complexity jump, and subsequently highlight useful properties following from assuming well-formed ATs.

4 Defeats on Assumptions

In this section we investigate defeats defined on assumptions (P, D), and their relation to defeats on arguments, as a precursor to rephrasing semantics on assumptions.

It will be useful to classify defeats into two categories: one contains successful contradictory rebuts and the other category, which we call individual defeats, contains all other types of defeats.

Definition 13. *Let $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$ be an AT, and A and B two arguments in T. We say that A individually defeats B if A*

- *successfully undercuts B,*
- *contrary rebuts B, or*
- *successfully undermines B.*

The reason for considering these two classes will become clear when considering their corresponding definitions on assumptions and their respective complexity. We define individual defeats on assumptions as follows.

Definition 14. *Let $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$ be an AT, and (P, D) be an assumption in T. We say that (P, D) individually defeats an $x \in \mathcal{K}_p \cup \mathcal{R}_d$ if*

- *$x = p \in \mathcal{K}_p$ and*
  - *there is a contrary of p in $Th_T(P, D)$ (contrary undermine) or*
  - *there is a contradictory of p in $Th_T(P', D)$ with $P' = \{p' \in P \mid p' \not< p\}$ (contradictory undermine),*
- *$x = r \in \mathcal{R}_d$ and*
  - *$\overline{n(r)} \cap Th_T(P, D) \neq \emptyset$ (undercut),*
  - *there is a contrary of $head(r)$ in $Th_T(P, D)$ (contrary rebut).*

These conditions reflect individual defeats on arguments. Assumption (P, D) undermines an ordinary premise if one can derive a contrary of the premise or a contradictory using not less preferred premises. Undercuts and contrary rebuts on assumptions directly reflect undercuts and contrary rebuts on arguments (which do not require preference handling).

Next we show that an assumption (P, D) individually defeating x represents a set of arguments $\mathcal{S}$ for which the following holds: for each argument B containing x (as an ordinary premise or a defeasible rule) there is an argument $A \in \mathcal{S}$ individually defeating B. This motivates the name: the defeat targets a single defeasible element and does not require (more) context information on the defeated argument.

Proposition 5. *Let $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$ be an AT, (P, D) be an assumption in T, and $x \in \mathcal{K}_p \cup \mathcal{R}_d$. There is an argument A based on (P, D) s.t. A individually defeats an argument B in T whenever $x \in \mathtt{Prem_d}(B) \cup \mathtt{DefRules}(B)$ iff (P, D) individually defeats x.*

What remains is the case of contradictory rebut. It turns out that phrasing contradictory rebuts on assumptions leads to a more complex condition.

Definition 15. *Let $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$ be an AT, and (P, D) and (P', D') be assumptions in T. Further, let $W = Th_T(P'', D) \cup Th_T(P, D'')$ with*

- *$P'' = \{p \in P \mid \exists p' \in P', p \not< p'\}$ and*
- *$D'' = \{r \in D \mid \exists r' \in D', r \not< r'\}$.*

We say (P, D) contradictory rebuts (P', D') on $r' \in D'$ if

- *P' is non-empty and there is a contradictory of $head(r')$ in W, or*
- *P' is empty and there is a contradictory of $head(r')$ in $Th_T(P, D'')$.*

Sets P'' and D'' are chosen in a way to accommodate the preference induced by the elitist lifting (for each element there is one not more preferred). The case distinction reflects the weakest-link principle (if P' is empty only defeasible rules are taken into account; D' is non-empty).

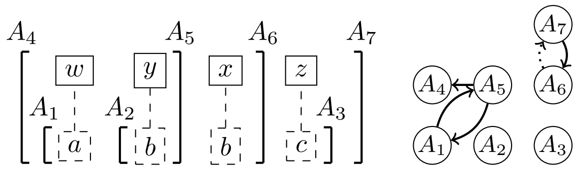

Figure 2: AT with $\mathcal{R}_s = \emptyset$ and corresponding AF

For individual defeats it holds that if an assumption individually defeats a defeasible element x, then this assumption (and its corresponding arguments) defeats all assumptions (arguments) individually containing the defeasible element x. In contrast, contradictory defeats are defined on assumptions contradictory rebutting a *specific* assumption (and as we will see, making no claim on sub-assumptions, for instance), which is why checking contradictory rebuts is computationally harder than checking individual attacks.

Example 3. *Let $T = (\mathcal{L}, \mathcal{R}, n, \overline{}, \mathcal{K}, \leq)$ be an AT with $\mathcal{L} = \{a, b, c, w, x, y, z\}$, $\mathcal{K}_p = \{a, b, c\}$, $\mathcal{K}_n = \emptyset$, $\overline{x} = \{z\}$, $\overline{z} = \{x\}$, $\overline{a} = \{y\}$, $\overline{y} = \{a\}$, $\mathcal{R}_s = \emptyset$, and $\mathcal{R}_d = \{(r_1 : b \Rightarrow x), (r_2 : b \Rightarrow y), (r_3 : c \Rightarrow z), (r_4 : a \Rightarrow w)\}$. Further, let $a \leq_p b$, $b \leq_p c$, and $r_1 \leq_d r_3$. The AT and corresponding AF are shown in Figure 2. Argument A_5 successfully contradictory undermines A_1 (and A_4 on A_1) and A_1 successfully contradictory rebuts A_5. The latter holds because $\texttt{Prem}_\texttt{d}(A_1) = \{a\} \lhd \{b\} = \texttt{Prem}_\texttt{d}(A_5)$ due to $a <_p b$, but $\texttt{DefRules}(A_1) = \emptyset$ and $\emptyset \ntriangleleft \{r_2\} = \texttt{DefRules}(A_5)$. Argument A_7 successfully contradictory rebuts A_6, but A_6 does not successfully contradictory rebut A_7, due to $A_6 \prec A_7$ ($b <_p c$ and $r_1 <_d r_3$). Consider $P = \{b, c\}$ and $D = \{r_1, r_2\}$. It holds that (P, D) individually defeats a (because we can derive the contradictory of a from (P, D) and neither b nor c is less preferred to a). For $P' = \{a, c\}$ and $D' = \{r_3, r_4\}$ it holds that (P, D) contradictory rebuts (P', D') on r_3, since with $P'' = \{b, c\} = P$ and $D'' = \{r_1, r_2\} = D$ we can derive a contradictory of $head(r_3) = z$ (namely x) by Definition 15. However, (P, D) does not contradictory rebut $(P^*, D^*) = (\{c\}, \{r_3\})$ on r_3 (the ingredients for A_7) since, again by Definition 15, neither from $(\{c\}, D)$ nor from $(P, \{r_2\})$ one can derive x (one requires b and defeasible rule r_1).*

That is, an assumption might contradictory rebut an assumption (P, D), but not $(P', D') \sqsubseteq (P, D)$. Next we state that if (P, D) contrary rebuts (P', D'), then there is an argument based on (P, D) successfully contradictory rebutting arguments containing exactly (P', D') as their defeasible part.

Proposition 6. *Let $T = (\mathcal{L}, \mathcal{R}, n, \overline{}, \mathcal{K}, \leq)$ be an AT, and (P, D) and (P', D') be assumptions in T. If (P, D) contradictory rebuts (P', D') on r then all arguments B in T with*

- $\texttt{Prem}_\texttt{d}(B) = P'$,
- $\texttt{DefRules}(B) = D'$, *and*
- $\texttt{TopRule}(B) = r$.

are successfully contradictory rebutted on B by an argument A based on (P, D).

5 Rephrasing Stability on Assumptions

In this section we make use of assumptions and defeats to define stable semantics on assumptions such that "stable assumptions" have a useful correspondence to stable extensions. In contrast to the case without preferences (Lehtonen, Wallner, and Järvisalo 2020), where a direct correspondence was possible, it turns out that defeasible rules, and even more, the *combination* of strict and defeasible rules, together with preferential reasoning, is a particular challenge for computation. General translation results (Modgil and Prakken 2018) that reduce an ASPIC$^+$ AT to an AT without strict rules or without defeasible rules are not directly usable in our case.[1]

We first show auxiliary results. First, a useful property is that if an argument attacks (but not necessarily defeats) another argument, this attack is present as a defeat, in a certain well-behaved manner, in case the underlying AT is well-formed (in particular closed under transposition). In case no strict rules exist at all, a rather strong condition holds: there is a defeat among subarguments of the two arguments involved in an attack. In case there are strict rules, which are closed under transposition, it can be the case that the defeat is "propagated" to a superargument of one of the two arguments (more precisely a superargument of a subargument of one of the two), but nevertheless in a controlled manner, as stated in the following lemma.

Lemma 7. *Let T be an AT with $\mathcal{R}_s$ closed under transposition, and A and B two arguments in T with $\texttt{TopRule}(A) \in \mathcal{R}_s$. If A unsuccessfully contradictory rebuts or unsuccessfully contradictory undermines B, then there is an argument C in T which defeats A and for every $C' \in \texttt{Sub}(C)$ with $\texttt{TopRule}(C')$ being defeasible or $C' \in \mathcal{K}_p$ it holds that $C' \in \texttt{Sub}(A) \cup \texttt{Sub}(B)$.*

Based on the lemma, the next proposition formalizes that attacks in well-formed ATs, do not, in a sense, disappear.

Proposition 8. *Let $T = (\mathcal{L}, \mathcal{R}, n, \overline{}, \mathcal{K}, \leq)$ be an AT, and argument A in T attacks argument B in T.*

- *If $\mathcal{R}_s = \emptyset$, then A defeats B or some $B' \in \texttt{Sub}(B)$ defeats A.*
- *If $\mathcal{R}_s$ is closed under transposition, then there is an argument C based on $(\texttt{Prem}_\texttt{d}(A) \cup \texttt{Prem}_\texttt{d}(B), \texttt{DefRules}(A) \cup \texttt{DefRules}(B))$ s.t. C defeats A or B.*

Example 4. *In the AT from Example 1, A_7 contradictory rebuts A_6, but A_7 does not defeat A_6. No subargument of A_6 defeats A_7. This exemplifies the fact that in presence of both strict and defeasible rules the first item of Proposition 8 does not hold. A simple example with only strict rules is obtained by considering two ordinary premises a and b, with $a <_p b$ and a strict rule $a \rightarrow x$ with x and b being contradictory. Then there are no defeats (the argument concluding x is strictly weaker than argument b); however, there is an attack based on contradictory undermining. Crucially the sets of strict rules are not closed under transposition.*

[1] While a translation of strict rules to defeasible ones has been proposed (Li and Parsons 2015), it turned out not to be valid (Li 2019) and hence not applicable for our case.

Another useful ingredient is a property related to subargument closure (see Proposition 1): if $\mathcal{E}$ is complete in an AT, then for each argument A whose defeasible elements are in $\mathcal{E}$ (i.e., $\texttt{defPart}(A) \subseteq \texttt{defPart}(\mathcal{E})$) it holds that $A \in \mathcal{E}$. In contrast to closure under subarguments, this kind of closure property does not hold in general (see extended version for an example); however, it does hold in case there are either no strict or no defeasible rules for complete semantics, and for well-formed ATs for stable semantics. Motivation for this property is that a correspondence of a (P, D) and an extension can be defined in a way that the extension contains all arguments that can be constructed using defeasible elements in (P, D) and strict components (strict rules, axioms) in an AT. Formally, for a semantics σ a set of ATs is closed under defeasible elements if for each $\mathcal{E} \in \sigma(F)$, for the AF F corresponding to T, it holds that all arguments A with $\texttt{defPart}(A) \subseteq \texttt{defPart}(\mathcal{E})$ are in $\mathcal{E}$.

Proposition 9. *The set of ATs with $\mathcal{R}_s = \emptyset$ or $\mathcal{R}_d = \emptyset$ is closed under defeasible elements for complete semantics and the set of ATs with $\mathcal{R}_s$ closed under transposition is closed under defeasible elements for stable semantics.*

We are now in a position to define stability on assumptions. We define w-stable assumptions for the case of well-formed ATs, and s-stable assumption for when there are only strict rules.

Definition 16. *Let $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$ be a well-formed AT. Assumption (P, D) of T is w-stable if all D are applicable by (P, D) and*

1. *$\nexists x, y \in Th_T(P, D)$ such that x is a contrary of y, x and y are contradictory, or $x \in \{\overline{n(r)} \mid r \in D\}$,*
2. *(P, D) individually defeats all $p \in \mathcal{K}_p \setminus P$, and*
3. *(P, D) contradictory rebuts all (P', D') on some r' where*
   - *each rule in D' is applicable by (P', D'),*
   - *$D' \nsubseteq D$, and*
   - *$P' \subseteq P$ and $D' \subseteq D_U$, where D_U are the rules in $\mathcal{R}_d$ that are not individually defeated by (P, D).*

The first condition states that one cannot derive from this assumption x and y that are contrary or contradictory to each other, and that one cannot derive contraries/contradictories of names of rules in D. This ensures conflict-freeness. Individual defeats and contradictory rebuts among arguments based on a w-stable (P, D) are prevented (otherwise contrary or contradictory x and y are derivable, or contrary/contradictory of a rule name is derivable). The second condition is direct: if violated there is an argument (an ordinary premise) not defeated. The last condition takes care of undercuts and rebuts "outside" the set to be stable: D_U "removes" all possible arguments that are undercut or contrary rebutted (by individual defeat), and if there is an argument still possible, then this argument has P' and D' (exactly) as its defeasible elements, and it has to be the case that (P, D) contradictory rebuts (P', D').

Example 5. *Considering the AT from Example 3, it holds that (P, D) from the example is not a w-stable assumption (a corresponding stable extension containing all arguments based on (P, D) also does not exist): condition 1 is satisfied directly, and condition 2, as well. However, condition 3 is not satisfied: while (P, D) does contradictory rebut (P', D') from the example, it holds that (P, D) does not contradictory rebut (P^*, D^*). A w-stable assumption is $(\{a, b, c\}, \{r_3, r_4\})$.*

We show the correspondence between w-stable assumptions and stable extensions.

Theorem 10. *Let $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$ be a well-formed AT, and $F = (\mathcal{A}, \mathcal{D})$ the corresponding AF to T.*

- *If (P, D) is a w-stable assumption in T, then $\mathcal{E} = \{A \mid A$ based on (P, D) in $T\}$ is a stable extension of F.*
- *If $\mathcal{E}$ is a stable extension of F, then (P, D) is a w-stable assumption of T with $P = \texttt{Prem}_\texttt{d}(\mathcal{E})$ and $D = \texttt{DefRules}(\mathcal{E})$.*

In the theorem, we see that closure under defeasible elements is utilized directly (e.g., via the construction of extensions from an assumption). Well-formedness is most visible in proving the correspondence. For an intuition, the proof of Proposition 8 (and Lemma 7) makes use of transposition: otherwise there might be a conflict derived on an assumption, but not present in a corresponding AF, since, e.g., if A attacks B, $A \prec B$, and the top rule of A is strict. Without transposition such a conflict might not materialize as a defeat in the AT (with transposition one may "chain" transposed rules to obtain a defeat).

We move on to the case of no defeasible rules. Here only individual defeats are present since rebuts require defeasible rules.

Definition 17. *Let $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$ be an AT with $\mathcal{R}_d = \emptyset$. We say that $(P, \emptyset)$, $P \subseteq \mathcal{K}_p$, is s-stable in T if*

- *$(P, \emptyset)$ does not individually defeat any $p \in P$, and*
- *$(P, \emptyset)$ individually defeats all $p \in \mathcal{K}_p \setminus P$.*

If clear from the context, we write (P) instead of $(P, \emptyset)$.

Proposition 11. *Let $T = (\mathcal{L}, \mathcal{R}, n, ^-, \mathcal{K}, \leq)$ be an AT with $\mathcal{R}_d = \emptyset$, and $F = (\mathcal{A}, \mathcal{D})$ the corresponding AF to T.*

- *If (P) is an s-stable assumption in T, then $\mathcal{E} = \{A \mid A$ based on $(P, \emptyset)$ in $T\}$ is a stable extension of F.*
- *If $\mathcal{E}$ is a stable extension of F, then (P) is an s-stable assumption of T with $P = \texttt{Prem}_\texttt{d}(\mathcal{E})$.*

The preceding results directly lead to the result that checking credulous or skeptical justification is possible via inspecting stable assumptions instead of extensions.

Proposition 12. *Given an AT T which is well-formed (or which has empty $\mathcal{R}_d$), and an $x \in \mathcal{L}$, it holds that*

- *x is credulously justified in T under stable semantics iff there is a w-stable (s-stable) assumption (P, D) in T with $x \in Th_T(P, D)$, and*
- *x is skeptically justified in T under stable semantics iff in all w-stable (s-stable) assumptions (P, D) in T we have $x \in Th_T(P, D)$.*

We end this section with complexity results for verifying whether an assumption is s- or w-stable, which form the basis for the membership results of Theorem 3 and Theorem 4.

Algorithm 1 Credulous justification

Require: Well-formed AT T and queried atom $s \in \mathcal{L}$
Ensure: return YES if s is credulously justified in T under stable semantics, NO otherwise
1: $\pi \leftarrow \pi_{1,2}(T) \cup \{\leftarrow \texttt{not derived}(s)\}$
2: **while** π has an answer set M **do**
3: **if** $\pi \cup \pi_{\neg 3}(M)$ has no answer sets **then return** YES
4: **else** $\pi \leftarrow \pi \cup \pi_r(M)$
5: **return** NO

Algorithm 2 Skeptical justification

Require: Well-formed AT T and queried atom $s \in \mathcal{L}$
Ensure: return YES if s is skeptically justified in T under stable semantics, NO otherwise
1: $\pi \leftarrow \pi_{1,2}(T) \cup \{\leftarrow \texttt{derived(s)}\}$
2: **while** π has an answer set M **do**
3: **if** $\pi \cup \pi_{\neg 3}(M)$ has no answer sets **then return** NO
4: **else** $\pi \leftarrow \pi \cup \pi_r(M)$
5: **return** YES

Theorem 13. *Verifying whether an assumption is s-stable is in P, and verifying whether an assumption is w-stable is coNP-complete.*

6 Algorithm for Stable Semantics

Our approach to credulous acceptance under stable semantics is outlined as Algorithm 1. (The ASP encodings used are detailed in the extended version available online.) One ASP solver provides candidate solutions corresponding to assumptions (P, D) that derive the queried atom, satisfy Items 1–2 of Definition 16 and the applicability of every $r \in D$ (Lines 1–2). Another ASP solver checks for counterexamples to the solution candidate (Line 3). A counterexample is an assumption (P', D') such that (P, D) does not contradictory rebut it while the other conditions of Item 3 of Definition 16 hold. If there is no counterexample, the candidate is w-stable and thus the query is credulously justified (Line 3).

We employ two encodings, $\pi_{1,2}(T)$ for candidate generation and $\pi_{\neg 3}$ for counterexample finding, with the following properties: (P, D) derives s and satisfies the first two conditions in T iff there is an answer set M of $\pi_{1,2}(T) \cup \{\leftarrow \texttt{not derived(s)}\}$ with $P \cup D = \{p \in (\mathcal{K}_p \cup \mathcal{R}_d) \mid \mathbf{in}(p) \in M\}$. The $\{\leftarrow \texttt{not derived(s)}\}$ is a constraint ruling out answers where s is not derivable from the guessed assumption set. The answer set M also indicates the premises and rules that (P, D) does not defeat individually. For verifying if the answer set corresponds to a w-stable assumption, it holds that (P', D') is a counterexample in T iff there is an answer set M' of $\pi_{\neg 3}(M)$ with $P' \cup D' = \{p \in (\mathcal{K}_p \cup \mathcal{R}_d) \mid \mathbf{in}(p) \in M'\}$. If a counterexample is found, the current π is refined, via π_r, to exclude each $(P'', D'') \sqsubseteq (P, D)$ (Line 5); the encoding $\pi_r(M)$ enforces that some defeasible element that is not in the candidate in M needs to be in any future candidate. Excluding all subassumptions is valid, since if (P, D) does not contradictory rebut a counterexample (P', D'), then no subassumption of (P, D) can contradictory rebut (P', D') either (follows from Definition 15).

The algorithm for skeptical reasoning (Algorithm 2) follows with relatively minor changes from Algorithm 1. In short, it searches for a counterexample, namely, for a stable extension that does not contain the query s. For this, the constraint in Line 1 is changed to rule out answers where s is derivable from the guessed assumption set. Then the algorithm returns NO if it finds a suitable counterexample and YES otherwise.

7 Experiments

We implemented (available at https://bitbucket.org/coreo-group/aspforaspic/) the ASP-based CEGAR algorithms using the incremental Python interface of Clingo v5.5.1 (Gebser et al. 2019) under default parameters. The experiments were run on 2.60-GHz Intel Xeon E5-2670 8-core 64-GB machines with CentOS 7 under a per-instance 600-second time and 16-GB memory limit.

We generated ATs with $N = 100, 200, \ldots 800$ atoms (i.e., members of $\mathcal{L}$ excluding the names for defeasible rules) as follows, selecting one queried non-premise atom per framework: all atoms aside from axioms and 10% of defeasible rules were assigned a contradictory or asymmetric contrary (each with equal probability); 5% of all atoms are axioms and 20% of atoms are premises. We varied the number of rules deriving each atom (rpa) and the sizes of rule bodies (rs): for each non-premise atom, the number of rules deriving the atom was chosen at random from $[1, 5]$ or $[1, 10]$, as was the number of atoms in the body of each rule body. When the head has a contradictory, a rule deriving it was chosen to be strict with a 10% probability and the atoms in the rule body were selected from the sentences that have a contradictory (closure under transposition requires each atom present in a strict rule to have a contradictory). Closure under transposition was enforced by creating the required additional strict rules. We generated preference relations over premises and defeasible rules for each framework by choosing for both a random permutation $(x_i)_{0<i\leq n}$ of the elements and for each $j < i$ setting x_i to be preferred to x_j with probability 30%. For each N and different combinations of rpa and rs we generated 10 frameworks.

Tables 1 and 2 give the number of timeouts and mean runtimes (in parentheses, with timeouts included as 600 s) of our approach for each N and choice of rs, rpa for credulous and skeptical reasoning, respectively. The empirical hardness of the instances depend on the parameters and reasoning tasks. Skeptical justification is empirically harder than credulous on all of the parameter families tested. The instances with many smaller rules seems the hardest for both reasoning tasks while the instances with fewer and smaller rules are comparatively easier.

In terms of systems for direct runtime comparison, there are few options currently. The Tweety library (Thimm 2017) offers one possible point of comparison. However, Tweety employs a translation to Dung AFs, and as its first step explicitly generates the arguments from a given AT. We observed that already the argument construction step fails (due

	#timeouts (mean runtime (s))					
N	rs=5, rpa=5		rs=5, rpa=10		rs=10, rpa=10	
100	0	(1)	0	(16)	0	(6)
200	0	(8)	4	(321)	0	(49)
300	1	(85)	6	(428)	0	(182)
400	0	(64)	4	(491)	3	(474)
500	4	(316)	10	(600)	10	(600)
600	0	(222)	10	(600)	10	(600)
700	1	(405)	10	(600)	10	(600)
800	3	(556)	10	(600)	10	(600)

Table 1: Timeouts and runtimes on credulous reasoning.

	#timeouts (mean runtime (s))					
N	rs=5, rpa=5		rs=5, rpa=10		rs=10, rpa=10	
100	0	(2)	0	(74)	0	(7)
200	0	(51)	8	(509)	0	(114)
300	1	(190)	10	(600)	1	(279)
400	7	(445)	9	(582)	5	(531)
500	9	(567)	10	(600)	10	(600)
600	8	(524)	10	(600)	10	(600)
700	9	(580)	10	(600)	10	(600)
800	10	(600)	10	(600)	10	(600)

Table 2: Timeouts and runtimes on skeptical reasoning.

to time or memory out) for all but seven of the benchmark instances (five of the succeeding instances had $N = 100$ and two $N = 200$). This further emphasizes the benefits of our AT-level ASP-based approach compared to approaches resorting to AF translation.

8 Related Work

For overviews on computational approaches to structured argumentation see, e.g., surveys by Dvořák and Dunne (2018) and Cerutti et al. (2018). Regarding rephrasing (sets of) arguments—or argument structures—in different forms, in assumption-based argumentation (Bondarenko et al. 1997; Cyras et al. 2018) semantics have been defined in terms of extensions of arguments and on subsets of assumptions, and recently extended with preferences, e.g., in ABA$^+$ (Čyras and Toni 2016). In ASPIC$^+$, preferred subtheories (Brewka 1989) and specific instantiations of ASPIC$^+$ have a semantic correspondence (Modgil and Prakken 2013). There are also connections between repairs of ontological knowledge bases and extensions of AFs instantiated from knowledge bases (Croitoru and Vesic 2013). By studying outcomes of rule-based systems and their instantiated arguments and attacks, certain subparts of the rule-base are connected to semantics of the resulting AF (Amgoud and Besnard 2019). Structured argumentation frameworks, including a fragment of ASPIC$^+$ (Heyninck and Straßer 2021) by making use of properties implying satisfaction of rationality postulates, have been connected to maximal consistent subset reasoning (see, e.g., the recent survey by Arieli, Borg, and Heyninck (2019)). In contrast, we consider ASPIC$^+$ with atomic strict and defeasible rules, contraries and contradictories, weakest-link principle, elitist lifting, and all forms of attacks and defeats defined on these based on (Modgil and Prakken 2018) (undermining, undercut, and rebut, with their contrary and contradictory versions), and present a rephrasing of stability in terms of defeasible elements of the given AT which is designed for computational purposes. We show correspondences for the fragments of (i) AT without defeasible rules and (ii) well-formed ATs. Conditions underlying well-formed ATs have been used earlier for showing satisfaction of properties regarding rationality. We show that these are also viable for computational purposes.

Computational complexity for structured argumentation includes several results, e.g., for assumption-based argumentation (Dimopoulos, Nebel, and Toni 2002; Cyras, Heinrich, and Toni 2021; Karamlou, Cyras, and Toni 2019; Lehtonen, Wallner, and Järvisalo 2021), for deductive argumentation, e.g., (Wooldridge, Dunne, and Parsons 2006; Hirsch and Gorogiannis 2010), and for DeLP, e.g., (Alfano et al. 2021). For ASPIC$^+$, complexity results were obtained for the case without preferences, upon which we build (Lehtonen, Wallner, and Järvisalo 2020). We present novel complexity results for ASPIC$^+$ with preferences.

Cerutti et al. (2018) provide a survey on systems for structured argumentation. Specific to ASPIC$^+$, the systems Tweety (Thimm 2017) and TOAST (Snaith and Reed 2012) implement reasoning for specific tasks. We were unable to obtain the source code for TOAST (Snaith and Reed 2012) from the authors for a further potential direct comparison. Further systems implementing reasoning in contexts drawing inspiration from ASPIC$^+$ (as considered here) include EPR (Visser 2008) and Arg2P (Calegari et al. 2022).

9 Conclusions

As a non-trivial extension of a recently proposed approach to reasoning in ASPIC$^+$, we established both complexity results and formal foundations for an incremental ASP approach to stable conclusions under the weakest-link principle. In terms of complexity, as the main contributions we established completeness for the second level of the polynomial hierarchy for both credulous and skeptical acceptance, thereby witnessing a jump in complexity due to inclusion of preferences. Pertaining to this complexity class, we proposed a counterexample-guided abstraction refinement style approach using incremental ASP solving for the task, scaling up to hundreds of atoms in practice. The approach circumvents a potential exponential blow-up intrinsic to approaches translating ASPIC$^+$ reasoning to abstract argumentation via formally rephrasing the semantics in terms of subsets of defeasible elements.

Our work opens up directions for further work, including extending the approach to further variants of preferential reasoning in ASPIC$^+$ and structured argumentation (Beirlaen et al. 2018; Young, Modgil, and Rodrigues 2016; Dyrkolbotn, Pedersen, and Broersen 2018; Heyninck and Straßer 2019), in terms of both complexity analysis and algorithms, as well as investigating computational properties of further fragments of ASPIC$^+$.

Acknowledgements

This work has been financially supported in part by Academy of Finland (grant 322869), University of Helsinki Doctoral Programme in Computer Science DoCS, and by the Austrian Science Fund (FWF) P35632. The authors wish to thank the Finnish Computing Competence Infrastructure (FCCI) for supporting this project with computational and data storage resources.

References

Alfano, G.; Greco, S.; Parisi, F.; Simari, G. I.; and Simari, G. R. 2021. Incremental computation for structured argumentation over dynamic DeLP knowledge bases. *Artif. Intell.* 300:103553.

Amgoud, L., and Besnard, P. 2019. A formal characterization of the outcomes of rule-based argumentation systems. *Knowledge and Information Systems* 61(1):543–588.

Arieli, O.; Borg, A.; and Heyninck, J. 2019. A review of the relations between logical argumentation and reasoning with maximal consistency. *Ann. Math. Artif. Intell.* 87(3):187–226.

Atkinson, K.; Baroni, P.; Giacomin, M.; Hunter, A.; Prakken, H.; Reed, C.; Simari, G. R.; Thimm, M.; and Villata, S. 2017. Towards artificial argumentation. *AI Magazine* 38(3):25–36.

Baroni, P.; Gabbay, D.; Giacomin, M.; and van der Torre, L., eds. 2018. *Handbook of Formal Argumentation*. College Publications.

Beirlaen, M.; Heyninck, J.; Pardo, P.; and Straßer, C. 2018. Argument strength in formal argumentation. *IfCoLog Journal of Logics and their Applications* 5(3):629–676.

Besnard, P., and Hunter, A. 2008. *Elements of Argumentation*. MIT Press.

Besnard, P., and Hunter, A. 2018. A review of argumentation based on deductive arguments. In Baroni et al. (2018). chapter 9, 437–484.

Bondarenko, A.; Dung, P. M.; Kowalski, R. A.; and Toni, F. 1997. An abstract, argumentation-theoretic approach to default reasoning. *Artif. Intell.* 93:63–101.

Brewka, G.; Delgrande, J. P.; Romero, J.; and Schaub, T. 2015. asprin: Customizing answer set preferences without a headache. In *AAAI*, 1467–1474. AAAI Press.

Brewka, G. 1989. Preferred subtheories: An extended logical framework for default reasoning. In *IJCAI*, 1043–1048. Morgan Kaufmann.

Calegari, R.; Omicini, A.; Pisano, G.; and Sartor, G. 2022. Arg2P: An argumentation framework for explainable intelligent systems. *J. Log. Comput.* 32(2):369–401.

Caminada, M. 2018. Rationality postulates: Applying argumentation theory for non-monotonic reasoning. In Baroni et al. (2018). chapter 15, 771–795.

Cerutti, F.; Gaggl, S. A.; Thimm, M.; and Wallner, J. P. 2018. Foundations of implementations for formal argumentation. In Baroni et al. (2018). chapter 15, 688–767.

Croitoru, M., and Vesic, S. 2013. What can argumentation do for inconsistent ontology query answering? In *SUM*, volume 8078 of *LNCS*, 15–29. Springer.

Čyras, K., and Toni, F. 2016. ABA+: Assumption-based argumentation with preferences. In *KR*, 553–556. AAAI Press.

Cyras, K.; Fan, X.; Schulz, C.; and Toni, F. 2018. Assumption-based argumentation: Disputes, explanations, preferences. In Baroni et al. (2018). chapter 7, 365–408.

Cyras, K.; Heinrich, Q.; and Toni, F. 2021. Computational complexity of flat and generic assumption-based argumentation, with and without probabilities. *Artif. Intell.* 293:103449.

Dimopoulos, Y.; Nebel, B.; and Toni, F. 2002. On the computational complexity of assumption-based argumentation for default reasoning. *Artif. Intell.* 141(1/2):57–78.

Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.* 77(2):321–358.

Dvořák, W., and Dunne, P. E. 2018. Computational problems in formal argumentation and their complexity. In Baroni et al. (2018). chapter 13, 631–688.

Dyrkolbotn, S. K.; Pedersen, T.; and Broersen, J. M. 2018. On elitist lifting and consistency in structured argumentation. *IfCoLog Journal of Logics and their Applications* 5(3):709–746.

Gabbay, D.; Giacomin, M.; Simari, G. R.; and Thimm, M., eds. 2021. *Handbook of Formal Argumentation*, volume 2. College Publications.

García, A. J., and Simari, G. R. 2014. Defeasible logic programming: Delp-servers, contextual queries, and explanations for answers. *Argument Comput.* 5(1):63–88.

García, A. J., and Simari, G. R. 2018. Argumentation based on logic programming. In Baroni et al. (2018). chapter 8, 409–435.

Gebser, M.; Kaminski, R.; Kaufmann, B.; and Schaub, T. 2019. Multi-shot ASP solving with clingo. *Theory Pract. Log. Program.* 19(1):27–82.

Heyninck, J., and Straßer, C. 2019. A fully rational argumentation system for preordered defeasible rules. In *AAMAS*, 1704–1712. IFAAMAS.

Heyninck, J., and Straßer, C. 2021. Rationality and maximal consistent sets for a fragment of ASPIC+ without undercut. *Argument & Compututation* 12(1):3–47.

Hirsch, R., and Gorogiannis, N. 2010. The complexity of the warranted formula problem in propositional argumentation. *J. Log. Comput.* 20(2):481–499.

Karamlou, A.; Cyras, K.; and Toni, F. 2019. Complexity results and algorithms for bipolar argumentation. In *AAMAS*, 1713–1721. IFAAMAS.

Lehtonen, T.; Wallner, J. P.; and Järvisalo, M. 2020. An answer set programming approach to argumentative reasoning in the ASPIC+ framework. In *KR*, 636–646. IJCAI.org.

Lehtonen, T.; Wallner, J. P.; and Järvisalo, M. 2021. Declarative algorithms and complexity results for assumption-based argumentation. *J. Artif. Intell. Res.* 71:265–318.

Li, Z., and Parsons, S. 2015. On argumentation with purely defeasible rules. In *SUM*, volume 9310 of *LNCS*, 330–343. Springer.

Li, Z. 2019. *A Purely Defeasible Argumentation Framework*. Ph.D. Dissertation, City University of New York.

Modgil, S., and Prakken, H. 2013. A general account of argumentation with preferences. *Artif. Intell.* 195:361–397.

Modgil, S., and Prakken, H. 2018. Abstract rule-based argumentation. In Baroni et al. (2018). chapter 6, 287–364.

Niemelä, I. 1999. Logic programs with stable model semantics as a constraint programming paradigm. *Ann. Math. Artif. Intell.* 25(3-4):241–273.

Prakken, H.; Wyner, A. Z.; Bench-Capon, T. J. M.; and Atkinson, K. 2015. A formalization of argumentation schemes for legal case-based reasoning in ASPIC+. *J. Log. Comput.* 25(5):1141–1166.

Prakken, H. 2010. An abstract framework for argumentation with structured arguments. *Argument & Computation* 1(2):93–124.

Prakken, H. 2012. Reconstructing Popov v. Hayashi in a framework for argumentation with structured arguments and Dungean semantics. *Artif. Intell. Law.* 20(1):57–82.

Schraagen, M.; Odekerken, D.; Testerink, B.; and Bex, F. 2018. Argumentation-driven information extraction for online crime reports. In *CIKM*, volume 2482 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Snaith, M., and Reed, C. 2012. TOAST: online ASPIC$^+$ implementation. In *COMMA*, volume 245 of *FAIA*, 509–510. IOS Press.

Thimm, M. 2017. The Tweety library collection for logical aspects of artificial intelligence and knowledge representation. *Künstliche Intelligenz* 31(1):93–97.

Toniolo, A.; Norman, T. J.; Etuk, A.; Cerutti, F.; Ouyang, W. R.; Srivastava, M. B.; Oren, N.; Dropps, T.; Allen, J. A.; and Sullivan, P. 2015. Supporting reasoning with different types of evidence in intelligence analysis. In *AAMAS*, 781–789. ACM.

Visser, W. 2008. Implementation of argument-based practical reasoning. Master's thesis, Utrecht University.

Wooldridge, M. J.; Dunne, P. E.; and Parsons, S. 2006. On the complexity of linking deductive and abstract argument systems. In *AAAI*, 299–304. AAAI Press.

Young, A. P.; Modgil, S.; and Rodrigues, O. 2016. Prioritised default logic as rational argumentation. In *AAMAS*, 626–634. ACM.

Yun, B., and Croitoru, M. 2016. An argumentation workflow for reasoning in ontology based data access. In *COMMA*, volume 287 of *FAIA*, 61–68. IOS Press.