The Computational Complexity of Ideal Semantics

Paul E. Dunne

Dept. of Computer Science, The University of Liverpool, Liverpool, United Kingdom

Abstract

We analyse the computational complexity of the recently proposed *ideal semantics* within both abstract argumentation frameworks (AFs) and assumption-based argumentation frameworks (ABFs). It is shown that while typically less tractable than credulous admissibility semantics, the natural decision problems arising with this extension-based model can, perhaps surprisingly, be decided more efficiently than sceptical preferred semantics. In particular the task of *finding* the unique ideal *extension* is easier than that of *deciding* if a given argument is accepted under the sceptical semantics. We provide efficient algorithmic approaches for the class of *bipartite* argumentation frameworks and, finally, present a number of technical results which offer strong indications that typical problems in ideal argumentation are complete for the class $P_{||}^{\mathcal{C}}$ of languages decidable by polynomial time algorithms allowed to make non-adaptive queries to a $\mathcal{C}$ oracle, where $\mathcal{C}$ is an upper bound on the computational complexity of deciding credulous acceptance: $\mathcal{C} = \text{NP}$ for AFs and logic programming (LP) instantiations of ABFs; $\mathcal{C} = \Sigma_2^p$ for ABFs modelling default theories.

Key words: Computational properties of argumentation; abstract argumentation frameworks assumption-based argumentation; computational complexity;

1 Introduction

Argumentation models have provided a fruitful source of ideas and technologies within both theoretical studies and applications of AI. A recent overview of these contributions may be found in the survey of Bench-Capon and Dunne [3]. Two important models which have received considerable attention over the last ten years are the abstract argumentation frameworks (AFs) of Dung [17] and the related assumption based frameworks (ABFs) of Bondarenko *et al.* [5]. Both approaches provide interpretations for intuitive notions of "collection of justified arguments" as subsets satisfying particular criteria with respect to the underlying framework. In Dung's model the concept of "argument" is regarded as an atomic entity whose principal feature of interest concerns those other arguments with which it is incompatible (such incompatability being described by the so-called *attack relation*).

The formalism adopted in Bondarenko *et al.* develops a rationale capturing incompatability by treating an argument's structure in terms of an assertion which is the outcome of a formal derivation process within some logical theory. In this way two arguments are incompatible if the assertion supported by one is inconsistent with the premises from which the other is derived.

Importing terminology from non-monotonic logic – one of the early and still important application domains of argumentation techniques – collections of justified arguments (in both schemes) are referred to as *extensions*. A variety of different semantics defining the criteria that a set must satisfy in order to constitute an extension of a particular form have been proposed, e.g. grounded, preferred, and stable. With respect to such semantics specific arguments may be viewed as *credulously* accepted (a member of at least one set sanctioned by the semantics) or *sceptically* accepted (a member of *every* set sanctioned by the semantics). The extension based semantics defined by *ideal extensions* were introduced by Dung, Mancarella and Toni [19,20] as an alternative sceptical basis for defining collections of justified arguments in the frameworks promoted by Dung [17] and Bondarenko *et al.* [5].

Our principal concern in this article is in classifying the computational complexity of a number of natural problems related to ideal semantics in both the abstract argumentation frameworks of [17] and the assumption-based approach of [5]: thereby addressing a question raised in the study of ideal semantics presented in [20]. In total our results present a complexity-theoretic analysis of ideal semantics of a similar level of detail to that which has already been achieved for the more widely studied preferred and stable semantics, e.g. in the work of Dimopoulos and Torres [15], Dunne and Bench-Capon [23,24], Dimopoulos, Nebel and Toni [12–14] and Dunne [21].

The problems we consider include both *decision* questions and those related to the *construction* of ideal extensions. Thus,

a. Given an argument x is it accepted under the ideal semantics?
b. Given a *set* of arguments, S
   b1. Is S a *subset* of the (unique) maximal ideal set (subsequently called the ideal *extension*)?, i.e. without, necessarily, being an ideal set itself.
   b2. Is S, itself, an ideal *set*?
   b3. Is S the (unique) ideal *extension*?
c. Is the ideal extension empty?
d. Does the ideal extension coincide with the set of all *sceptically accepted* arguments?
e. Given an AF or ABF, construct its ideal extension.

For the case of AFs and *flat* ABFs with credulous reasoning problems decidable in some complexity class $\mathcal{C}$, we obtain bounds for these problems ranging from $\mathcal{C}$–hard, co$\mathcal{C}$–hard, and Diff$\mathcal{C}$–hard (the last of these being the class of languages expressible as the intersection of a language $L_1 \in \mathcal{C}$ with a language $L_2 \in \text{co}\mathcal{C}$) through to an exact $\text{FP}^{\mathcal{C}}_{||}$–completeness classification for the construction problem defined in (e).

These preliminary results leave a gap between lower (hardness) bounds and upper bounds for a number of the decision questions. We subsequently present strong evidence that problems (a), (b1), (b3) and (c) are not contained in any complexity class strictly below $\text{P}^{\mathcal{C}}_{||}$: specifically that all of these problems are $\text{P}^{\mathcal{C}}_{||}$–hard via *randomized reductions* which are correct with probability approaching 1.

In the remainder of this paper, background definitions are given in Section 2.1 together with formal definitions of the problems introduced in (a)–(e) above. In the first part of the paper, forming Section 2, we concentrate on complexity and algorithmic properties of AFs. [1] In Section 2.2, two technical lemmata are given which characterise properties of ideal sets (Lemma 1) and of *arguments* belonging to the ideal extension (Lemma 2). The complexity of decision questions is considered in Section 2.3 while Section 2.4 provides details of efficient solution approaches for the special case of *bipartite* argumentation frameworks, a class whose properties have previously been studied in [21]. Our main technical result for AFs is presented in Section 2.5 wherein an exact classification for the complexity of *finding* the ideal extension is given. One consequence of this result is that (under the usual complexity-theoretic assumptions) *constructing* the ideal extension of a given argumentation framework is, in general, easier than *deciding* if one of its arguments is sceptically accepted. In Section 2.6 we apply a number of techniques originating from work of [7,8,34] which provide strong evidence that the upper bounds resulting from Section 2.5 are optimal.

We then, in Section 3, turn our attention to complexity questions as they arise for ideal semantics in assumption-based frameworks, thus extending the range of known complexity properties for ABFs from the pioneering studies of Dimopoulos, Nebel and Toni in [12–14]. In Section 3.1 we review the basic elements of ABFs and the formulation of analogues to concepts in AFs given earlier in Section 2.1. We then recall the translations of divers non-monotonic reasoning systems – amongst which we focus on *Logic Programming* (LP) and Default Logic (DL) – as originally presented in [5]. The complexity of decision problems in the ideal semantics for a number of such settings is considered in Section 3.3. Concluding remarks are presented in Section 4.

[1] The results in this first section have been reported in Dunne [22]: the current article includes full proofs of these together with their development to ABF settings.

2 Abstract Argumentation Frameworks

2.1 Backgrounds and review of AF concepts

The following concepts were introduced in Dung [17].

Definition 1 *An* argumentation framework *(*AF*) is a pair* $\mathcal{H} = \langle \mathcal{X}, \mathcal{A} \rangle$*, in which* $\mathcal{X}$ *is a finite set of* arguments *and* $\mathcal{A} \subset \mathcal{X} \times \mathcal{X}$ *is the* attack relationship *for* $\mathcal{H}$*. A pair* $\langle x, y \rangle \in \mathcal{A}$ *is referred to as '*y is attacked by x*' or '*x attacks y*'. For* R*,* S *subsets of arguments in the* AF $\mathcal{H}(\mathcal{X}, \mathcal{A})$*, we say that* $s \in S$ *is* attacked *by* R *– written* $attacks(R, s)$ *– if there is some* $r \in R$ *such that* $\langle r, s \rangle \in \mathcal{A}$*. For subsets* R *and* S *of* $\mathcal{X}$ *we write* $attacks(R, S)$ *if there is some* $s \in S$ *for which* $attacks(R, s)$ *holds;* $x \in \mathcal{X}$ *is* acceptable with respect to S *if for every* $y \in \mathcal{X}$ *that attacks* x *there is some* $z \in S$ *that attacks* y*. A subset,* S*, is* conflict-free *if no argument in* S *is attacked by any other argument in* S*. A conflict-free set* S *is* admissible *if every* $y \in S$ *is acceptable w.r.t* S *and* S *is a* preferred extension *if it is a maximal (with respect to* $\subseteq$*) admissible set. A subset,* S*, is a* stable extension *if* S *is conflict free and every* $y \notin S$ *is attacked by* S*. An* AF*,* $\mathcal{H}$ *is* coherent *if every preferred extension in* $\mathcal{H}$ *is also a stable extension.*

The subset S *is an* ideal set *([19,20]) of* $\mathcal{H}$ *if* S *is admissible* and *a subset of every preferred extension of* $\mathcal{H}$*;* S *is the ideal* extension *if it is the maximal such set.* [2]

The AF*,* $\mathcal{H}$ *is* cohesive *if its maximal ideal extension coincides with the intersection of all preferred extensions of* $\mathcal{H}$*.* [3]

For $S \subseteq \mathcal{X}$*,*

$$S^- =_{\text{def}} \{ p : \exists\, q \in S \text{ such that } \langle p, q \rangle \in \mathcal{A} \}$$
$$S^+ =_{\text{def}} \{ p : \exists\, q \in S \text{ such that } \langle q, p \rangle \in \mathcal{A} \}$$

The various semantics motivate the general decision problems of Table 1 that have been considered in [15,23] w.r.t. AFs and [12–14] in ABFs. Subsequently s denotes one of the (extension) type semantics {ADM,PR,ST,IDL,IE} corresponding to admissible sets, preferred extensions, stable extensions, ideal *sets*, and the ideal *extension*. [4] For a given semantics s and AF, $\mathcal{H}(\mathcal{X}, \mathcal{A})$ we use $\mathcal{E}_s$ to denote the set of

[2] Dung *et al.* [19,20] show that there is a unique maximal ideal set in every AF, ABF.
[3] The term cohesive is introduced here: although the concept is defined in [19,20], no explicit terminology is used to describe such frameworks therein.
[4] These, of course, are far from exhaustive: in addition to the ideal semantics with which the present paper is concerned, our selection is intended to cover the principal cases for which complexity-theoretic issues have been addressed.

all subsets of $\mathcal{X}$ that satisfy the conditions specified by s. Informally, the canonical decision problems are *Verification* (VER), *Credulous Acceptance* (CA) and *Sceptical Acceptance* (SA), so that for example VER_s, refers to the decision problem of verifying that a given set of arguments satisfies the conditions of the semantics s, i.e. that the set is in the collection $\mathcal{E}_s$. The formal definitions of these problems for AFs is presented in Table 1.

Table 1
Decision Problems in AFs

Problem Name	Instance	Question
Verification (VER_s)	$\mathcal{H}(\mathcal{X}, \mathcal{A})$; $S \subseteq \mathcal{X}$	Is $S \in \mathcal{E}_s(\mathcal{H})$?
Credulous Acceptance (CA_s)	$\mathcal{H}(\mathcal{X}, \mathcal{A})$; $x \in \mathcal{X}$	$\exists\, S \in \mathcal{E}_s(\mathcal{H})$ for which $x \in S$?
Sceptical Acceptance (SA_s)	$\mathcal{H}(\mathcal{X}, \mathcal{A})$; $x \in \mathcal{X}$	$\forall\, T \in \mathcal{E}_s(\mathcal{H})$ is $x \in T$?
Existence (EXISTS_s)	$\mathcal{H}(\mathcal{X}, \mathcal{A})$	Is $\mathcal{E}_s(\mathcal{H}) \neq \emptyset$?
Emptiness ($\text{VER}_s^{\emptyset}$)	$\mathcal{H}(\mathcal{X}, \mathcal{A})$	Is $\mathcal{E}_s(\mathcal{H}) = \{\emptyset\}$?

In order to avoid excessive repetition, when we subsequently refer to an argument x as credulously accepted, unless explicitly stated otherwise, this is with respect to admissibility, i.e. CA_{ADM}. Similarly "sceptically accepted" should be understood as with respect to preferred extensions, i.e. SA_{PR}.

Table 2 summarises known complexity results w.r.t these problems.

Table 2
Computational Complexity w.r.t. semantics s

s	VER_s	CA_s	SA_s	EXIST_s	$\text{VER}_s^{\emptyset}$
ADM	P ([17])	NP-c ([15])	Trivial ([17])	Trivial ([17])	coNP-c ([15])
PR	coNP-c ([15])	NP-c ([15])	Π_2^p-c ([23])	Trivial ([17])	coNP-c ([15])
ST	P ([17])	NP-c ([15])	coNP-c / D^p-c	NP-c ([15])	Trivial

Remarks

(1) For a complexity class $\mathcal{C}$, $\mathcal{C} - c$ denotes $\mathcal{C}$-completeness.
(2) Cases which are described as "trivial" are either those for which the property in question always holds such as existence of preferred extensions, or for which it never holds (or holds only in extreme cases), e.g. the set of stable extensions for $\langle \mathcal{X}, \mathcal{A} \rangle$ is $\{\emptyset\}$ if and only if $\mathcal{X} = \emptyset$.
(3) The two distinct classifications for SA_{ST} arise from the two possible interpretations of sceptical acceptance w.r.t. stable extensions for AFs without any, i.e. if one regards $x \in \cap_{S \in \text{ST}(\mathcal{H})} S$ as holding even when $\mathcal{E}_{\text{ST}}(\mathcal{H}) = \emptyset$ then the decision problem is coNP–complete ([15] via commentary of [23, p. 189]). If,

however, one requires $\mathcal{H}$ to have *at least one* stable extension as a precondition for x to be sceptically accepted the decision problem becomes D^p-complete. [5]

We consider a number of decision problems relating to properties of ideal extensions in argumentation frameworks as described in Table 3.

Table 3
Decision questions for Ideal Semantics

	Problem Name	Instance	Question
A.	VER_{IDL}	$\mathcal{H}(\mathcal{X}, \mathcal{A})$; $S \subseteq \mathcal{X}$	Is S an ideal *set*?
B.	VER_{IE}	$\mathcal{H}(\mathcal{X}, \mathcal{A})$; $S \subseteq \mathcal{X}$	Is S the ideal *extension*?
C.	$\text{VER}_{\text{IE}}^{\emptyset}$	$\mathcal{H}(\mathcal{X}, \mathcal{A})$	Is the ideal extension empty?
D.	CA_{IE}	$\mathcal{H}(\mathcal{X}, \mathcal{A})$; $x \in \mathcal{X}$	Is x in the ideal extension?
E.	CS	$\mathcal{H}(\mathcal{X}, \mathcal{A})$	Is $\mathcal{H}(\mathcal{X}, \mathcal{A})$ *cohesive*?

We also consider *search* (so-called *function problems*) where the aim is not simply to verify that a given set has a specific property but to *construct* an example. In particular we examine the function problem FIE in which, given an AF $\mathcal{H}(\mathcal{X}, \mathcal{A})$, it is required to return the ideal extension of $\mathcal{H}$.

We recall that D^p is the class of decision problems, L, whose positive instances are characterised as those belonging to $L_1 \cap L_2$ where $L_1 \in \text{NP}$ and $L_2 \in \text{coNP}$. The problem SAT-UNSAT whose instances are pairs of 3-CNF formulae $\langle \Phi_1, \Phi_2 \rangle$ accepted if Φ_1 is satisfiable *and* Φ_2 is unsatisfiable has been shown to be complete for this class [32, p. 413]. This class can be interpreted as those decision problems which may be solved by a (deterministic) polynomial time algorithm which is allowed to make at most two calls upon an NP *oracle*. More generally, the complexity classes P^{NP} and FP^{NP} (sometimes denoted Δ_2^p and $\text{F}\Delta_2^p$) consist of those decision problems (respectively function problems) that can be solved by a (deterministic) polynomial time algorithm provided with access to an NP *oracle* (calls upon which take a single step so that only polynomially many invocations of this oracle are allowed). [6] An important (presumed) subset of P^{NP} and its associated function class is defined by distinguishing whether oracle calls are *adaptive* – i.e. the exact formulation of the next oracle query may be dependent on the answers received to previous questions – or whether such queries are *non-adaptive*, i.e. the form of the questions to be put to the oracle is predetermined allowing all of these to be performed in parallel. The latter class has been considered in Wagner [36,37], Jenner and Toran [27]. Under the standard complexity-theoretic assumptions, it is

[5] Although not directly relevant to the topic of the current article, since we are unaware of any previous published proof of this claim we present such in Appendix B.
[6] We refer the reader to e.g. [32, pp. 415–423] for further background.

conjectured that,

$$\text{P} \subset \left\{ \begin{array}{c} \text{NP} \\ \\ \text{co-NP} \end{array} \right\} \subset \text{D}^p \subset \text{P}_{||}^{\text{NP}} \subset \text{P}^{\text{NP}} \subset \left\{ \begin{array}{c} \Sigma_2^p \\ \\ \Pi_2^p \end{array} \right\}$$

We prove the following complexity classifications.

a. VER_{IDL} is coNP–complete.
b. CA_{IDL} is coNP–hard via $\leq_m^p$-reducibility.
c. $\text{VER}_{\text{IE}}^{\emptyset}$ is NP–hard via $\leq_m^p$-reducibility.
d. VER_{IE} is D^p–hard via $\leq_m^p$-reducibility.
e. CS is Σ_2^p–complete.
f. FIE is $\text{FP}_{||}^{\text{NP}}$–complete.
g. Problems (A)–(E) of Table 3 and FIE are polynomial time solvable for *bipartite* frameworks.
h. Problems (B)–(E) of Table 3 are $\text{P}_{||}^{\text{NP}}$–complete via *randomized* reductions.

2.2 Characteristic properties of ideal sets

The upper bound proofs exploit a characterisation of ideal sets in terms of credulous acceptability presented in Lemma 1. Lemma 2 gives a necessary and sufficient condition for a given *argument* to be a member of the ideal extension.

Lemma 1 *Let $\mathcal{H}(\mathcal{X}, \mathcal{A})$ be a* AF *and $S \subseteq \mathcal{X}$. Then S defines an ideal set of $\mathcal{H}$ if and only if both of the conditions below are satisfied:*

I1. $S \in \mathcal{E}_{\text{ADM}}(\mathcal{H})$*, i.e. S is an admissible set of arguments in $\mathcal{H}$.*
I2. For every argument $p \in S^-$, there is no *admissible set of $\mathcal{H}$ that contains p, i.e.* $\forall\, p \in S^- \;\; \neg\text{CA}_{\text{ADM}}(\mathcal{H}, p)$.

Proof:
($\Rightarrow$) Suppose that $S \subseteq \mathcal{X}$ is an ideal set of $\mathcal{H}$. It is immediate from the definition of ideal set that S is admissible so (I1) holds. Furthermore, were it the case that (I2) failed to hold, then there would be some admissible set, T, of $\mathcal{H}$ for which $T \cap S^- \neq \emptyset$, and thus some preferred extension, R, with $R \cap S^- \neq \emptyset$. For this preferred extension, however, one cannot have $S \subseteq R$, thereby contradicting the assumption that S is an ideal set.

($\Leftarrow$) Let S be an admissible set for which no argument in S^- is credulously accepted. We show that S is a subset of every preferred extension of $\mathcal{H}$ and, thus, an ideal set. Consider any preferred extension, R of $\mathcal{H}$. We first claim that the set $S \cup R$

must be conflict free: the only way in which this could fail to be true is if there are arguments $s \in S$ and $r \in R$ such that $\langle r, s\rangle \in \mathcal{A}$ or $\langle s, r\rangle \in \mathcal{A}$. In the former case $r \in S^-$ which contradicts the assumption that no argument in S^- is credulously accepted. In the latter case, since R is a preferred extension, there must be some argument $q \in R$ that defends r against the attack by s, i.e. $\langle q, s\rangle \in \mathcal{A}$ and $q \in R$: again this gives $q \in S^-$ and would contradict the assumption that no argument in S^- were credulously accepted. The set $S \cup R$ is thus conflict-free. It is furthermore, admissible: any argument in $\mathcal{X}$ attacking $S \cup R$ either attacks an argument in S (and so is counterattacked by an argument in S since S is admissible) or attacks an argument in R (and, again, is counterattacked by an argument in R since R is a preferred extension). The set R, however, is a *maximal* admissible set and thus $S \cup R = R$, i.e. $S \subseteq R$ as required. □

Lemma 2 *Let $\mathcal{H}(\mathcal{X}, \mathcal{A})$ be an* AF *and let $\mathcal{M} \subseteq \mathcal{X}$ be its ideal extension. Then $x \in \mathcal{X}$ is a member of $\mathcal{M}$ if and only if both of the conditions below are satisfied:*

M1. No attacker of x is credulously accepted, i.e. $\forall\, y \in \{x\}^-\ \neg \mathrm{CA}_{\mathrm{ADM}}(\mathcal{H}, y)$.
M2. For each attacker y of x, at least one attacker z of y is in $\mathcal{M}$, i.e. $\forall\, y \in \{x\}^- : \{y\}^- \cap \mathcal{M} \neq \emptyset$.

Proof:
($\Rightarrow$) Suppose that $x \in \mathcal{M}$, the ideal extension of $\mathcal{H}$. Since $\mathcal{M}$ is an ideal set, from Lemma 1, no attacker of $\mathcal{M}$ can be credulously accepted and, in particular, no attacker of x can be credulously accepted. Any such attack, $y \in \{x\}^-$, must, however, be counterattacked by at least one argument of $\mathcal{M}$ since $\mathcal{M}$ is admissible. The only available counterattacks on $y \in \{x\}^-$ are those in the set $\{y\}^-$, hence $\{y\}^- \cap \mathcal{M} \neq \emptyset$.

($\Leftarrow$) Suppose that $x \in \mathcal{X}$ is such that no attacker of x is credulously accepted and that for each such attacker, y, some counterattacker, z of y is in $\mathcal{M}$. We show that $\mathcal{M} \cup \{x\}$ forms an ideal set, from which it follows that $x \in \mathcal{M}$ since $\mathcal{M}$ is maximal. Consider the set $\mathcal{M} \cup \{x\}$. To see that $\mathcal{M} \cup \{x\}$ is admissible, first observe that it is conflict-free: if, for $p \in \mathcal{M}$, we have $\langle p, x\rangle \in \mathcal{A}$ then p is credulously accepted (by the admissibility of $\mathcal{M}$) contradicting the property (M1); similarly if $\langle x, p\rangle \in \mathcal{A}$ for some $p \in \mathcal{M}$ then as $\mathcal{M}$ is admissible we find $q \in \mathcal{M}$ with $\langle q, x\rangle \in \mathcal{A}$ resulting in a similar contradiction. Thus $\mathcal{M} \cup \{x\}$ is conflict-free. This set, however, also defends itself against any attack. For consider any argument y that attacks $\mathcal{M} \cup \{x\}$: either y attacks $\mathcal{M}$ and so is counterattacked by some $z \in \mathcal{M}$; alternatively y attacks x. Now since $y \in \{x\}^-$ we can identify $z \in \{y\}^- \cap \mathcal{M}$ which counterattacks y. In summary, $\mathcal{M} \cup \{x\}$ is admissible. Since $\mathcal{M}$ is an ideal set we know from Lemma 1 that no attacker of $\mathcal{M}$ is credulously accepted. From the properties assumed of x, it is also the case that no attacker of x is credulously accepted. It follows that $\mathcal{M} \cup \{x\}$ is an admissible set none of whose attackers is credulously accepted, i.e. from Lemma 1, $\mathcal{M} \cup \{x\}$ is an ideal set. The set $\mathcal{M}$ is, however, already *maximal* so that $\mathcal{M} \cup \{x\} = \mathcal{M}$, i.e. $x \in \mathcal{M}$ as required. □

2.3 Preliminary complexity results on ideal semantics in AF*s*

In this section we derive some initial results on the complexity of the decision problems for ideal semantics described in Table 3. In a number of cases these leave a gap between upper and lower bounds, however, the constructions form the basis for the analyses of Section 2.6 in which approaches to obtaining exact classifications are developed.

Theorem 2 VER$_{\mathrm{IDL}}$ *is co*NP*–complete.*

Proof: Given an instance $\langle \mathcal{H}(\mathcal{X}, \mathcal{A}), S \rangle$ of VER$_{\mathrm{IDL}}$ we can decide if this should be accepted by checking

$$\mathrm{VER}_{\mathrm{ADM}}(\mathcal{H}, S) \ \wedge \ \bigwedge_{q \in S^-} \neg \mathrm{CA}_{\mathrm{ADM}}(\mathcal{H}, q)$$

Correctness follows from Lemma 1 and since CA$_{\mathrm{ADM}}(\mathcal{H}, q)$ is decidable in NP, its complement is decidable by a coNP algorithm. [7]

To prove VER$_{\mathrm{IDL}}$ is coNP–hard we reduce from CNF-UNSAT (without loss of generality, restricted to instances which are 3-CNF). Given a 3-CNF formula

$$\Phi(z_1, \ldots, z_n) \ = \ \bigwedge_{i=1}^{m} C_i \ = \ \bigwedge_{i=1}^{m} (z_{i,1} \vee z_{i,2} \vee z_{i,3})$$

as an instance of UNSAT we form an instance $\langle \mathcal{F}_\Phi, S \rangle$ of VER$_{\mathrm{IDL}}$ as follows. First construct the AF $\mathcal{H}_\Phi$ (described in Appendix A.1) from the CNF Φ. In this, via [15, Thm. 5.1, p. 227], the argument Φ is credulously accepted if and only if the CNF, $\Phi(Z_n)$ is satisfiable, i.e. Φ is *not* credulously accepted if and only if $\Phi(Z_n)$ is unsatisfiable. The AF, $\mathcal{F}_\Phi$, is formed from $\mathcal{H}_\Phi$ by adding an argument Ψ together with attacks

$$\{ \ \langle \Psi, z_i \rangle, \ \langle \Psi, \neg z_i \rangle \ : \ 1 \le i \le n \} \ \cup \ \{ \ \langle \Phi, \Psi \rangle, \ \langle \Psi, \Phi \rangle \}$$

The instance of VER$_{\mathrm{IDL}}$ is completed by setting $S = \{\Psi\}$.

We claim $\langle \mathcal{F}_\Phi, \{\Psi\} \rangle$ is accepted as an instance of VER$_{\mathrm{IDL}}$ if and only if Φ is unsatisfiable.

[7] The form $\bigwedge_{q \in S^-} \neg \mathrm{CA}_{\mathrm{ADM}}(\mathcal{H}, q)$ is equivalent to $\forall\, T \ \mathrm{VER}_{\mathrm{ADM}}(\mathcal{H}, T) \Rightarrow (T \cap S^- = \emptyset)$ so that it is not necessary to use $|S|$ distinct coNP tests.

First observe that $\{\Psi\}$ *is* an admissible set: its only attacker is the argument Φ which Ψ counterattacks. Thus, via Lemma 1, in order to complete the proof it suffices to observe that

$$\neg\textsc{ca}_{\text{ADM}}(\mathcal{F}_{\Phi}, \Phi) \;\Leftrightarrow\; \neg\textsc{ca}_{\text{ADM}}(\mathcal{H}_{\Phi}, \Phi) \;\Leftrightarrow\; \textsc{unsat}(\Phi)$$

$\square$

Corollary 1 $\textsc{ca}_{\text{IDL}}$ *is co*NP*–hard.*

Proof: It suffices to note that $\langle \mathcal{F}_{\Phi}, \Psi\rangle$ with $\mathcal{F}_{\Phi}$ the AF defined in Thm. 2, defines a positive instance of $\textsc{ca}_{\text{IDL}}$ if and only if $\Phi(Z_n)$ is unsatisfiable. $\square$

Corollary 2 $\textsc{ver}^{\emptyset}_{\text{IE}}$ *is* NP*–hard*

Proof: The AF $\mathcal{F}_{\Phi}$ defined in Thm. 2 has an empty ideal extension if and only if $\Phi(Z_n)$ is satisfiable. $\square$

Theorem 3 $\textsc{ver}_{\text{IE}}$ *is* D^p*–hard.*

Proof: Given $\langle \Phi_1(Z_n), \Phi_2(Y_n)\rangle$ as an instance of SAT-UNSAT, form $\mathcal{F}_{\langle\Phi_1,\Phi_2\rangle}$ as the AF containing the frameworks $\mathcal{F}_{\Phi_1}$ and $\mathcal{F}_{\Phi_2}$ described in the proof of Thm. 2 where we use Ψ_1 and Ψ_2 to denote the arguments added to $\mathcal{H}_{\Phi_1}$ and $\mathcal{H}_{\Phi_2}$ respectively. The instance $\langle \mathcal{F}_{\langle\Phi_1,\Phi_2\rangle}, \{\Psi_2\}\rangle$ of $\textsc{ver}_{\text{IE}}$ is accepted if and only if $\langle \Phi_1, \Phi_2\rangle$ is accepted as an instance of SAT-UNSAT. To see this note that there are exactly four possibilities for the ideal extension, $\mathcal{M}$, of $\mathcal{F}_{\langle\Phi_1,\Phi_2\rangle}$: $\mathcal{M} = \emptyset$ (both Φ_1 and Φ_2 are satisfiable); $\mathcal{M} = \{\Psi_1, \Psi_2\}$ (neither formula is satisfiable); $\mathcal{M} = \{\Psi_1\}$ (Φ_1 is unsatisfiable and Φ_2 is satisfiable; $\mathcal{M} = \{\Psi_2\}$ (Φ_1 is satisfiable and Φ_2 is unsatisfiable). Only the final case corresponds with the set given in the constructed instance. $\square$

Theorem 4 CS *is* Σ_2^p*–complete.*

Proof: For membership in Σ_2^p, $\mathcal{H}(\mathcal{X}, \mathcal{A})$ is a cohesive system if and only if

$$\textsc{ver}_{\text{ADM}}\left(\mathcal{H}(\mathcal{X},\mathcal{A}),\ \bigcap_{S\in\mathcal{E}_{\text{PR}}} S\right)$$

which can be tested by checking

$$\exists\, S\ \textsc{ver}_{\text{IDL}}(\mathcal{H}, S) \ \wedge \bigwedge_{x\in\mathcal{X}\setminus S} \neg\textsc{sa}_{\text{PR}}(\mathcal{H}, x) \tag{1}$$

That is, there is a subset (S) of $\mathcal{X}$ which defines an ideal set of $\mathcal{H}$ and for which no argument outside S is in every preferred extension. [8] From Thm. 2, VER_{IDL} is in coNP; in addition since $\text{SA}_{\text{PR}} \in \Pi_2^p$ its complement $\neg\text{SA}_{\text{PR}}$ is in Σ_2^p hence (1) gives a Σ_2^p test for CS. [9]

For Σ_2^p–hardness, we use a reduction from QSAT_2^Σ, instances of which comprise a CNF formula $\Phi(Y_n, Z_n)$ over disjoint sets of propositional variables. Such instances being accepted if and only if there is an instantiation α of X_n for which every instantiation, β, of Y_n fails to satisfy Φ, i.e. $\exists\, \alpha\, \forall\, \beta\;\; \neg\Phi(\alpha, \beta)$.

We use the reduction presented in Dunne and Bench-Capon [23] from the complementary problem – QSAT_2^Π, details of which are presented in Appendix A.2.

Given an instance $\Phi(Y_n, Z_n)$ of QSAT_2^Σ, consider the AF $\mathcal{G}_\Phi$, defined from this as described in Appendix A.2. Noting that the ideal extension of $\mathcal{G}_\Phi$ is the empty set it suffices to show the intersection of all preferred extensions is empty if and only if the CNF from which it is defined is accepted as an instance of QSAT_2^Σ.

Notice that every S containing *at most one* element from each of the pairs $\{y_i, \neg y_i\}$ is admissible. Furthermore, if S_α is such a set containing *exactly one* representative from each of these pairs (corresponding to an instantiation α of Y_n) then S_α is a preferred extension if and only if there is no instantiation β of Z_n under which $\Phi(\alpha, \beta) = \top$. In summary, the preferred extensions of $\mathcal{G}_\Phi$ have the form $S_\alpha \cup T_\alpha$ with

$$T_\alpha \;=\; \begin{cases} \emptyset & \text{if } \;\forall\beta\;\Phi(\alpha, \beta) = \bot \\ \{\Phi\} \;\cup\; R_\beta & \text{if } \;\exists\beta\;\Phi(\alpha, \beta) = \top \end{cases}$$

where R_β denotes the subset of $\{\, z_i, \neg z_i \;: 1 \le i \le n\}$ induced by the instantiation $\beta = \langle b_1, \ldots, b_n \rangle$ of Z_n, i.e. $z_i \in R_\beta \;\Leftrightarrow\; b_i = \top$.

The set $\{\Phi\}$ is not admissible and the argument Φ occurs in *every* preferred extension if and only if for every instantiation α of Y_n there is some instantiation, β, of Z_n, for which $\Phi(\alpha, \beta) = \top$. In other words, the intersection of all preferred sets is *non-empty* if and only if $\Phi(Y_n, Z_n)$ is *not* accepted as an instance of QSAT_2^Σ. We deduce that CS is Σ_2^p-complete in consequence. □

Corollary 3 *The property of* coherence *is neither necessary nor sufficient for an* AF *to be cohesive.*

[8] Notice that any such S would form the ideal *extension* of $\mathcal{H}$.

[9] Note that we could extrapolate the existence part of the $\exists\forall$ structure implicit in $\neg\text{SA}(\mathcal{H}, x)$ by "guessing" a set U_x to associate with each $x \notin S$ in the scope of the opening existential quantifier. With this approach, the test $\neg\text{SA}_{\text{PR}}(\mathcal{H}, x)$ is replaced by verifying that U_x is a preferred extension of $\mathcal{H}$ (coNP) and that $x \notin U_x$.

Proof: From [23, Corollary 18], the AF $\mathcal{G}_\Phi$ used in the proof of Thm. 4 is coherent if and only the argument Φ is sceptically accepted. Recalling that $\mathcal{G}_\Phi$ has an empty ideal extension, regardless of whether Φ is accepted as an instance of QSAT_2^Π, it follows that $\mathcal{G}_\Phi$ is coherent if and only if it is *not* cohesive. □

2.4 Frameworks with Efficient Algorithms

We recall that *bipartite* AFs, $\langle \mathcal{X}, \mathcal{A} \rangle$ are those for which $\mathcal{X}$ may be partitioned into two sets – $\mathcal{Y}$ and $\mathcal{Z}$ – both of which are conflict-free in $\langle \mathcal{X}, \mathcal{A} \rangle$. We use the notation $\mathcal{B}(\mathcal{Y}, \mathcal{Z}, \mathcal{A})$ for such frameworks.

Theorem 5 *If* $\mathcal{B}(\mathcal{Y}, \mathcal{Z}, \mathcal{A})$ *is a bipartite* AF *then* $\mathcal{B}(\mathcal{Y}, \mathcal{Z}, \mathcal{A})$ *is* cohesive.

Proof: Consider any bipartite AF, $\mathcal{B}(\mathcal{Y}, \mathcal{Z}, \mathcal{A})$, and let

$$\mathcal{M} = \bigcap_{S \subseteq \mathcal{Y} \cup \mathcal{Z} \,:\, S \in \mathcal{E}_{\text{PR}}(\mathcal{B}), \text{ i.e. } S \text{ is a preferred extension of } \mathcal{B}} S$$

Thus $\mathcal{M}$ is the set of sceptically accepted arguments of $\mathcal{B}$. We can define a partition of each of the sets $\mathcal{Y}$ and $\mathcal{Z}$ into three subsets as follows:

$$\begin{aligned}
\mathcal{Y}_{\text{SA}} &= \mathcal{Y} \cap \mathcal{M} \\
\mathcal{Y}_{\text{CA}} &= \{ y \in \mathcal{Y} \; : \; \text{CA}_{\text{ADM}}(\mathcal{B}, y) \} \setminus \mathcal{M} \\
\mathcal{Y}_{\text{OUT}} &= \{ y \in \mathcal{Y} \; : \; \neg\text{CA}_{\text{ADM}}(\mathcal{B}, y) \} \\
\mathcal{Z}_{\text{SA}} &= \mathcal{Z} \cap \mathcal{M} \\
\mathcal{Z}_{\text{CA}} &= \{ z \in \mathcal{Z} \; : \; \text{CA}_{\text{ADM}}(\mathcal{B}, z) \} \setminus \mathcal{M} \\
\mathcal{Z}_{\text{OUT}} &= \{ z \in \mathcal{Z} \; : \; \neg\text{CA}_{\text{ADM}}(\mathcal{B}, z) \}
\end{aligned}$$

Notice that since every argument in $\mathcal{M}$ is sceptically accepted and from the fact that $\mathcal{B}$ is coherent – so that every preferred extension of $\mathcal{B}$ is also a stable extension – we must have

$$\begin{aligned}
\mathcal{Y}_{\text{SA}}^- \subseteq \mathcal{Z}_{\text{OUT}} \;\; &; \;\; \mathcal{Y}_{\text{SA}}^+ \subseteq \mathcal{Z}_{\text{OUT}} \\
\mathcal{Z}_{\text{SA}}^- \subseteq \mathcal{Y}_{\text{OUT}} \;\; &; \;\; \mathcal{Z}_{\text{SA}}^+ \subseteq \mathcal{Y}_{\text{OUT}}
\end{aligned}$$

In addition,

$$\forall\, y \in \mathcal{Y}_{\mathrm{CA}}\; \exists\, z \in \mathcal{Z}_{\mathrm{CA}}\;\; \langle z, y\rangle \in \mathcal{A}$$
$$\forall\, z \in \mathcal{Z}_{\mathrm{CA}}\; \exists\, y \in \mathcal{Y}_{\mathrm{CA}}\;\; \langle y, z\rangle \in \mathcal{A}$$

To see this, suppose without loss of generality, that $\mathcal{Z}_{\mathrm{CA}}$ does not attack $y \in \mathcal{Y}_{\mathrm{CA}}$: then since the set $\mathcal{Z}_{\mathrm{SA}} \cup \mathcal{Z}_{\mathrm{CA}}$ does not attack y the only arguments which could attack y are those in the set $\mathcal{Z}_{\mathrm{OUT}}$, i.e. no attacker of y is credulously accepted. Now, since $\mathcal{B}$ is coherent, such a situation would mean that y was sceptically accepted, thereby contradicting the maximality of $\mathcal{Y}_{\mathrm{SA}}$.

To complete the proof it suffices to argue that the set $\mathcal{M}$ is *admissible*. Notice that both of the sets $\mathcal{M} \cup \mathcal{Y}_{\mathrm{CA}}$ and $\mathcal{M} \cup \mathcal{Z}_{\mathrm{CA}}$ are *preferred extensions* of $\mathcal{B}$: the set $\mathcal{Y}_{\mathrm{SA}} \cup \mathcal{Y}_{\mathrm{CA}}$ is the maximal subset of $\mathcal{Y}$ which is admissible, however, any preferred extension containing $\mathcal{Y}_{\mathrm{SA}} \cup \mathcal{Y}_{\mathrm{CA}}$ must have $\mathcal{Z}_{\mathrm{SA}}$ as a subset, i.e. there is a preferred extension, $P_{\mathcal{Y}}$, of which $\mathcal{M} \cup \mathcal{Y}_{\mathrm{CA}}$ is a subset. The set $\mathcal{M} \cup \mathcal{Y}_{\mathrm{CA}}$ cannot be a *strict* subset of $P_{\mathcal{Y}}$ otherwise we would have $P_{\mathcal{Y}} \cap \mathcal{Z}_{\mathrm{CA}} \neq \emptyset$ and $P_{\mathcal{Y}}$ is not conflict-free, or $P_{\mathcal{Y}} \cap (\mathcal{Y}_{\mathrm{OUT}} \cup \mathcal{Z}_{\mathrm{OUT}}) \neq \emptyset$ contradicting the property that no argument in $\mathcal{Y}_{\mathrm{OUT}} \cup \mathcal{Z}_{\mathrm{OUT}}$ is credulously accepted. In summary, we have identified two preferred extensions $\mathcal{M} \cup \mathcal{Y}_{\mathrm{CA}}$ and $\mathcal{M} \cup \mathcal{Z}_{\mathrm{CA}}$ of $\mathcal{B}$. That $\mathcal{M}$ is admissible will follow from the fact that both $\mathcal{Y}_{\mathrm{SA}}$ and $\mathcal{Z}_{\mathrm{SA}}$ are admissible. Suppose $\mathcal{Y}_{\mathrm{SA}}$ is not admissible: this could only happen if there were an argument $z \in \mathcal{Z}_{\mathrm{OUT}}$ which attacked $\mathcal{Y}_{\mathrm{SA}}$ and for which $z \notin \mathcal{Y}^{+}_{\mathrm{SA}}$. In this case, however, the same attack would be undefended in the set $\mathcal{M} \cup \mathcal{Z}_{\mathrm{CA}}$ contradicting the fact this latter set is a preferred extension. By an identical argument we see that $\mathcal{Z}_{\mathrm{SA}}$ is admissible and now, by a similar argument to that of Lemma 1 it follows that $\mathcal{M} = \mathcal{Y}_{\mathrm{SA}} \cup \mathcal{Z}_{\mathrm{SA}}$ is admissible.
□

Corollary 4 *Let $\mathcal{B}(\mathcal{Y}, \mathcal{Z}, \mathcal{A})$ be a bipartite* AF. *The ideal extension of $\mathcal{B}(\mathcal{Y}, \mathcal{Z}, \mathcal{A})$ may be constructed in polynomial time.*

Proof: From Thm. 5 the ideal extension of $\mathcal{B}$ corresponds with the set of all sceptically accepted arguments of $\mathcal{B}$. Applying the methods described in [21, Thm. 6] this set can be identified in polynomial time. □

We note that as a consequence of Thm. 5 and Corollary 4 in the case of bipartite AFs, the decision problems $\mathrm{VER}_{\mathrm{IDL}}$, $\mathrm{VER}^{\emptyset}_{\mathrm{IE}}$, $\mathrm{VER}_{\mathrm{IE}}$ and $\mathrm{CA}_{\mathrm{IDL}}$ are all in P and CS is trivial.

A second class of restricted frameworks for which efficient decision methods exist are *bounded treewidth* AFs as described in [21, Sect. 7].

Theorem 6 *Let $\mathcal{H}(\mathcal{X}, \mathcal{A})$ be an* AF *and $tw(\mathcal{H})$ denote the treedwidth of $\mathcal{H}$.*

a. CS *is* fixed-parameter tractable *([16]) with respect to $tw(\mathcal{H})$.*
b. The special case, $\text{VER}^{\emptyset}_{\text{IE}}$ *of deciding whether the ideal extension is empty, is fixed-parameter tractable with respect to $tw(\mathcal{H})$.*

Proof: Both parts follow from Courcelle's Theorem [10,11,2] by defining sentences of *monadic second-order logic* [10] that describe the properties.

For CS the expression verifying membership of CS in Σ_2^p in the proof of Thm. 4 leads to such a sentence after expanding the predicates $\text{VER}_{\text{IE}}(\mathcal{H}, S)$ and $\neg\text{SA}_{\text{PR}}(\mathcal{H}, x)$.

An MSOL sentence describing $\text{VER}^{\emptyset}_{\text{IE}}$ is given by

$$\begin{array}{lll} \forall\, S\, \exists\, T\;\; (S = \emptyset) & \vee & \neg\text{VER}_{\text{ADM}}(\mathcal{H}, S) \\ & \vee & \text{VER}_{\text{ADM}}(\mathcal{H}, T) \wedge (\exists t \exists s\; (s \in S)(t \in T)\langle t, s\rangle \in \mathcal{A}) \end{array}$$

□

2.5 Finding the Ideal Extension

The analysis of properties of the ideal extension in bipartite frameworks and the polynomial time method for constructing this suggest an approach to constructing the ideal extension in arbitrary AFs. In this section we show that if the set of credulously accepted within $\mathcal{X}$ has already been identified then this suffices efficiently to build the ideal extension. An immediate corollary is that decision questions concerning the ideal extension are in the class $\text{P}^{\text{NP}}_{||}$. We shall, subsequently, in Section 2.6 argue that this upper bound is optimal.

Theorem 7 FIE *is* $\text{FP}^{\text{NP}}_{||}$*-complete.*

Proof: We first present the argument that FIE is $\text{FP}^{\text{NP}}_{||}$–hard.

The following function problem is easily seen to be complete for $\text{FP}^{\text{NP}}_{||}$.
Sat Collection SC
Instance: $\Xi = \langle \varphi_1, \varphi_2, \ldots, \varphi_r\rangle$ a collection of 3-CNF formulae.
Problem: Compute the r-bit value $\chi(\Xi) = c_1c_2c_3\cdots c_r \in [0, 2^r - 1]$ in which $c_j = 1$ if and only if φ_j is satisfiable.

Given an instance, $\Xi = \langle \varphi_1, \varphi_2, \ldots, \varphi_r\rangle$ of SC form the AF consisting of the r instantiations of $\mathcal{F}_{\varphi_i}$. Letting $\mathcal{M}$ denote the set of arguments forming the ideal

[10] Simplified descriptions of this may be found in [2] or [21, Sect. 7].

extension of this framework, from Thm. 3, it follows that $\mathcal{M} \subseteq \{\Psi_1, \Psi_2, \ldots, \Psi_r\}$ (where Ψ_i is the argument added to $\mathcal{H}_{\varphi_i}$). In addition, $\Psi_i \notin \mathcal{M}$ if and only if φ_i is satisfiable. It follows that $\chi(\Xi)$ can be computed directly given $\mathcal{M}$, and thus FIE is $\text{FP}^{\text{NP}}_{||}$–hard.

To see that FIE $\in \text{FP}^{\text{NP}}_{||}$ let $\mathcal{H}(\mathcal{X}, \mathcal{A})$ be an AF and consider the following partition of $\mathcal{X}$ (similar to that described in the proof of Thm. 5),

$$\begin{aligned}
\mathcal{X}_{\text{OUT}} &= \{\, x \in \mathcal{X} \;:\; \neg\text{CA}_{\text{ADM}}(\mathcal{H}, x)\} \\
\mathcal{X}_{\text{PSA}} &= \{\, x \in \mathcal{X} \;:\; \{x\}^- \cup \{x\}^+ \subseteq \mathcal{X}_{\text{OUT}}\} \setminus \mathcal{X}_{\text{OUT}} \\
\mathcal{X}_{\text{CA}} &= \{\, x \in \mathcal{X} \;:\; \text{CA}_{\text{ADM}}(\mathcal{H}, x)\} \setminus \mathcal{X}_{\text{PSA}}
\end{aligned}$$

This partition satisfies $\mathcal{X}^-_{\text{PSA}} \subseteq \mathcal{X}_{\text{OUT}}$ and $\mathcal{X}^+_{\text{PSA}} \subseteq \mathcal{X}_{\text{OUT}}$. In addition,

$$\forall\, y \in \mathcal{X}_{\text{CA}}\; \exists\, z \in \mathcal{X}_{\text{CA}}\;\; (\langle y, z\rangle \in \mathcal{A} \text{ or } \langle z, y\rangle \in \mathcal{A})$$

for were this not the case for some $x \in \mathcal{X}_{\text{CA}}$ then x would be in $\mathcal{X}_{\text{PSA}}$ as all of its attackers and attacked arguments would belong to $\mathcal{X}_{\text{OUT}}$.[11]

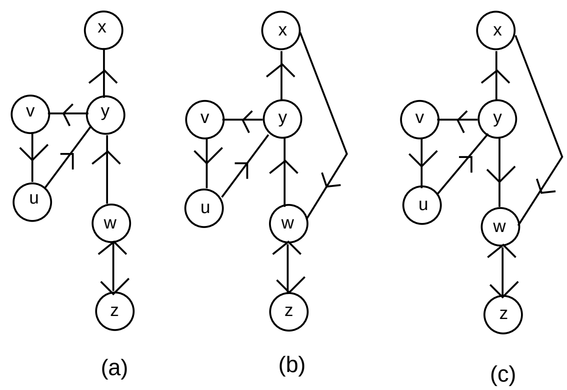

Fig. 1.
a. $\mathcal{X}_{\text{SA}} = \emptyset \subset \{x, v\} = \mathcal{X}_{\text{PSA}}$; $\mathcal{X}_{\text{OUT}} = \{y, u\}$; $\mathcal{X}_{\text{CA}} = \{w, z\}$.
b. $\mathcal{X}_{\text{PSA}} = \{v\}$; $\mathcal{X}_{\text{OUT}} = \{y, u\}$; $\mathcal{X}_{\text{CA}} = \{x, w, z\}$.
c. $\mathcal{X}_{\text{PSA}} = \{z\}$; $\mathcal{X}_{\text{OUT}} = \{x, y, u, v, w\}$; $\mathcal{X}_{\text{CA}} = \emptyset$.

[11] In *coherent* systems $\mathcal{X}_{\text{PSA}}$ is exactly the set of all sceptically accepted arguments, $\mathcal{X}_{\text{SA}}$. In general, however, $\mathcal{X}_{\text{SA}}$ will be a *subset* of $\mathcal{X}_{\text{PSA}}$. The further conditions for membership in $\mathcal{X}_{\text{PSA}}$ are required to distinguish examples such as those illustrated in Fig. 1: in Fig. 1 (b) we have $x \in \mathcal{X}_{\text{CA}}$ rather than $x \in \mathcal{X}_{\text{PSA}}$ despite $\{x\}^- \subseteq \mathcal{X}_{\text{OUT}}$ (on account of the attack $\langle x, w\rangle$; in Fig. 1 (c), although $\{x\}^- \cup \{x\}^+ = \{y, w\} \subset \mathcal{X}_{\text{OUT}}$ since x itself is in $\mathcal{X}_{\text{OUT}}$ it cannot be placed in $\mathcal{X}_{\text{PSA}}$.

With the partition of $\mathcal{X}$ just defined we may construct a *bipartite* framework – $\mathcal{B}(\mathcal{X}_{\text{PSA}}, \mathcal{X}_{\text{OUT}}, \mathcal{F})$ – in which the set of attacks, $\mathcal{F}$, is

$$\mathcal{F} \quad =_{\text{def}} \quad \mathcal{A} \setminus \{ \langle y, z \rangle \ : \ y \in \mathcal{X}_{\text{CA}} \cup \mathcal{X}_{\text{OUT}} \ \text{ and } \ z \in \mathcal{X}_{\text{CA}} \cup \mathcal{X}_{\text{OUT}} \ \}$$

(Note that $\mathcal{B}(\mathcal{X}_{\text{PSA}}, \mathcal{X}_{\text{OUT}}, \mathcal{F})$ *is* bipartite since $\mathcal{X}_{\text{PSA}}$ is conflict-free and $\mathcal{F}$ contains no attacks involving two arguments from $\mathcal{X}_{\text{OUT}}$).

The $\text{FP}_{||}^{\text{NP}}$ upper bound now follows from the following observations:

O1. The partition $\langle \mathcal{X}_{\text{PSA}}, \mathcal{X}_{\text{CA}}, \mathcal{X}_{\text{OUT}} \rangle$ can be constructed using $|\mathcal{X}|$ calls (made in parallel, i.e. non-adaptively) to an NP oracle that decides $\text{CA}(\mathcal{H}, x)$ (one for each $x \in \mathcal{X}$). Each x on which the oracle returns **false** is placed in the set $\mathcal{X}_{\text{OUT}}$ otherwise x is placed into a set $\mathcal{Y}$. The correct partition of $\mathcal{Y}$ into $\mathcal{X}_{\text{PSA}}$ and $\mathcal{X}_{\text{CA}}$ is found by by identifying those arguments in $y \in \mathcal{Y}$ for which $\{y\}^- \cup \{y\}^+ \subseteq \mathcal{X}_{\text{OUT}}$, this set forming $\mathcal{X}_{\text{PSA}}$.

O2. Given the partition $\langle \mathcal{X}_{\text{PSA}}, \mathcal{X}_{\text{CA}}, \mathcal{X}_{\text{OUT}} \rangle$ the bipartite graph $\mathcal{B}(\mathcal{X}_{\text{PSA}}, \mathcal{X}_{\text{OUT}}, \mathcal{F})$ described above, can be constructed from $\mathcal{H}(\mathcal{X}, \mathcal{A})$ by a (deterministic) polynomial time algorithm.

O3. The ideal extension of $\mathcal{H}$ is the maximal admissible subset of $\mathcal{X}_{\text{PSA}}$ in the bipartite graph $\mathcal{B}(\mathcal{X}_{\text{PSA}}, \mathcal{X}_{\text{OUT}}, \mathcal{F})$: this follows from the characterisation proved in Lemma 1.

Using the algorithm of [21, Thm. 6(a)] this set can be found in polynomial time and thus $\text{FIE} \in \text{FP}_{||}^{\text{NP}}$ as claimed. □

The technique employed to established $\text{FP}_{||}^{\text{NP}}$–hardness in proving Thm. 7 can be used to demonstrate $\text{P}_{||}^{\text{NP}}$–hardness for a number of (admittedly rather artificial) decision problems concerning properties of the ideal extension. For example,

Corollary 5 *Let* PARITY-IE *be the decision problem which given an* AF*,* $\mathcal{H}$*, returns* **true** *if and only if the ideal extension of* $\mathcal{H}$ *contains an* odd *number of arguments. The problem* PARITY-IE *is* $\text{P}_{||}^{\text{NP}}$*–complete.*

Proof: Membership is immediate from the construction of Thm. 7. Hardness follows from the fact – [36, Cor. 12.4, p. 274] – that determining the parity of the number of satisfiable formulae in a collection $\langle \Phi_1, \ldots, \Phi_m \rangle$ of given CNFs is $\text{P}_{||}^{\text{NP}}$–hard and the reduction from SC of Thm. 7. □

More generally, for any predicate over collections of CNF formulae related to the cardinality of the set of satisfiable formulae *and* which is $\text{P}_{||}^{\text{NP}}$–hard, the corresponding predicate with respect to the ideal extension of a given AF can also be proven $\text{P}_{||}^{\text{NP}}$–hard using the approach of Corollary 5

Corollary 6 *$\mathcal{H}(\mathcal{X},\mathcal{A})$ is cohesive if every argument $x \in \mathcal{X}$ is credulously accepted.*

Proof: If $\forall x \in \mathcal{X}$ $\text{CA}_{\text{ADM}}(\mathcal{H},x)$ then $\mathcal{X}_{\text{OUT}} = \emptyset$. In such cases, the only arguments, x, that could belong to $\mathcal{X}_{\text{PSA}}$ are those for which $\{x\}^+ \cup \{x\}^- = \emptyset$, i.e. arguments which are "isolated" in $\mathcal{H}(\mathcal{X},\mathcal{A})$. Such arguments are sceptically accepted and form an admissible subset of $\mathcal{X}$. Furthermore no argument in $\mathcal{X} \setminus \mathcal{X}_{\text{PSA}}$ can be sceptically accepted (since each of these attacks or is attacked by has at least one credulously accepted argument). It follows that $\mathcal{X}_{\text{PSA}} = \{\ x\ :\ \text{SA}_{\text{PR}}(\mathcal{H},x)\}$ and $\mathcal{X}_{\text{PSA}}$ is the ideal extension, i.e. $\mathcal{H}$ is cohesive. ▫

Combining the results and noting the equivalences

$$\text{CA}_{\text{PR}} \equiv \text{CA}_{\text{ADM}} \quad ; \quad \text{CA}_{\text{IDL}} \equiv \text{CA}_{\text{IE}} \equiv \text{SA}_{\text{IE}}$$

we obtain the picture of the relative complexities in Table 4.

Table 4
Relative Complexity of Testing Acceptability.

Decision problem	Lower bound	Upper Bound
CA_{PR}	NP–hard	NP
CA_{IDL}	coNP–hard	$\text{P}_{\|\|}^{\text{NP}}$
SA_{PR}	Π_2^p–hard	Π_2^p

Similarly, Table 5 considers checking whether a given *set* of arguments collectively satisfies the requirements of a given semantics or is a maximal such set, i.e. the various cases of the verification problem

Table 5
Deciding set and maximality properties

Semantics	Lower bound	Upper Bound
ADM	P	P
IDL	coNP–hard	coNP
PR	coNP–hard	coNP
IE	D^p–hard	$\text{P}_{\|\|}^{\text{NP}}$

We note in passing that the problem of deciding if S is the maximal set of *sceptically* accepted arguments, although not previously considered, is easily shown to be complete for the complexity class D_2^p of languages L expressible as the intersection of a language $L_1 \in \Sigma_2^p$ and $L_2 \in \Pi_2^p$.

In total the classifications given by the these tables reinforce the case that $\mathrm{CA}_{\mathrm{ADM}}$ is easier than $\mathrm{CA}_{\mathrm{IDL}}$ which, in turn, is easier than $\mathrm{SA}_{\mathrm{PR}}$.

2.6 Reducing the complexity gaps

In [7], Chang and Kadin introduce the concepts of a language having the properties OP_2 and OP_ω where OP is one of the Boolean operators $\{\mathrm{AND}, \mathrm{OR}\}$. Formally,

Definition 8 *([7, pp. 175–76] Let L be a language, i.e. a set of finite words over an alphabet. The languages, $\mathrm{AND}_k(L)$ and $\mathrm{OR}_k(L)$ ($k \geq 1$) are*

$$\begin{aligned}\mathrm{AND}_k(L) &=_{\mathrm{def}} \{\langle w_1, w_2, \ldots, w_k\rangle \ : \ \forall\, 1 \leq i \leq k\ w_i \in L\}\\ \mathrm{OR}_k(L) &=_{\mathrm{def}} \{\langle w_1, w_2, \ldots, w_k\rangle \ : \ \exists\, 1 \leq i \leq k\ w_i \in L\}\end{aligned}$$

The languages $\mathrm{AND}_\omega(L)$ and $\mathrm{OR}_\omega(L)$ are,

$$\mathrm{AND}_\omega(L) =_{\mathrm{def}} \bigcup_{k\geq 1} \mathrm{AND}_k(L) \quad ; \quad \mathrm{OR}_\omega(L) =_{\mathrm{def}} \bigcup_{k\geq 1} \mathrm{OR}_k(L)$$

A language, L, is said to have property OP_k (resp. OP_ω) if $\mathrm{OP}_k(L) \ \leq_m^p \ L$ (resp. $\mathrm{OP}_\omega(L) \ \leq_m^p \ L$).

The reason why these language operations are of interest is the following result.

Fact 9 *([7, Thm. 9, p. 182])*
A language L is $\mathrm{P}_{||}^{\mathrm{NP}}$–complete (via $\leq_m^p$ reducibility) if and only if all of the following hold.

F1. $L \in \mathrm{P}_{||}^{\mathrm{NP}}$.
F2. L *is* NP*–hard* and L *is co*NP*–hard.*
F3. L *has property* AND_2.
F4. L *has property* OR_ω.

As a consequence of Fact 9, we have,

Theorem 10

a. *If $\mathrm{CA}_{\mathrm{IDL}}$ is NP–hard then $\mathrm{CA}_{\mathrm{IDL}}$ is $\mathrm{P}_{||}^{\mathrm{NP}}$–complete.*
b. *If $\mathrm{CA}_{\mathrm{IDL}} \in$ coNP then $\mathrm{VER}_{\mathrm{IE}}$ is D^p–complete.*
c. *If $\mathrm{CA}_{\mathrm{IDL}} \in$ coNP then $\mathrm{VER}_{\mathrm{IE}}^{\emptyset}$ is NP–complete.*
d. *If $\mathrm{VER}_{\mathrm{IE}}$ has property OR_ω then $\mathrm{VER}_{\mathrm{IE}}$ is $\mathrm{P}_{||}^{\mathrm{NP}}$–complete.*

Proof:
a. With the assumption that $\textsc{ca}_{\mathrm{IDL}}$ is $\textsc{np}$–hard, $\textsc{ca}_{\mathrm{IDL}}$ would satisfy conditions (F1) and (F2) of Fact 9. To complete the argument it suffices to show that $\textsc{ca}_{\mathrm{IDL}}$ already has property $\textsc{and}_2$ and property $\textsc{or}_\omega$. For the first of these consider any instance $\langle\langle\mathcal{H}_1,x\rangle,\langle\mathcal{H}_2,y\rangle\rangle$ of $\textsc{and}_2(\textsc{ca}_{\mathrm{IDL}})$. Form the $\textsc{af}$, $\mathcal{H}$, consisting of copies of $\mathcal{H}_1$ and $\mathcal{H}_2$ together with three additional arguments $\{z_x,z_y,z\}$. Now adding the attacks $\{\langle x,z_x\rangle,\langle y,z_y\rangle,\langle z_x,z\rangle,\langle z_y,z\rangle\}$, via Lemma 2, $\langle\mathcal{H},z\rangle$ is accepted as an instance of $\textsc{ca}_{\mathrm{IDL}}$ if and only if $\textsc{ca}_{\mathrm{IDL}}(\langle\mathcal{H}_1,x\rangle)\wedge\textsc{ca}_{\mathrm{IDL}}(\langle\mathcal{H}_2,y\rangle)$. To see that $\textsc{ca}_{\mathrm{IDL}}$ has property $\textsc{or}_\omega$ consider an instance $\langle\ \langle\mathcal{H}_1,x_1\rangle,\langle\mathcal{H}_2,x_2\rangle,\ldots,\langle\mathcal{H}_m,x_m\rangle\ \rangle$ of $\textsc{or}_\omega(\textsc{ca}_{\mathrm{IDL}})$. Form an $\textsc{af}$, $\mathcal{H}$, from these m frameworks, adding two new arguments, $\{y,z\}$. The instance is completed by adding the attacks $\{\langle x_i,y\rangle\ :\ 1\leq i\leq m\}$ and the attack $\langle y,z\rangle$. Again, via Lemma 2, $\langle\mathcal{H},z\rangle$ is accepted as an instance of $\textsc{ca}_{\mathrm{IDL}}$ if and only if $\vee_{i=1}^{m}\ \textsc{ca}_{\mathrm{IDL}}(\langle\mathcal{H}_i,x_i\rangle)$.
b. It has already been shown that $\textsc{ver}_{\mathrm{IE}}$ is $\textsc{d}^p$–hard. Consider the languages,

$$
\begin{aligned}
L_1 &=_{\mathrm{def}}\ \{\langle\mathcal{H},S\rangle\ :\ \forall\, x\in S,\ \textsc{ca}_{\mathrm{IDL}}(\mathcal{H},x)\ \}\\
L_2 &=_{\mathrm{def}}\ \{\langle\mathcal{H},S\rangle\ :\ \forall\, x\notin S,\ \neg\textsc{ca}_{\mathrm{IDL}}(\mathcal{H},x)\ \}
\end{aligned}
$$

We have $L_1\ \in$ co$\textsc{np}$ (by the assumption $\textsc{ca}_{\mathrm{IDL}}$ is in co$\textsc{np}$ and by the straighforward generalisation of (a) that shows $\textsc{ca}_{\mathrm{IDL}}$ has property $\textsc{and}_\omega$). In addition, $L_2\in\textsc{np}$ (from the premise $\textsc{ca}_{\mathrm{IDL}}\in$ co$\textsc{np}$ and the fact that $\neg\textsc{ca}_{\mathrm{IDL}}$ has property $\textsc{and}_\omega$ since $\textsc{ca}_{\mathrm{IDL}}$ has property $\textsc{or}_\omega$). With these choices of L_1 and L_2, $\langle\mathcal{H},S\rangle$ is accepted as an instance $\textsc{ver}_{\mathrm{IE}}$ if and only if $\langle\mathcal{H},S\rangle\in L_1\cap L_2$ so that $\textsc{ver}_{\mathrm{IE}}\in\textsc{d}^p$.
c. Easy consequence of (b).
d. It has already been shown that $\textsc{ver}_{\mathrm{IE}}$ satisfies (F1) and (F2) of Fact 9. In addition, $\textsc{ver}_{\mathrm{IE}}$ has property $\textsc{and}_\omega$ (thus, trivially, also $\textsc{and}_2$): given an instance $\langle\ \langle\mathcal{H}_1,S_1\rangle,\langle\mathcal{H}_2,S_2\rangle,\ldots,\langle\mathcal{H}_m,S_m\rangle\ \rangle$ of $\textsc{and}_\omega(\textsc{ver}_{\mathrm{IE}})$ fix $\mathcal{H}$ to consist of the m frameworks $\langle\mathcal{H}_1,\mathcal{H}_2,\ldots,\mathcal{H}_m\rangle$ and S as $\cup_{i=1}^{m}\ S_i$. With these, $\langle\mathcal{H},S\rangle$ is accepted as an instance of $\textsc{ver}_{\mathrm{IE}}$ if and only if $\wedge_{i=1}^{m}\ \textsc{ver}_{\mathrm{IE}}(\mathcal{H}_i,S_i)$. It follows that were $\textsc{ver}_{\mathrm{IE}}$ to have property $\textsc{or}_\omega$, then $\textsc{ver}_{\mathrm{IE}}$ would be $\textsc{p}_{||}^{\mathrm{NP}}$-complete via Fact 9. □

We may interpret Thm. 10 as focusing the issue of obtaining exact classifications in terms of $\textsc{ca}_{\mathrm{IDL}}$. If $\textsc{ca}_{\mathrm{IDL}}\in$ co$\textsc{np}$ (so that, with the usual assumption of $\textsc{np}\neq$co$\textsc{np}$, $\textsc{ca}_{\mathrm{IDL}}$ would not be $\textsc{np}$–hard) then we obtain exact classifications of the complexity of $\{\textsc{ca}_{\mathrm{IDL}},\textsc{ver}_{\mathrm{IE}},\textsc{ver}_{\mathrm{IE}}^{\emptyset}\}$ as $\{\mathrm{co}\textsc{np},\textsc{d}^p,\textsc{np}\}$–complete. On the other hand, an alternative hypothesis, in the event of $\textsc{ca}_{\mathrm{IDL}}\ \not\in$ co$\textsc{np}$, is that suggested by Thm. 10 (a): that $\textsc{ca}_{\mathrm{IDL}}$ is $\textsc{p}_{||}^{\mathrm{NP}}$–complete, a result which would follow by demonstrating $\textsc{ca}_{\mathrm{IDL}}$ to be $\textsc{np}$–hard.

In fact, there is strong evidence that $\textsc{ca}_{\mathrm{IDL}}\ \not\in$ co$\textsc{np}$ and, using one suite of techniques is more likely to be complete within $\textsc{p}_{||}^{\mathrm{NP}}$. Our formal justification of these claims rests on a number of technical analyses using results of Chang *et al.* [8], which in turn develop ideas of [1,4,34]. Two key concepts in our further analyses of CA_{IDL} are,

a. The so-called *Unique Satisfiability* problem (USAT).
b. Randomized reductions between languages.

Unique Satisfiability (USAT)
Instance: CNF formula $\Phi(X_n)$ with propositional variables $\langle x_1, \ldots, x_n \rangle$.
Question: Does $\Phi(X_n)$ have *exactly one* satisfying instantiation?

Determining the exact complexity of USAT remains an open problem. It is known that USAT $\in$ D^p and while Blass and Gurevich [4] show it to be coNP–hard [12], USAT has only be shown to be complete for D^p using a *randomized* reduction technique of Valiant and Vazirani [34]. Two concepts of such reductions are studied in Chang *et al.* [8] specifically with respect to USAT via the following general definition.

Definition 11 *Let L_1 and L_2 be languages and $\delta \in [0, 1]$. We say that L_1* randomly reduces to L_2 *(denoted $L_1 \leq_m^{rp} L_2$)* with probability δ *if there is a polynomial time computable function, f, and polynomial bound q with f mapping pairs $\langle x, z \rangle$ – x an instance of L_1 and z an element of $\langle 0, 1 \rangle^{q(|x|)}$ – to instances, y, of L_2, such that for z drawn uniformly at random from $\langle 0, 1 \rangle^{q(|x|}$*

$$x \in L_1 \ \Rightarrow \ Prob[f(x, z) \in L_2] \ \geq \ \delta$$
$$x \notin L_1 \ \Rightarrow \ Prob[f(x, z) \notin L_2] \ = \ 1$$

We have the following properties of USAT and randomized reductions:

Fact 12

a. SAT $\leq_m^{rp}$ USAT *with probability* $1/(4n)$*. ([34, Lemma 2.1, p. 88])*
b. If $L_1 \leq_m^{rp} L_2$ with probability $1/p(n)$ for some polynomially bounded function, p, and L_2 has property OR_ω *then $L_1 \leq_m^{rp} L_2$ with probability $1 - 2^{-n}$. ([8, Fact 1, p. 361]* [13] *)*

A relationship between unique satisfiability (USAT) and CA_{IDL} is established in the following theorem. Notice that the reduction we describe is *deterministic*, i.e. not randomized.

[12] The reader should note that [28, p. 93] has a typographical slip whereby Blass and Gurevich's result is described as proving USAT to be NP–hard.
[13] The bound actually stated in [8] is for arbitrary exponentially decreasing functions, i.e. not just 2^{-n}.

Theorem 13 USAT $\leq_m^p$ CA$_{\text{IDL}}$.

Proof: Given an instance $\Phi(Z_n)$ of USAT construct an AF, $\mathcal{K}(\mathcal{X},\mathcal{A})$ as follows. First form the system $\mathcal{F}_\Phi$ described in Thm. 2, but without the attack $\langle\Psi,\Phi\rangle$ contained in this and with attacks $\langle C_j, C_j\rangle$ for each clause of Φ.[14] We then add a further $n+1$ arguments, $\{y_1,\ldots,y_n,x\}$ and attacks

$$\{\langle z_i,\ y_i\rangle,\ \langle\neg z_i, y_i\rangle\ :\ 1\leq i\leq n\}\ \cup\ \{\langle y_i, x\rangle\ :\ 1\ \leq\ i\ \leq\ n\}$$

The instance of CA$_{\text{IDL}}$ is $\langle\mathcal{K}(\mathcal{X},\mathcal{A}),x\rangle$ and the resulting AF is illustrated in Fig. 2.

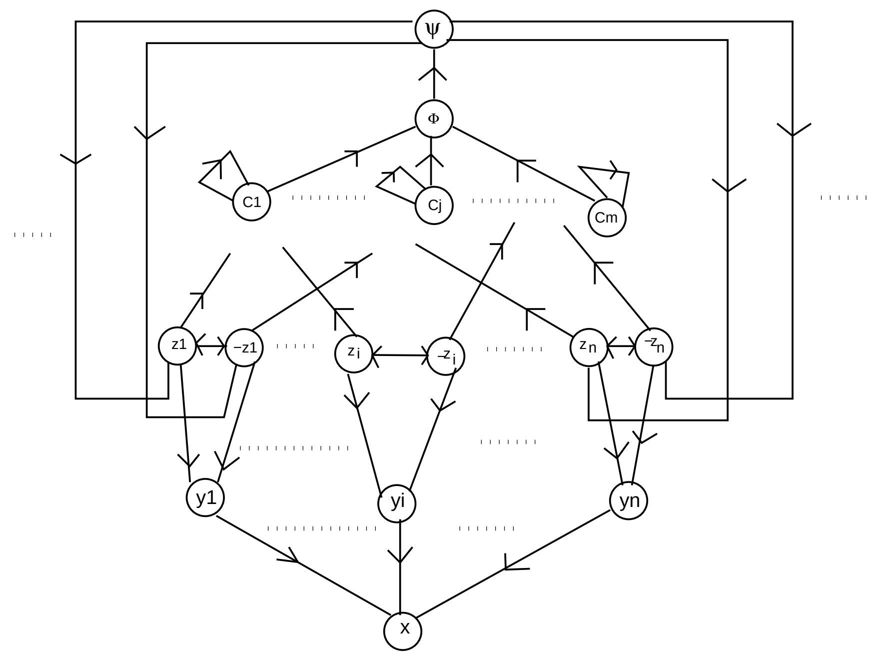

Fig. 2. The Argumentation Framework $\mathcal{K}_\Phi$

We now claim that $\Phi(Z_n)$ has a unique satisfying instantiation if and only if x is a member of $\mathcal{M}_\mathcal{K}$ the ideal extension of $\mathcal{K}(\mathcal{X},\mathcal{A})$.

Suppose first that $\Phi(Z_n)$ does *not* have a unique satisfying instantiation. If Φ is unsatisfiable – i.e. the number of satisfying assignments is zero – then all of the arguments forming the sub-system, $\mathcal{F}_\Phi$, fail to be credulously accepted, in particular, each of the arguments z_i and $\neg z_i$ fail to be so accepted. It easily follows that $x\ \not\in\ \mathcal{M}_\mathcal{K}$ since no defence to the attack on x by y_i is possible. There remains the possibility that $\Phi(Z_n)$ has two or more satisfying assignments. Suppose

[14] We make these arguments self-attacking purely for ease of presentation: the required effect - that no argument C_j is ever credulously accepted – can be achieved without self-attacks simply by adding two arguments d_j and e_j for each clause together with attacks $\{\langle C_j,d_j\rangle,\langle d_j,e_j\rangle,\langle e_j,C_j\rangle\}$.

$\alpha = \langle a_1, a_2, \ldots, a_n\rangle$ and $\beta = \langle b_1, b_2, \ldots, b_n\rangle$ are such that $\Phi(\alpha) = \Phi(\beta) = \top$ and $\alpha \neq \beta$. Without loss of generality, we may assume that $a_1 \neq b_1$ (since $\alpha \neq \beta$ there must be at least one variable of Z_n that is assigned differing values in each). In this case *both* z_1 and $\neg z_1$ are credulously accepted so that neither can belong to $\mathcal{M}_{\mathcal{K}}$: from Lemma 2 condition (M1) gives $z_1 \notin \mathcal{M}_{\mathcal{K}}$ (since $\neg z_1$ is credulously accepted) and $\neg z_1 \notin \mathcal{M}_{\mathcal{K}}$ (since z_1 is credulously accepted). It now follows that $x \notin \mathcal{K}_{\mathcal{M}}$ via (M2) of Lemma 2: neither attacker of y_1, an argument which attacks x, belongs to $\mathcal{M}_{\mathcal{K}}$. We deduce that if $\Phi(Z_n)$ is not a positive instance of USAT then $\langle \mathcal{K}, x\rangle$ is not a positive instance of CA_{IDL}.

One the other hand suppose that $\alpha = \langle a_1, a_2, \ldots, a_n\rangle$ defines the unique satisfying instantiation of $\Phi(Z_n)$. Consider the following subset of $\mathcal{X}$:

$$\mathcal{M} \;=\; \bigcup_{i\,:\,a_i=\top} \{z_i\} \;\cup\; \bigcup_{i\,:\,a_i=\bot} \{\neg z_i\} \;\cup\; \{\Phi, x\}$$

Certainly $\mathcal{M}$ is admissible: since α satisfies $\Phi(Z_n)$ each C_j and y is attacked by some z or $\neg z$ in $\mathcal{M}$ and thus all of the attacks on Φ and x are counterattacked. Similarly Φ defends arguments against the attacks by Ψ. It is also the case, however, that no admissible set of $\mathcal{K}$ contains an attacker of $\mathcal{M}$. No admissible set can contain C_j (since these arguments are self-attacking), Ψ (since the only defenders of the attack by Φ are C_j arguments) or y_k ($1 \leq k \leq n$) (since these require Ψ as a defence against $\{z_k, \neg z_k\}$). Furthermore for $z_i \in \mathcal{M}$ an admissible set containing $\neg z_i$ would only be possible if there were a satisfying assignment of Φ under which $\neg z_i = \top$: this would contradict the assumption the Φ had exactly one satisfying instantiation.

We deduce that $\Phi(Z_n)$ has a unique satisfying instantiation if and only if x is in the ideal extension of $\mathcal{K}(\mathcal{X}, \mathcal{A})$. $\square$

Combining Thms. 10 and 13 with Facts 9 and 12 gives the following corollaries.

Corollary 7 USAT $\leq^p_m$ $\neg\text{VER}^{\emptyset}_{\text{IE}}$

Proof: The AF $\mathcal{K}_\Phi$ of Thm. 13 has a *non-empty* ideal extension, $\mathcal{M}_{\mathcal{K}}$, if and only if $x \in \mathcal{M}_{\mathcal{K}}$. $\square$

Corollary 8 CA_{IDL} *is complete for* $\text{P}^{\text{NP}}_{||}$ *under* $\leq^{rp}_m$ *with probability* $1 - 2^{-n}$.

Proof: The decision problem $\text{OR}_\omega(\text{SAT-UNSAT})$ is $\text{P}^{\text{NP}}_{||}$–complete under (standard, deterministic) $\leq^p_m$ reductions. We thus obtain

$$\text{OR}_\omega(\text{SAT-UNSAT}) \;\leq^{rp}_m\; \text{SAT-UNSAT} \text{ with probability } 1/n$$

(as observed in [8, Lemma 1, p. 365], simply choose, uniformly at random, one of the n sub-problems $\langle \Phi_i, \Psi_i \rangle$ in the instance $\langle \langle \Phi_1, \Psi_1 \rangle, \ldots, \langle \Phi_n, \Psi_n \rangle \rangle$ of $\mathrm{OR}_\omega(\mathrm{SAT\text{-}UNSAT})$.

Now, via [34], $\mathrm{SAT\text{-}UNSAT} \ \leq_m^{rp} \ \mathrm{USAT}$ with probability $1/(4n)$ so that, combining these randomized reductions,

$$\mathrm{OR}_\omega(\mathrm{SAT\text{-}UNSAT}) \ \leq_m^{rp} \ \mathrm{USAT} \text{ with probability } 1/(4n^2)$$

Now applying the (deterministic) reduction of Thm. 13 shows

$$\mathrm{OR}_\omega(\mathrm{SAT\text{-}UNSAT}) \ \leq_m^{rp} \ \mathrm{CA}_{\mathrm{IDL}} \text{ with probability } 1/(4n^2)$$

As demonstrated in the proof of Thm. 10(a), $\mathrm{CA}_{\mathrm{IDL}}$ has property OR_ω so that via Fact 12(b) we obtain,

$$\mathrm{OR}_\omega(\mathrm{SAT\text{-}UNSAT}) \ \leq_m^{rp} \ \mathrm{CA}_{\mathrm{IDL}} \text{ with probability } 1 - 2^{-n}$$

Since we know that $\mathrm{CA}_{\mathrm{IDL}} \in \mathrm{P}_{||}^{\mathrm{NP}}$ this completes the proof. □

Corollary 9 $\mathrm{VER}_{\mathrm{IE}}^{\emptyset}$ *is complete for* $\mathrm{P}_{||}^{\mathrm{NP}}$ *via* $\leq_m^{rp}$ *with probability* $1 - 2^{-n}$.

Proof: We may apply a similar argument to that of Corollary 8 to obtain $\mathrm{OR}_\omega(\mathrm{SAT\text{-}UNSAT}) \leq_m^{rp} \neg\mathrm{VER}_{\mathrm{IE}}^{\emptyset}$ with probability $1/(4n^2)$. Since $\mathrm{VER}_{\mathrm{IE}}^{\emptyset}$ has property AND_ω so its complement, has property OR_ω. The corollary now follows via Fact 12(b) and the fact that $\mathrm{P}_{||}^{\mathrm{NP}}$ is closed under complementation. □

Corollary 10 $\mathrm{VER}_{\mathrm{IE}}$ *is complete for* $\mathrm{P}_{||}^{\mathrm{NP}}$ *via* $\leq_m^{rp}$ *with probability* $1 - 2^{-n}$.

Proof: Easy consequence of Corollary 9: $\mathrm{VER}_{\mathrm{IE}}^{\emptyset}$ is a special case of $\mathrm{VER}_{\mathrm{IE}}$. □

To conclude we observe that although $\mathrm{USAT} \leq_m^p \mathrm{CA}_{\mathrm{IDL}}$ it is unlikely to be the case that these decision problems have equivalent complexity, i.e. that $\mathrm{CA}_{\mathrm{IDL}} \leq_m^p \mathrm{USAT}$.

Corollary 11 *If* $\mathrm{CA}_{\mathrm{IDL}} \leq_m^p \mathrm{USAT}$ *(note* deterministic *reduction) then the Polynomial Hierarchy (*PH*) collapses to* Σ_3^p*, i.e.*

$$\mathrm{CA}_{\mathrm{IDL}} \leq_m^p \mathrm{USAT} \quad \Rightarrow \quad \bigcup_{k \geq 3} \Sigma_k^p \cup \bigcup_{k \geq 3} \Pi_k^p \ \subseteq \ \Sigma_3^p$$

Proof: Suppose it is the case that $\text{CA}_{\text{IDL}} \leq^p_m \text{USAT}$. We then have

$$\begin{array}{llll} \text{OR}_\omega(\text{USAT}) & \leq^p_m & \text{OR}_\omega(\text{CA}_{\text{IDL}}) & \text{by Thm. 13} \\ & \leq^p_m & \text{CA}_{\text{IDL}} & \text{since } \text{CA}_{\text{IDL}} \text{ has property } \text{OR}_\omega \\ & \leq^p_m & \text{USAT} & \text{by premise} \end{array}$$

So that USAT would have property OR_ω: [8, Thm. 5, p. 364] demonstrates that this leads to the collapse stated. □

Now, noting that $\leq^p_m$ can be interpreted as "$\leq^{rp}_m$ with probability 1", we can reconsider the lower bounds of Tables 4 and 5 using hardness via $\leq^{rp}_m$ (with "high" probability) instead of hardness via (deterministic) $\leq^p_m$, as shown in Table 6.

Table 6
Complexity of ideal semantics relative to randomized reductions

Decision Problem	Complexity	$\leq^{rp}_m$ probability
CA_{ADM}	NP–complete	1
CA_{IDL}	$\text{P}^{\text{NP}}_{\|\|}$–complete	$1-2^{-n}$
SA_{PR}	Π^p_2–complete	1
VER_{ADM}	P	–
VER_{IDL}	coNP–complete	1
VER_{PR}	coNP–complete	1
$\text{VER}^{\emptyset}_{\text{PR}}$	coNP–complete	1
VER_{IE}	$\text{P}^{\text{NP}}_{\|\|}$–complete	$1-2^{-n}$
$\text{VER}^{\emptyset}_{\text{IE}}$	$\text{P}^{\text{NP}}_{\|\|}$–complete	$1-2^{-n}$

3 Ideal Semantics in Assumption-based frameworks

The formalism of abstract assumption-based argumentation frameworks (ABFs) is described in Bondarenko *et al.* [5] and offers an alternative but related approach to the AF mechanisms of Dung [17]. Whereas the concept of argument and attack within AFs does not attempt to analyse issues of argument *structure* or the rationale underpinning attacks between arguments, ABFs view arguments as a statements justified through some formal logical deductive system with the concept of attack being that the conclusion of one argument is incompatible with the premises supporting another. We now consider similar complexity issues to those examined in the preceding sections for AFs, concentrating on the model of ABFs. Although there is some variation in the exact specification of decision problems, as before, the canonical questions of interest concern Verification and Credulous Acceptance.

We first review the basic elements of ABFs in Section 3.1 including the formal description of the verification and credulous reasoning problems. Following, in Section 3.2 we describe the translations, from [5] of divers non-classical logics into corresponding ABF contexts and summarise the contribution of [12–14] in which the computational complexity of credulous and sceptical reasoning under preferred, stable and admissible semantics was considered.

Finally in Section 3.3 we consider the computational complexity of ideal semantics in ABFs using a number of the settings described in Section 3.2.

3.1 Review of Elements from Assumption-based argumentation frameworks

In the sequel $\langle L, R\rangle$ is a *deductive system*, i.e. L is a formal language whose elements are denumerable – e.g. well-formed propositional formulae – and R is a set of *inference rules*, which we consider to be of the form

$$\frac{\alpha_1, \alpha_2, \ldots, \alpha_n}{\beta}$$

with $\alpha_i \in L$, $\beta \in L$ and $n \geq 0$. We refer to any $T \subseteq L$ as a *theory*. Given $T \subseteq L$ in the system $\langle L, R\rangle$ the subset $Th(T)$ of L of *derivable* sentences for T in $\langle L, R\rangle$ has $\alpha \in Th(T)$ if there is a (finite) sequence $\beta_1, \ldots, \beta_m$ ($m > 0$) with which for every i ($1 \leq i \leq m$) either $\beta_i \in T$ or there is a rule [15] $\beta_i \leftarrow \alpha_1, \alpha_2, \ldots, \alpha_n \in R$ and $\{\alpha_1, \ldots, \alpha_n\} \subseteq \{\beta_1, \ldots, \beta_{i-1}\}$. We write $T \vdash \alpha$ if $\alpha \in Th(T)$. Such systems are *monotonic*, i.e. if $T \vdash \alpha$ then $T' \vdash \alpha$ for any $T' \supseteq T$.

As shown by several examples in [5] coupling the well-studied formalism of deductive systems with the novel concepts of *assumptions* and *contrary* provides an (argumentation founded) approach giving a unified treatment of a wide variety of non-classical logics.

Definition 14 *For a deductive system $\langle L, R\rangle$ an* assumption-based framework *w.r.t. $\langle L, R\rangle$ is a triple $\langle T, A, ^{-}\rangle$ where $T \subseteq L$ is a theory, $A \subseteq L$ a non-empty set of* assumptions*, and $^{-} : A \to L$ a mapping that associates with each assumption $\alpha \in A$ its* contrary*, denoted $\overline{\alpha}$.*

Although the contrary mapping can be instantiated as classical negation, it is *not* limited to this sense only. Instead of the atomic notion of argument from [17], the

[15] For ease of readability we use the form $\beta \leftarrow \alpha_1, \ldots, \alpha_n$ inside text.

objects of interest within ABFs are subsets of *assumptions* that define *extensions* of the theory T according to various semantics. In this way, the following mirrors Defn. 1 for AFs, presenting analogous ideas in ABFs.

Definition 15 *Let $\langle T, A, ^{-}\rangle$ be an* ABF *w.r.t. some deductive system $\langle L, R\rangle$ and $\Delta \subseteq A$. For $\varphi \in L$, we write $\Delta \models \varphi$ as a shorthand for $\varphi \in Th(T \cup \Delta)$.*

The set of assumptions Δ attacks *an assumption $\alpha \in A$ if $\Delta \models \overline{\alpha}$; Δ attacks a* set *of assumptions Δ' if $\Delta \models \overline{\alpha}$ for some $\alpha \in \Delta'$. We write $att(\Delta, \alpha)$ to denote the attack relation over $2^A \times A$, and similarly (albeit with a slight abuse of notation) use $att(\Delta, \Delta')$ for attacks by Δ on a set of assumptions Δ'. The assumption set Δ is said to be* closed *if $\Delta = \{\alpha \in A \ : \ \Delta \models \alpha\}$, i.e. a closed assumption set cannot derive any assumption other than those already contained in it. Those frameworks whose supporting deductive systems are such that* every *set of assumptions is closed are called* flat frameworks.

A set $\Delta \subseteq A$ is conflict-free *if $\neg att(\Delta, \Delta)$; Δ is an* admissible *set of assumptions if it is closed, conflict-free, and for every* closed *assumption set Δ' if $att(\Delta', \Delta)$ then $att(\Delta, \Delta')$; Δ is a preferred extension if it is a maximal admissible set. A set Δ is a* stable *extension if it is closed, conflict-free, and for every $\alpha \in A \setminus \Delta$, $att(\Delta, \alpha)$, i.e $\Delta \models \overline{\alpha}$.*

Following [19,20], Δ is an ideal set *if Δ it is admissible and a subset of every preferred extension; Δ is (the) ideal* extension *if it is the maximal such set.* [16]

Corresponding to the decision problems for AFs considered earlier, we have the following formulations in ABFs. Note that the underlying deductive system $\langle L, R\rangle$ is reflected in the problem *name* rather than explicitly as part of the *instance*. We use $\mathcal{E}_s(\langle T, A, ^{-}\rangle)$ to denote the subsets of assumptions satisfying the criteria of semantics s in the ABF $\langle T, A, ^{-}\rangle$.

Table 7
Decision Problems in ABFs

Problem Name	Instance	Question
Verification ($\text{VER}_s^{\langle L,R\rangle}$)	$\langle T, A, ^{-}\rangle$; $\Delta \subseteq A$	Is $\Delta \in \mathcal{E}_s(\langle T, A, ^{-}\rangle)$?
Credulous Acceptance ($\text{CA}_s^{\langle L,R\rangle}$)	$\langle T, A, ^{-}\rangle$; $\varphi \in L$	$\exists\ \Delta \in \mathcal{E}_s(\langle T, A, ^{-}\rangle)$ for which $\Delta \models \varphi$?
Sceptical Acceptance ($\text{SA}_s^{\langle L,R\rangle}$)	$\langle T, A, ^{-}\rangle$; $\varphi \in L$	$\forall\ \Delta \in \mathcal{E}_s(\langle T, A, ^{-}\rangle)$ does $\Delta \models \varphi$?
Existence ($\text{EXISTS}_s^{\langle L,R\rangle}$)	$\langle T, A, ^{-}\rangle$	Is $\mathcal{E}_s(\langle T, A, ^{-}\rangle) \neq \emptyset$?
Emptiness ($\text{VER}_s^{\langle L,R\rangle,\emptyset}$)	$\langle T, A, ^{-}\rangle$	Is $\mathcal{E}_s(\langle T, A, ^{-}\rangle) = \{\emptyset\}$?

The formulations of $\text{CA}_s^{\langle L,R\rangle}$ and $\text{SA}_s^{\langle L,R\rangle}$ are rather more general than might seem to be the natural analogue: rather than asking whether a given *assumption*, α belongs to at least one (or every) set in $\mathcal{E}_s$ (so directly paralleling the form of CA_s and

[16] We recall that [19,20] have shown every AF, ABF has a uniquely defined ideal extension.

SA$_s$ from AFs), the decision problems ask whether a given *sentence*, $\varphi \in L$, can be deduced via at least one (resp. every) set in $\mathcal{E}_s$. In complexity terms for flat frameworks and $\varphi = \alpha \in A$, the given form is equivalent to the natural analogue: $\Delta \models \varphi = \alpha$ if and only if $\alpha \in \Delta$.

3.2 *Instantiations of* ABFs *modeling default reasoning and their complexity*

In this section we reprise the translations from a range of reasoning formalisms into equivalent ABFs. Our presentation summarises the descriptions from [5,14].

3.2.1 *Logic Programming* – LP

We recall that a (normal) logic program, T, comprises a set of clauses of the form $\alpha \leftarrow \beta_1, \ldots, \beta_m$ where α is a ground atom from some underyling *Herbrand base* ($\mathcal{HB}$) and β_i is a literal from $Lits = \mathcal{HB} \cup \mathcal{HB}_{not}$ where $\mathcal{HB}_{not} = \{not\ \alpha : \alpha \in \mathcal{HB}\}$. Given a normal logic program, T, the corresponding ABF is $\langle T, \mathcal{HB}_{not}, ^{-} \rangle$ where for each assumption $not\ \alpha \in \mathcal{HB}_{not}$ its contrary $\overline{not\ \alpha} = \alpha$. The underlying deductive system $\langle L, R \rangle$ corresponds to Horn logic derivability, where following [25], assumptions $not\ \alpha$ are regarded as new atoms α^*.

3.2.2 *Default Logic* – DL

Given a deductive system for classical first-order logic, $\langle L_0, R_0 \rangle$ Reiter [33] defines a *default theory* as a pair $\langle W, D \rangle$ wherein $W \subseteq L_0$ and D is a set of *default rules*

$$\frac{\alpha\ :\ M\beta_1, \ldots, M\beta_n}{\gamma} \quad where\ \alpha,\ \beta_1, \ldots, \beta_n,\ \gamma \in L_0,\ n \geq 0$$

Informally a default rule may be interpreted as "it is reasonable to assume γ if we know (or have proved that) α is the case and have no basis on which to suppose any $\neg\beta_i$ ($1 \leq i \leq n$) holds", e.g. in the standard example of default reasoning "it is reasonable to assume that *Tweety* can fly if we know that *Tweety* is a bird and have no basis to suppose either that *Tweety* is a penguin or that *Tweety* can not fly." is expressed via the default rule

$$\frac{bird(Tweety)\ :\ M\neg\, penguin(Tweety),\ M\ flies(Tweety)}{flies(Tweety)}.$$

For the default theory $\langle W, D\rangle$ the related ABF $\langle T, A, ^{-}\rangle$ is formulated in the deductive system $\langle L, R\rangle$ with

$$\begin{array}{lcl} L & = & L_0 \cup \{M\alpha \;:\; \alpha \in L_0\} \\ R & = & R_0 \cup D \\ T & = & W \\ A & = & \{M\beta \;:\; \beta \in L_0 \text{ and } M\beta \text{ occurs in some default rule of } D\} \\ \overline{M\alpha} & = & \neg\,\alpha \end{array}$$

3.2.3 Autoepistemic Logic – AEL

In AEL the starting point is a deductive system, $\langle L, R\rangle$, in which L is a modal language with modal operator **B**, and R an inference scheme of classical logic for L: the interpretation of $\mathbf{B}\alpha$ being that "α is *believed*". For a theory $T \subseteq L$ of AEL, the corresponding ABF, $\langle T, A, ^{-}\rangle$ w.r.t. $\langle L, R\rangle$ has $A = \{\mathbf{B}\alpha \;:\; \alpha \in L\} \cup \{\neg\mathbf{B}\alpha \;:\; \alpha \in L\}$. The contrary mapping has $\overline{\mathbf{B}\alpha} = \neg\mathbf{B}\alpha$ and $\overline{\neg\mathbf{B}\alpha} = \alpha$.

Both Reiter [33] and Konolige [29] have observed that default rules of the form $\gamma \leftarrow \alpha \;:\; M\beta_1, \ldots, M\beta_n$ can be regarded as AEL inference rules of the form $\gamma \leftarrow \mathbf{B}\alpha,\, \neg\mathbf{B}\neg\beta_1, \ldots \neg\mathbf{B}\neg\beta_n$.

3.2.4 Summary of known complexity properties

A key contribution of [14] is in linking the computational complexity of the decision problems in Table 7 to that of the *derivability problem* in the supporting deductive system $\langle L, R\rangle$, i.e. the computational complexity of deciding given $\Delta \subseteq A$ and $\varphi \in L$ whether $\Delta \models \varphi$. If $(\Delta \models ?\varphi)$ is decidable some class $\mathcal{C}$ then a number of generic upper bounds can be demonstrated in terms of oracle computations provided with access to $\mathcal{C}$ oracles. In total from the complexity bounds on deciding $\Delta \models \varphi$ stated in Table 8 the upper bounds of Table 9 have been derived.

Table 8
Computational Complexity of deciding $\Delta \models \varphi$

Logic	Complexity of deciding $\Delta \models \varphi$
LP	P
DL	coNP–complete
AEL	coNP–complete

Although [14] does not address *lower* bounds for the verification problems, all of the bounds for the credulous and sceptical reasoning problems are shown to be

Table 9
Upper bounds for Decision Problems in ABFs from [14]

Problem	LP	DL	AEL
$\text{VER}_{\text{ADM}}^{\langle L,R\rangle}$	P	P^{NP}	Π_2^p
$\text{VER}_{\text{PR}}^{\langle L,R\rangle}$	coNP	Π_2^p	Π_3^p
$\text{VER}_{\text{ST}}^{\langle L,R\rangle}$	P	P^{NP}	Π_2^p
$\text{CA}_{\text{PR}}^{\langle L,R\rangle}$	NP	Σ_2^p	Σ_3^p
$\text{CA}_{\text{ST}}^{\langle L,R\rangle}$	NP	Σ_2^p	Σ_2^p
$\text{SA}_{\text{PR}}^{\langle L,R\rangle}$	Π_2^p	Π_3^p	Π_4^p
$\text{SA}_{\text{ST}}^{\langle L,R\rangle}$	coNP	Π_2^p	Π_2^p

tight, i.e. if $\mathcal{C}$ is an upper bound on $\text{CA}_s^{\langle L,R\rangle}$ or $\text{SA}_s^{\langle L,R\rangle}$ from Table 9 then the decision question is also $\mathcal{C}$–hard.

3.3 Complexity of Ideal Semantics in ABFS

Lemma 1 and Lemma 2 characterise properties of ideal sets and the ideal extension in AFs. Under certain restrictions both of these have counterparts for ABFs.

Lemma 3 *Let* $\mathcal{T} = \langle T, A, ^{-} \rangle$ *be an* ABF *w.r.t. the deductive system* $\langle L, R\rangle$*. If* $\mathcal{T}$ *is a* flat *framework then:*

a. $\forall\, \Delta \subseteq A$, $\text{VER}_{\text{IDL}}^{\langle L,R\rangle}(\mathcal{T}, \Delta)$ *if and only if*

$$\text{VER}_{\text{ADM}}^{\langle L,R\rangle}(\mathcal{T}, \Delta) \;\; \textit{and} \; \forall\, \Gamma \subseteq A \; : \; att(\Gamma, \Delta) \Rightarrow \neg\text{VER}_{\text{ADM}}^{\langle L,R\rangle}(\mathcal{T}, \Gamma)$$

b. Let $\Theta \subseteq A$ *be the ideal extension of* $\mathcal{T}$*. For all* $\alpha \in A$*,* $\alpha \in \Theta$ *if and only if*

$$\forall \Gamma \subseteq A \; : \; att(\Gamma, \alpha) \;\Rightarrow\; [\, \neg\text{VER}_{\text{ADM}}^{\langle L,R\rangle}(\mathcal{T}, \Gamma) \;\; \textit{and} \;\; att(\Theta, \Gamma)\,]$$

Proof: For (a), if Δ is an ideal set it is admissible by definition. Considering any $\Gamma \subseteq A$ for which $att(\Gamma, \Delta)$, were Γ to be admissible it would not be possible for $\Delta \subseteq \Gamma'$ for every preferred extension since a preferred extension containing Γ would fail to be conflict-free. Thus, if Δ is an ideal set then no subset, Γ, attacking it can be admissible. On the other hand let Δ be admissible and no Δ' attacking it be so. Consider any preferred extension, Γ of $\mathcal{T}$: from the fact that Γ is a preferred extension we have $\neg\, att(\Gamma, \Gamma)$; from the premise that Δ is admissible it further holds $\neg\, att(\Delta, \Delta)$. Consider the set $\Gamma \cup \Delta$.[17] We claim that $\neg att(\Gamma \cup \Delta, \Gamma \cup \Delta)$.

[17] It is this part of the argument that requires $\mathcal{T}$ to be a *flat* framework: otherwise, even if Γ and Δ are *both* closed, we cannot infer that $\Gamma \cup \Delta$ will also be so.

Suppose this were not so, i.e. $att(\Gamma \cup \Delta, \Gamma \cup \Delta)$. Then either $att(\Gamma \cup \Delta, \Delta)$ or $att(\Gamma \cup \Delta, \Gamma)$. In the former case

$$\begin{array}{rcll} att(\Gamma \cup \Delta, \Delta) & \Rightarrow & att(\Delta, \Gamma \cup \Delta) & \text{since } \Delta \text{ is admissible} \\ & \Rightarrow & att(\Delta, \Gamma) & \text{since } \neg att(\Delta, \Delta) \\ & \Rightarrow & att(\Gamma, \Delta) & \text{since } \Gamma \text{ is admissible} \\ & \Rightarrow & \neg \ \mathrm{VER}^{\langle L,R \rangle}_{\mathrm{ADM}}(\mathcal{T}, \Gamma) & \text{from premise that no attacker of } \Delta \text{ is admissible} \end{array}$$

This contradicts the choice of Γ as a preferred extension.

In the latter case,

$$\begin{array}{rcll} att(\Gamma \cup \Delta, \Gamma) & \Rightarrow & att(\Gamma, \Gamma \cup \Delta) & \text{since } \Gamma \text{ is admissible} \\ & \Rightarrow & att(\Gamma, \Delta) & \text{since } \neg att(\Gamma, \Gamma) \\ & \Rightarrow & \neg \ \mathrm{VER}^{\langle L,R \rangle}_{\mathrm{ADM}}(\mathcal{T}, \Gamma) & \text{from premise that no attacker of } \Delta \text{ is admissible} \end{array}$$

again contradicting the choice of Γ as a preferred extension.

It follows that $\Gamma \cup \Delta$ is conflict-free. This set, however, is also admissible:

$$att(\Gamma', \Delta \cup \Gamma) \ \Rightarrow \ [att(\Gamma', \Gamma) \ \text{ or } att(\Gamma', \Delta)] \ \Rightarrow \ att(\Gamma \cup \Delta, \Gamma')$$

Since Γ is a preferred extension $\Gamma \cup \Delta \ = \ \Gamma$, i.e. $\Delta \subseteq \Gamma$ and we deduce that Δ is an ideal set.

For (b), if $\alpha \in \Theta$, then it is immediate from part (a) that no $\Gamma \subseteq A$ for which $att(\Gamma, \alpha)$ is admissible. In addition, since Θ is admissible, $att(\Gamma, \alpha) \Rightarrow att(\Gamma, \Theta)$ thus $att(\Theta, \Gamma)$. Conversely, if no attacker, Γ of α is admissible and every such attacker is attacked by Θ then $\Theta \cup \{\alpha\}$ is conflict-free. To see this, assume the contrary and that $att(\Theta \cup \{\alpha\}, \Theta \cup \{\alpha\})$. Either $att(\Theta \cup \{\alpha\}, \{\alpha\})$ or $att(\Theta \cup \{\alpha\}, \Theta)$. In the first case we have,

$$\begin{array}{rcll} att(\Theta \cup \{\alpha\}, \{\alpha\}) & \Rightarrow & att(\Theta, \Theta \cup \{\alpha\}) & \text{from premise} \\ & \Rightarrow & att(\Theta, \{\alpha\}) & \text{since } \neg att(\Theta, \Theta) \text{ from admissibility of } \Theta \\ & \Rightarrow & \neg \mathrm{VER}^{\langle L,R \rangle}_{\mathrm{ADM}}(\mathcal{T}, \Theta) & \text{from premise} \end{array}$$

contradicting Θ being the ideal extension. In the second case, since

$$att(\Theta \cup \{\alpha\}, \Theta) \Rightarrow att(\Theta, \Theta \cup \{\alpha\})$$

we derive a similar contradiction. It is easy to see that $\Theta \cup \{\alpha\}$ is also admissible since every set attacking it is counterattacked by Θ. Furthermore no attacker Γ of $\Theta \cup \{\alpha\}$ is admissible: either $att(\Gamma, \Theta)$ and Γ is not an admissible set via (a); or $att(\Gamma, \{\alpha\})$ and Γ is not admissible from the premises on α. We deduce that $\Theta \cup \{\alpha\}$ is an ideal set, and with Θ being the ideal extension, i.e. maximal such set, so $\alpha \in \Theta$ as claimed. $\square$

Corollary 12 *If $\mathcal{T} = \langle T, A, ^{-} \rangle$ is a flat framework in $\langle L, R\rangle$ and deciding $\Delta \models \varphi$ is in some complexity class $\mathcal{C}$, then $\textsc{ver}_{\text{IDL}}^{\langle L,R\rangle} \in$ co$\text{NP}^{\mathcal{C}}$.*

Proof: Given an instance $\langle \mathcal{T}, \Delta\rangle$ of $\textsc{ver}_{\text{IDL}}^{\langle L,R\rangle}$ we may check $\textsc{ver}_{\text{ADM}}^{\langle L,R\rangle}(\mathcal{T}, \Delta)$ via a $\text{P}^{\mathcal{C}}$ computation, as described in [14, Thm. 4]. We can then test for every $\Gamma \subseteq A$ that should $att(\Gamma, \Delta)$ ($|\Delta|$ calls to a $\mathcal{C}$ oracle deciding $\Gamma \models \overline{\alpha}$ for $\alpha \in \Delta$), then Γ is not admissible ($\text{P}^{\mathcal{C}}$). In total the algorithm is implemented in $\text{coNP}^{\text{P}^{\mathcal{C}}} = \text{coNP}^{\mathcal{C}}$ as claimed. $\square$

Corollary 13 *The problem of verifying that Δ is an ideal set can be decided in*

a. coNP *for* LP *instantiations of* ABFs*.*
b. Π_2^p *for* DL *instantiations of* ABFs*.*

Proof: All instances of ABFs instantiating LP or DL describe flat frameworks. Thus both bounds follow from Corollary 12 and Table 8 giving $\text{coNP}^{\text{P}} = \text{coNP}$ for LP instances, and $\text{coNP}^{\text{coNP}} = \text{coNP}^{\text{NP}} = \Pi_2^p$ for DL instances. $\square$

In contrast to the upper bounds given in Corollary 13, where it is *not* possible to assume flatness, the characterisation from Lemma 3 may fail. In such cases one has the general upper bound,

Lemma 4 *Let $\mathcal{T} = \langle T, A, ^{-} \rangle$ be an* ABF *with underlying deductive system $\langle L, R\rangle$ and for which $\Delta \models \varphi$ is decidable in $\mathcal{C}$.*

$$\textsc{ver}_{\text{IDL}}^{\langle L,R\rangle} \in co\text{NP}^{\text{NP}^{\text{NP}^{\mathcal{C}}}}$$

Proof: Given an instance $\langle \mathcal{T}, \Delta\rangle$ of $\textsc{ver}_{\text{IDL}}^{\langle L,R\rangle}$ for $\mathcal{T} = \langle T, A, ^{-} \rangle$ as in the Lemma statement, this can be decided by checking $\textsc{ver}_{\text{ADM}}^{\langle L,R\rangle}(\mathcal{T}, \Delta)$ ($\text{coNP}^{\mathcal{C}}$, via [14, Thm. 3]) and then testing

$$\forall\, \Gamma \subseteq A \;\; (\neg \textsc{ver}_{\text{PR}}^{\langle L,R\rangle}(\mathcal{T}, \Gamma) \;\vee\; \Delta \subseteq \Gamma)$$

That is, every preferred extension of $\mathcal{T}$ contains Δ. Again via, [14, Thm. 3], $\neg\textsc{ver}_{\text{PR}}^{\langle L,R\rangle} \in \textsc{np}^{\text{NP}^c}$ so the test described can be completed in $\text{co}\textsc{np}^{\text{NP}^{\text{NP}^c}}$ as claimed. ▫

Corollary 14 *For* AEL *instantiations of* ABF*s the problem of verifying that Δ is an ideal set is in Π_4^p.*

Proof: Those frameworks describing AEL instantiations can fail to be flat, hence the upper bound is immediate from Table 8 and Lemma 4. ▫

The characterisation of ideal sets and properties of the ideal extension are not the only properties of AFs that carry across to flat frameworks. With some minor variations it turns out that the *construction* process for building the ideal extension – described in Thm. 7 – can also be adapted to flat frameworks. We first describe the algorithm for flat frameworks in Algorithm 1 and prove its correctness in Thm. 16. Finally its run-time is analysed in terms of upper bounds on $\textsc{ca}_{\text{ADM}}^{\langle L,R\rangle}$ and that of deciding $\Delta \models \varphi$ in Thm. 17.

Algorithm 1 Construction of the Ideal Extension in flat ABFs

1: **function** FIND-IDEAL-EXTENSION ($\mathcal{T} = \langle T, A, ^- \rangle$)
2: $A_{out} := \{\alpha \in A : \neg\textsc{ca}_{\text{ADM}}^{\langle L,R\rangle}(\mathcal{T}, \alpha)\}$;
3: $A_{in} := A \setminus A_{out}$;
4: $A_{\text{CA}} := \{ \alpha \in A_{in} : A_{in} \models \overline{\alpha} \}$;
5: $A_{\text{PSA}} := A_{in} \setminus A_{\text{CA}}$;
6: $\Gamma := A_{\text{PSA}}$;
7: **repeat**
8: $\quad\Gamma_{in} := \Gamma$;
9: $\quad\Xi := \{ \alpha \in A_{out} : \neg (\Gamma \models \overline{\alpha}) \}$;
10: $\quad\Delta := \{ \gamma \in \Gamma_{in} : \Xi \cup A_{\text{CA}} \models \overline{\gamma} \}$;
11: $\quad\Gamma := \Gamma \setminus \Delta$;
12: **until** $\Gamma_{in} = \Gamma$
13: **return** Γ;

Prior formally to proving its correctness, some discussion of this algorithm may be helpful. In a similar manner to the mechanism described in obtaining the upper bound of Thm. 7, the algorithm builds a partition of the assumption set into three parts: A_{out} (the counterpart of $\mathcal{X}_{\text{OUT}}$ from Thm. 7) with the remaining assumptions (A_{in}) divided between those which cannot be sceptically accepted (the set A_{CA}) and those which *could* be sceptically accepted (the set A_{PSA}). The main loop, between ll. 7–12, progressively removes assumptions from A_{PSA} (until no change results), by identifying those assumptions which *cannot* be part of the ideal extension. The computational process, in effect, mirrors the the view of $\langle \mathcal{X}_{\text{PSA}}, \mathcal{X}_{\text{OUT}} \rangle$ as a bipartite subgraph of $\mathcal{H}$ adopted in Thm. 7 in its treatment of $\langle A_{\text{PSA}}, A_{out} \rangle$. The major difference is that the attack relation must be considered in terms of *sets* of assumptions, whereas in the AF algorithm it sufficed to deal with the interaction between individual arguments in $\mathcal{X}_{\text{OUT}}$ and arguments in $\mathcal{X}_{\text{PSA}}$.[18] Informally one may view the rationale of Algorithm 1 as implicitly considering a bipartite AF formed by $\langle A_{\text{PSA}}, 2^{A_{out}} \rangle$ with the attack relation containing $\langle \Delta, \alpha \rangle$ for Δ any (minimal) subset of A_{out} for which $\Delta \cup A_{\text{CA}} \models \overline{\alpha}$; together with $\langle A_{\text{PSA}}, \beta \rangle$ whenever (the current, in the sense of l. 8), $A_{\text{PSA}} \models \overline{\beta}$.

Theorem 16 *Given a flat framework, $\mathcal{T}$, Algorithm 1 returns its ideal extension.*

Proof: For $\mathcal{T} = \langle T, A, \overline{\ \ } \rangle$ a flat framework w.r.t. the deductive system $\langle L, R \rangle$ let $\Theta \subseteq A$ be its ideal extension. First observe that the subset A_{PSA} computed in line 5 of the algorithm is both conflict-free and is such that $\Theta \subseteq A_{\text{PSA}}$. To see this notice that $att(A_{\text{PSA}}, A_{\text{PSA}})$ would imply $A_{\text{PSA}} \models \overline{\gamma}$ and, hence, $A_{in} \models \overline{\gamma}$ contradicting $\gamma \in A_{\text{PSA}}$ (cf. lines 4–5). That $\Theta \subseteq A_{\text{PSA}}$, follows by observing $\Theta \cap A_{out} = \emptyset$: Θ is admissible, however, no assumption in A_{out} belongs to an admissible set (line 2). Supposing, to the contrary, that $\Theta \cap A_{\text{CA}} \neq \emptyset$ consider any $\alpha \in \Theta \cap A_{\text{CA}}$. By definition, $\alpha \in A_{\text{CA}}$ implies $A_{in} \models \overline{\alpha}$, hence $att(A_{in}, \Theta)$ from which $att(\Theta, A_{in})$ as Θ is admissible. We, therefore, can find some $\delta \in A_{in}$ such that $\Theta \models \overline{\delta}$ so that Θ attacks any preferred extension, Δ of $\mathcal{T}$ for which $\delta \in \Delta$. Such a preferred extension exists by virtue of $\text{CA}^{\langle L,R \rangle}_{\text{ADM}}(\mathcal{T}, \delta)$ but then we cannot have $\Theta \subseteq \Delta$ contradicting the choice of Θ as the ideal extension of $\mathcal{T}$. In summary A_{PSA} is conflict-free and $\Theta \subseteq A_{\text{PSA}}$. In addition $\neg att(A_{\text{PSA}}, A_{\text{CA}})$: if $A_{\text{PSA}} \models \overline{\gamma}$ for $\gamma \in A_{\text{CA}}$, then $att(A_{\text{PSA}}, \Gamma)$ for any preferred extension Γ with $\gamma \in \Gamma$ leading to the contradiction $att(A_{in}, A_{\text{PSA}})$ as a consequence of $att(\Gamma, A_{\text{PSA}})$.

To complete the proof of correctness it remains to show that the set Γ returned by Algorithm 1 is the maximal admissible subset of A_{PSA}. Certainly Γ is conflict-free (since A_{PSA} is conflict-free). Consider any $\Xi \subseteq A$ for which $att(\Xi, \Gamma)$. It must be the case that $\Xi \cap A_{out} \neq \emptyset$ for otherwise $\Xi \subseteq A_{\text{CA}}$ and $att(\Xi, \Gamma)$ would yield $\Xi \models \overline{\gamma}$ for $\gamma \in \Gamma$, i.e. $A_{in} \models \overline{\gamma}$ contradicting $\gamma \in A_{\text{PSA}}$. Let $\Xi_{out} = \Xi \cap A_{out}$. If $att(\Gamma, \Xi_{out})$ then, trivially, $att(\Gamma, \Xi)$, so Γ would only fail to be admissible if it is attacked by Ξ such that

$$\forall\, \xi \in \Xi_{out} \ \neg(\Gamma \models \overline{\xi})$$

This, however, contradicts Γ being the set returned by Algorithm 1: for consider the subset Δ for which $\Xi \models \overline{\delta}$ for each $\delta \in \Delta$ (from $att(\Xi, \Gamma)$ it is immediate that $|\Delta| \geq 1$). This is precisely the subset identified in line 10 and removed (from the current Γ) in line 11. We deduce, as a result that Γ is an admissible subset of

[18] The polynomial time algorithm of [21] identifying non-credulously accepted arguments in bipartite AFs, $\mathcal{B}(\mathcal{Y}, \mathcal{Z}, \mathcal{A})$ does so by repeatedly identifying those $y \in \mathcal{Y}$ (resp. $z \in \mathcal{Z}$) for which some $z \in \{y\}^-$ (resp. $y \in \{z\}^-$) is unattacked by $\mathcal{Y}$ (resp. $\mathcal{Z}$), cf. the construction in ll. 9–10 of Algorithm 1: line 9 identifies the set Ξ of assumptions in A_{out} that are unattacked (by A_{PSA}) and eliminates from A_{PSA} those assumptions attacked by Ξ.

A_{PSA}. It must, however, also be a *maximal* such set. For consider any non-empty $\Delta \subseteq A_{\text{PSA}} \setminus \Gamma$. We claim the set $\Gamma \cup \Delta$ is not admissible. To see this let

$$\Gamma_0,\ \Gamma_1\ , \ldots,\ \Gamma_k$$

be the sequence of sets $\Gamma_0 = A_{\text{PSA}}$, $\Gamma_k = \Gamma$, and $\Gamma_i \subset \Gamma_{i-1}$ ($\leq i \leq k$) over successive iterations of lines 7–11. From $\Delta \cap \Gamma = \emptyset$, we can identify a partition of Δ into r sets $\langle \Delta_1, \ldots, \Delta_r \rangle$ and a subsequence $\langle j_1, j_2, \ldots, j_r \rangle$ of $\langle 0, 1, \ldots, k-1 \rangle$ such that

$$\Delta_i \subseteq \Gamma_{j_i} \quad ; \quad \Delta_i \cap \Gamma_{j_{i+1}} = \emptyset$$

Thus, Δ_i is a subset of those assumptions removed from the current collection $\Gamma_{in} = \Gamma_{j_i}$ in line 11 of the algorithm. Without loss of generality we can focus on Δ_1 and Γ_{j_1}. By inspection we see that,

$$\begin{array}{lcll} \Xi & = & \{\alpha \in A_{out} \ : \ \neg(\Gamma_{j_1} \models \overline{\alpha})\} & \text{line 9} \\ \Delta_1 & \subseteq & \{\gamma \in \Gamma_{j_1} \ : \ \Xi \ \cup \ A_{\text{CA}} \models \overline{\gamma}\} & \text{line 10} \end{array}$$

Noting that $\Gamma \cup \Delta_1 \subseteq \Gamma_{j_1}$ and that

$$\Xi \ \subseteq \ \Xi' = \{\alpha \in A_{out} \ : \ \neg(\Gamma_{j_1} \models \overline{\alpha})\}$$

it follows that $att(A_{\text{CA}} \cup \Xi', \Gamma \cup \Delta_1)$ and, in particular, $att(A_{\text{CA}} \cup \Xi', \Delta_1)$. In summary, if it is the case that $\Gamma \cup \Delta$ *is* admissible, then this set must be able to counter the attack on Δ_1 by $A_{\text{CA}} \cup \Xi'$. Hence,

$$\begin{array}{rcl} \text{VER}_{\text{ADM}}^{\langle L,R \rangle}(\mathcal{T}, \Gamma \cup \Delta) & \Rightarrow & att(\Gamma \cup \Delta, A_{\text{CA}} \cup \Xi') \\ & \Rightarrow & att(\Gamma \cup \Delta, A_{\text{CA}}) \ \text{ or } att(\Gamma \cup \Delta, \Xi') \end{array}$$

From $\Gamma \cup \Delta \subseteq A_{\text{PSA}}$ and $\neg att(A_{\text{PSA}}, A_{\text{CA}})$ the only possibility is $att(\Gamma \cup \Delta, \Xi')$: the definition of Γ_{j_1} which is a superset of $\Gamma \cup \Delta$ shows that no such counterattack is possible. In total we deduce that $\Gamma \cup \Delta$ cannot be admissible and, hence Γ is the *maximal* admissible subset of A_{PSA}. It is easily seen that that no attacker of Γ defines an admissible set of assumptions since $att(\Delta, \Gamma) \Rightarrow \Delta \cap A_{out} \neq \emptyset$.

To summarise: Θ, the ideal extension, is a subset of A_{PSA}; Γ the set returned by Algorithm 1 is an admissible subset of A_{PSA} and no attacker, Δ of Γ is admissible, i.e by Lemma 3 (a) is an ideal set; Γ, however, is the *maximal* admissible subset of A_{PSA} none of whose attackers is admissible, so that $\Gamma = \Theta$ as claimed. □

Theorem 17 *Let $\mathcal{T} = \langle T, A, ^{-} \rangle$ be a flat framework w.r.t. $\langle L, R \rangle$ for which $\Delta \models \varphi$ is decidable in $\mathcal{C}$. Using Algorithm 1 the maximal ideal extension of $\mathcal{T}$ may be found in $\mathrm{FP}_{||}^{\mathrm{NP}^{\mathcal{C}}}$, i.e. the class of function problems that are solvable by polynomial time algorithms which make non-adaptive queries to an oracle in $\mathrm{NP}^{\mathcal{C}}$.*

Proof: The partition of $A = \{\alpha_1, \ldots, \alpha_n\}$ into $\langle A_{out},\ A_{in} \rangle$ can be obtained using the single parallel query which reports $z_1 z_2 \ldots z_n \in \langle \top, \bot \rangle^n$ with $z_i = \top \Leftrightarrow \mathrm{CA}_{\mathrm{ADM}}^{\langle L,R \rangle}(\mathcal{T}, \alpha_i)$. From [14, Thm. 8], in the case of flat frameworks with $\Delta \models \varphi$ decidable in $\mathcal{C}$, $\mathrm{CA}_{\mathrm{ADM}}^{\langle L,R \rangle} \in \mathrm{NP}^{\mathcal{C}}$ so that this partition is constructible in $\mathrm{FP}_{||}^{\mathrm{NP}^{\mathcal{C}}}$. The remaining stages of the algorithm require only a polynomial number of adaptive queries to $\mathcal{C}$ oracles: an easy generalisation of [27, Thm. 2.2, p. 379] gives $\mathrm{FP}^{\mathrm{NP}^{\mathcal{C}}[\log]} = \mathrm{FP}_{||}^{\mathrm{NP}^{\mathcal{C}}}$, (where $\mathrm{FP}^{\mathcal{D}[\log]}$ is the class of functions computable by polynomial time algorithms that make $O(\log n)$ queries to a $\mathcal{D}$ oracle on instances of size n), so that since $\mathrm{FP}^{\mathcal{C}} \subseteq \mathrm{FP}^{\mathrm{NP}^{\mathcal{C}}}$ the overall upper bound stated follows. □

Turning to the specific cases LP and DL, the following corollaries are immediate from Thm. 17 and Table 8.

Corollary 15

a. $\mathrm{FIE}^{\mathrm{LP}} \in \mathrm{FP}_{||}^{\mathrm{NP}}$.

b. $\mathrm{FIE}^{\mathrm{DL}} \in \mathrm{FP}_{||}^{\Sigma_2^p}$.

Corollary 16

a. The decision problems $\mathrm{VER}_{\mathrm{IE}}^{\mathrm{LP}}$, $\mathrm{VER}_{\mathrm{IE}}^{\mathrm{LP},\emptyset}$ *and* $\mathrm{CA}_{\mathrm{IDL}}^{\mathrm{LP}}$ *are all in* $\mathrm{P}_{||}^{\mathrm{NP}}$.

b. The decision problems $\mathrm{VER}_{\mathrm{IE}}^{\mathrm{DL}}$, $\mathrm{VER}_{\mathrm{IE}}^{\mathrm{DL},\emptyset}$ *and* $\mathrm{CA}_{\mathrm{IDL}}^{\mathrm{DL}}$ *are all in* $\mathrm{P}_{||}^{\Sigma_2^p}$.

The results of Corollaries 13, 15, and 16 establish upper bounds on each of the three variants of the verification problem, credulous reasoning, and the complexity of constructing ideal extensions. These upper bounds rely on properties of the supporting deductive system and the complexity of deciding $\Delta \models \varphi$. In order to address issues of lower bounds and the complexity of the other decision problems we consider the different deductive systems in turn.

Regarding LP instantiations we have the following theorem.

Theorem 18

a. $\mathrm{VER}_{\mathrm{IDL}}^{\mathrm{LP}}$ *is co*NP*–hard via* $\leq_m^p$.

b. $\mathrm{CA}_{\mathrm{IDL}}^{\mathrm{LP}}$ *is* $\mathrm{P}_{||}^{\mathrm{NP}}$*–hard via* $\leq_m^{rp}$ *with probability* $1 - 2^{-n}$.

c. $\mathrm{VER}_{\mathrm{IE}}^{\mathrm{LP}}$ *is* $\mathrm{P}_{||}^{\mathrm{NP}}$*–hard via* $\leq_m^{rp}$ *with probability* $1 - 2^{-n}$.

d. $\mathrm{VER}_{\mathrm{IE}}^{\mathrm{LP},\emptyset}$ *is* $\mathrm{P}_{||}^{\mathrm{NP}}$*–hard via* $\leq_m^{rp}$ *with probability* $1 - 2^{-n}$.

e. $\mathrm{FIE}^{\mathrm{LP}}$ *is* $\mathrm{FP}_{||}^{\mathrm{NP}}$*–hard via* $\leq_m^p$.

Proof: All of the lower bounds follow by giving a translation from arbitrary AFs $\mathcal{H}(\mathcal{X}, \mathcal{A})$ to a related LP setting, i.e. in effect, if L_{ID} is a problem defined in the ideal semantics for AFs with $L_{\mathrm{ID}}^{\mathrm{LP}}$ its counterpart in ABFs instantiating LP forms, then the translation we describe forms the basis of a proof that $L_{\mathrm{ID}} \leq_m^p L_{\mathrm{ID}}^{\mathrm{LP}}$.

The translation we describe from $\mathcal{H}(\mathcal{X}, \mathcal{A})$ to a logic program $T^{\mathcal{H}}$ is effectively that given in Dung [17, p. 348]. For an AF, $\mathcal{H}(\mathcal{X}, \mathcal{A})$ define $Lits^{\mathcal{H}}$ as the set of ground atoms

$$Lits^{\mathcal{H}} \;=\; \{\, d(x) \;:\; x \in \mathcal{X}\} \;\bigcup\; \{\, \mathit{not}\; d(x) \;:\; x \in \mathcal{X}\}$$

(although these sets are described as unary functions, for any fixed $\langle \mathcal{X}, \mathcal{A} \rangle$, $Lits^{\mathcal{H}}$ defines a set of $2|\mathcal{X}|$ propositional variables.)

The logic program, $T^{\mathcal{H}}$ has exactly the following rules:

$$T^{\mathcal{H}} \;=\; \bigcup_{x \in \mathcal{X}} \;\bigcup_{y \in \{x\}^-} \;\{\, d(x) \;\leftarrow\; \mathit{not}\; d(y)\,\}$$

Informally these assert that the argument x is "defeated" if any of its attackers ($y \in \{x\}^-$) is assumed *not* to be defeated. For the deductive system so defined we have the the ABF, $\mathcal{T}^{\mathcal{H}} \;=\; \langle T^{\mathcal{H}}, A^{\mathcal{H}}, ^{-} \rangle$ in which $A^{\mathcal{H}} \;=\; \{\, \mathit{not}\; d(x) \;:\; x \in \mathcal{X}\}$, $\overline{\mathit{not}\; d(x)} \;=\; d(x)$.

Rather than derive (a)–(e) separately, it is, in fact sufficient to prove there is a one-to-one correspondence between ideal sets (of *arguments*) in $\mathcal{H}(\mathcal{X}, \mathcal{A})$ and ideal sets (of *assumptions*) in $\mathcal{T}^{\mathcal{H}}$. For $S \subseteq \mathcal{X}$, $\Delta(S) \subseteq A^{\mathcal{H}}$ is the set of assumptions $\{\mathit{not}\; d(y) \;:\; y \in S\}$. Similarly, for $\Delta \subseteq A^{\mathcal{H}}$, $S(\Delta) \subseteq \mathcal{X}$ is the set of arguments $\{\, y \;:\; \mathit{not}\; d(y) \in \Delta\}$. It is easy to see that (a)–(e) are all immediate consequences of the following

$$S \text{ is an ideal set in } \mathcal{H} \text{ if and only if } \Delta(S) \text{ is an ideal set in } \mathcal{T}^{\mathcal{H}} \tag{2}$$

In order to establish (2), first suppose that S defines an ideal set within $\mathcal{H}$ and consider $\Delta(S) \subseteq A^{\mathcal{H}}$. The set $\Delta(S)$ is conflict-free since

$$\begin{aligned} att(\Delta(S),\Delta(S)) \;&\Leftrightarrow\; \{not\ d(x),\ not\ d(y)\} \subseteq \Delta(S) \\ &\qquad \text{and } (d(x) \leftarrow not\ d(y) \in T^{\mathcal{H}} \text{ or } d(y) \leftarrow not\ d(x) \in T^{\mathcal{H}}) \\ &\Leftrightarrow\; \{x,y\} \subseteq S \text{ and } (\langle y,x\rangle \in \mathcal{A} \text{ or } \langle x,y\rangle \in \mathcal{A}) \end{aligned}$$

contradicting the fact that S is conflict-free. Similarly $\Delta(S)$ must be admissible, for given any $\Gamma \subseteq A^{\mathcal{T}}$, $att(\Gamma,\Delta(S))$ if and only if $\Gamma \models \overline{not\ d(x)} = d(x)$ for some $not\ d(x) \in \Delta(S)$. From the definition of $T^{\mathcal{H}}$, Γ contains an assumption $not\ d(y)$ for which $d(x) \leftarrow not\ d(y)$ is a rule in $T^{\mathcal{H}}$, so that $\langle y,x\rangle \in \mathcal{A}$, i.e. $y \in \mathcal{X}$ attacks x in S. It follows that there is some $z \in S$ for which $\langle z,y\rangle \in \mathcal{A}$, hence some $not\ d(z) \in \Delta(S)$ for which $d(y) \leftarrow not\ d(z)$ is in $T^{\mathcal{H}}$, i.e. $\Delta_S \models \overline{not\ d(y)}$ and $att(\Delta(S),\Gamma)$. Thus $\Delta(S)$ is also admissible. We recall from Lemma 1 that, since S is an ideal set, no argument in S^- is credulously accepted. We show that, in consequence, no subset Γ of assumptions for which $att(\Gamma,\Delta(S))$ is admissible. This suffices to complete the first part of the proof via Lemma 3(a). So suppose $att(\Gamma,\Delta(S))$ and hence $\Gamma \models \overline{not\ d(x)}$ for $not\ d(x) \in \Delta(S)$, i.e. $not\ d(y) \in \Gamma$ and $d(x) \leftarrow not\ d(y) \in T^{\mathcal{H}}$ so that $\langle y,x\rangle \in \mathcal{A}$. It is easy to show, however, that if Γ were admissible then $S(\Gamma) = \{\, y \,:\, not\ d(y) \in \Gamma\} \subseteq \mathcal{X}$ would be admissible in $\mathcal{H}$: as $y \in S(\Gamma) \cap S^-$, this contradicts the fact that no attacker of S is credulously accepted. We deduce that if S is an ideal set in $\mathcal{H}$ then $\Delta(S)$ is an ideal set in $\mathcal{T}^{\mathcal{H}}$.

For the converse implication, let Δ be an ideal set of $\mathcal{T}^{\mathcal{H}}$ and consider the subset $S(\Delta)$ of $\mathcal{X}$. By similar arguments to those above $S(\Delta)$ is an admissible set since Δ is admissible. Now consider any attacker y of some $x \in S(\Delta)$. In $\mathcal{T}^{\mathcal{H}}$ any set of assumptions, Γ containing $not\ d(y)$ attacks $not\ d(x) \in \Delta$ by virtue of the rule $d(x) \leftarrow not\ d(y)$ in $T^{\mathcal{H}}$, thus, from Lemma 3(a) it follows that no such set can be admissible in $\mathcal{T}^{\mathcal{H}}$, i.e. the argument y is not credulously accepted in $\mathcal{H}$. Thus $S(\Delta)$ is an admissible set of $\mathcal{H}$ none of whose attackers is credulously accepted: from Lemma 1, $S(\Delta)$ is therefore an ideal set. □

Readers familiar with [15] may recognise that the form of $T^{\mathcal{H}}$ in the proof just given is that of a "reduced negative logic program" the structure employed in order to derive earlier complexity results in AFs. In principle we could also have derived *upper* bounds for problems on LP instantiations of ABFs, by transforming normal logic programs to reduced negative logic programs and invoking the "rule graph" construction of [15, Defn. 3.8, p. 219]. There are several reasons why we have *not* adopted such an approach: although translations from normal logic programs to reduced negative logic programs can always be carried out, e.g. [15, Propn. 4.2], there is some debate about how efficiently such processes can be effected: Linke [30] claims an exponential size increase may occur, an assertion disputed in Costantini *et al.* [9] which outlines a polynomial time process building on an algorithm of [6].

Irrespective of how efficiently transformations to negative logic programs can be performed, the upper bound mechanisms presented earlier offer a more general approach encompassing all flat frameworks, i.e. not simply LP cases. The exact form of the lower bound translation (and the fact the the upper bound construction is unconstrained) yield the following observation.

Corollary 17 *The lower bounds of Thm. 18 continue to hold even in the case of* LP *theories, T, in which* every *clause of T has the form $\alpha \leftarrow not\ \beta$, i.e. with a single atom in $\mathcal{HB}_{not}$ defining the premises of each rule.*

We now consider lower bounds for ABFs instantiating default logics. The proof methods are built on techniques from work of Gottlob [26] (which also feature in the treatment of default logics in [14]), the characterisation of ideal sets from Lemma 3, and the approach adopted in the proof of Thm 2.

Theorem 19 $\textsc{ver}^{\mathrm{DL}}_{\mathrm{IDL}}$ *is Π_2^p–complete.*

Proof: That $\textsc{ver}^{\mathrm{DL}}_{\mathrm{IDL}} \in \Pi_2^p$ has already been shown in Corollary 13 (b). To show that this problem is Π_2^p–hard, we reduce to the complementary problem, $\neg\ \textsc{ver}^{\mathrm{DL}}_{\mathrm{IDL}}$ from the Σ_2^p–complete problem $\textsc{qsat}_2^{\Sigma}$. We assume instances are 3-DNF formulae with product terms $\{P_1, P_2, \ldots, P_m\}$ each comprising exactly three literals defined over two (disjoint) sets of propositional variables $Y_n = \langle y_1, \ldots, y_n\rangle$; $Z_n = \langle z_1, \ldots, z_n\rangle$, an instance $\varphi(Y_n, Z_n)$ being accepted if there is some instantiation, α_Y, of Y_n, under which $\Phi(\alpha, Z_n) \equiv \top$.

Given, $\varphi(Y_n, Z_n)$ form the default theory (W, D) over the language whose literals are

$$L_0 = \{\ y_i,\ z_i,\ \neg\, y_i, \neg\, z_i,\ :\ 1 \le i \le n\} \ \cup\ \{\varphi, \neg\varphi, \psi, \neg\psi\}$$

[**Note:** The terms $\{\varphi, \neg\varphi, \psi, \neg\psi\}$ are treated as *literals* in this language, so that φ is effectively a "place-holder" for the DNF expression $P_1 \vee \cdots \vee P_m$ over literals of $Z_n \cup Y_n$, e.g. as in the second set of default rules below.]

We fix $W = \emptyset$ and D to contain the following default rules: [19]

$$\left\{ \frac{\top\ :\ My_i}{y_i}, \frac{\top\ :\ M\neg y_i}{\neg y_i}, \frac{\top\ :\ M\psi}{y_i}, \frac{\top\ :\ M\psi}{\neg\, y_i}\ :\ 1 \le i \le n \right\}$$

$$\left\{ \frac{\top\ :\ M\psi}{\neg\varphi}, \frac{\top\ :\ M\varphi}{\neg\psi}, \frac{\top\ :\ M(\neg P_1 \wedge \neg P_2 \wedge \cdots \wedge \neg P_m)}{\neg\varphi} \right\}$$

[19] Readers familiar with default logic will note that these do not define a *normal* set of defaults. There are, however, general translation mechanisms, cf. [26, p. 414], that can be used to build equivalent *semi-normal* default theories from arbitrary defaults. In order to minimise notational overheads we have eschewed such translation in our presentation.

The ABF instance formed from (W, D) is $\mathcal{T}_\varphi = \langle \emptyset, A_\varphi, ^- \rangle$ with

$$A_\varphi = \{ My_i,\ M\neg y_i \ :\ 1 \leq i \leq n\} \cup \{M\varphi,\ M\psi\} \cup \{ M(\neg P_1 \wedge \neg P_2 \wedge \cdots \wedge \neg P_m) \}$$

We recall that the contrary mapping is $\overline{M\alpha} = \neg\alpha$.

We claim that $\{M\psi\}$ does *not* define an ideal set within $\mathcal{T}_\varphi$ if and only if $\varphi(Y_n, Z_n)$ is a positive instance of QSAT_2^Σ. We first observe that $\{M\psi\}$ defines an admissible set of assumptions from A_φ. For consider any $\Delta \subset A_\varphi$ for which $att(\Delta, \{M\psi\})$. This can only happen if $\Delta \models \neg\psi$, from which it follows that $M\varphi \in \Delta$ and hence (since $\{M\psi\} \models \neg\varphi$) it follows that $att(\{M\psi\}, \Delta)$. Given that $\{M\psi\}$ is admissible, in order for it to fail to be an ideal set, via Lemma 3 (a), some $\Delta \subseteq A_\varphi$ for which $att(\Delta, M\psi)$ must define an admissible set within $\mathcal{T}$. We have aleady argued that any set attacking $\{M\psi\}$ must contain the assumption $M\varphi$ so it suffices to show that $\text{CA}_{\text{ADM}}^{\text{DL}}(\mathcal{T}, M\varphi)$ if and only if $\exists\alpha_Y \ :\ \varphi(\alpha_Y, Z_n) \equiv \top$.

Suppose first that $\text{CA}_{\text{ADM}}^{\text{DL}}(\mathcal{T}, M\varphi)$ and let $\Gamma \subseteq A_\varphi$ be a preferred extension of $\mathcal{T}$. for which $M\varphi \in \Gamma$. Certainly if $M\psi \in \Delta$ then $att(\Delta, \Gamma)$, however from $\{M\varphi\} \models \neg\psi$ such attacks are countered. From $M\varphi \in \Gamma$, any set of assumptions, Δ, for which $M(\neg P_1 \wedge \neg P_2 \wedge \cdots \wedge \neg P_m) \in \Delta$ will also attack Γ so that $\Gamma \models \neg(\neg P_1 \wedge \cdots \wedge \neg P_m) \equiv P_1 \vee \cdots \vee P_m$. In consequence, Γ must contain n assumptions (one for each pair $\{My_i, M\neg y_i\}$), Γ', such that $\Gamma' \models\ P_1 \vee \cdots \vee P_m$. It is easily seen that choosing $\alpha_Y = \langle a_1, \ldots, a_n \rangle$ to be the corresponding instantiation of Y_n, i.e. $a_i := \top \ \Leftrightarrow\ My_i \in \Gamma'$ gives $\varphi(\alpha_Y, Z_n) \equiv \top$.

For the converse implication supposing that some α_Y satisfies $\varphi(\alpha_Y, Z_n) \equiv \top$. Consider the set of assumptions $\Gamma = \{ My_i \ :\ a_i = \top\} \cup \{M\neg y_i \ :\ a_i = \perp\} \cup \{M\varphi\}$. It is certainly the case that $\neg att(\Gamma, \Gamma)$. Furthermore the only Δ attacking Γ are those containing $M\psi$ (which is counterattacked through $\{M\varphi\} \models \neg\psi$), $M\neg y_i$ (for those y_i with $a_i = \top$: counterattacked via $\{My_i\} \models y_i$); My_i (for those y_i with $a_i = \perp$: counterattacked by $\{M\neg y_i\} \models \neg y_i$), and, finally, $M(\neg P_1 \wedge \neg P_2 \wedge \cdots \wedge \neg P_m)$ (which is counterattacked by the premise that $\varphi(\alpha_Y, Z_n) \equiv \top$ so that $\Gamma \models P_1 \vee \cdots \vee P_m$.

In summary, $\text{CA}_{\text{ADM}}^{\text{DL}}(\mathcal{T}, M\varphi)$ if and only if $\varphi(Y_n, Z_n)$ is accepted as an instance of QSAT_2^Σ, hence $\neg\text{VER}_{\text{IDL}}^{\text{DL}}(\mathcal{T}, \{M\psi\})$ if and only if $\exists\alpha_Y\ \varphi(\alpha_Y, Z_n) \equiv \top$. We deduce that $\text{VER}_{\text{IDL}}^{\text{DL}}(\mathcal{T}, \{M\psi\})$ is Π_2^p–complete as claimed. □

Corollary 18

a. $\text{CA}_{\text{IDL}}^{\text{DL}}$ *is* Π_2^p*–hard.*

b. $\text{VER}_{\text{IE}}^{\text{DL},\emptyset}$ *is* Σ_2^p*–hard.*

c. $\textsc{ver}^{\mathrm{DL}}_{\mathrm{IE}}$ *is* $\textsc{d}^p_2$*–hard, where* $\textsc{d}^p_2$ *is the set of languages expressible as the intersection of a language in* Σ^p_2 *with a language in* Π^p_2*.*

Proof: Similar to the arguments used in Corollary 1, Corollary 2 and Thm. 3. For (a), $\textsc{ca}^{\mathrm{DL}}_{\mathrm{IDL}}(\mathcal{T}, M\psi)$ using the $\textsc{abf}$ of Thm. 19 if and only if $\varphi(Y_n, Z_n)$ defines a negative instance of $\textsc{qsat}^{\Sigma}_2$. To establish (b), $\mathcal{T}$ either has an empty ideal extension or its ideal extension is $\{M\psi\}$, so that $\neg\textsc{ver}^{\mathrm{DL},\emptyset}_{\mathrm{IE}}(\mathcal{T})$ if and only if $\textsc{ver}^{\mathrm{DL}}_{\mathrm{IDL}}(\mathcal{T}, \{M\psi\})$. Finally (c) follows by considering instances $\langle\varphi_1, \varphi_2\rangle$ of the canonical $\textsc{d}^p_2$–complete problem, $\textsc{qsat}^{\Sigma}_2$–$\textsc{qsat}^{\Pi}_2$, applying the translation to an $\textsc{abf}$ described above and choosing the ideal extension to be verified as $\{M\psi_2\}$. □

Corollary 19 $\textsc{fie}^{\mathrm{DL}}$ *is* $\textsc{fp}^{\Sigma^p_2}_{||}$*–complete.*

Proof: The upper bound has been proven in Corollary 15(b). The lower bound uses the construction given in Thm. 19 and a similar argument to that of Thm. 7. Instead of *Sat Collection* we use the $\textsc{fp}^{\Sigma^p_2}_{||}$–complete problem of computing the sequence of values describing if $\varphi_i(Y^i_n, Z^i_n)$ is accepted as instance of $\textsc{qsat}^{\Pi}_2$ for n separate instances $\langle\varphi_1, \varphi_2, \ldots, \varphi_n\rangle$. □

The lower bounds results of Corollary 18 (Π^p_2–hardness for credulous reasoning, Σ^p_2–hardness for verifying if the empty set defines the ideal extension, and $\textsc{d}^p_2$–hardness on verifying if an arbitrary set of assumptions is the ideal extension), as with the initial bounds proved in $\textsc{af}$s for the related problems (Corollary 1, Corollary 2 and Thm. 3) are some distance from the upper bound $\textsc{p}^{\Sigma^p_2}_{||}$ on these which is an immediate consequence of Corollary 15(b). In order to reduce this gap, in the $\textsc{af}$ setting, we made use of a number of structural complexity results from [7,8] (Facts 9, 12) together with properties of the unique satisfiability problem ($\textsc{usat}$) established in [34]. Given that our approach with $\textsc{dl}$ instantiatons of $\textsc{abf}$s mirrors many of the ideas applied in the analysis of $\textsc{af}$s, a natural tactic in closing this gap would be to exploit similar methods. To this end it is helpful to observe that the structural characterisation of $\textsc{p}^{\mathrm{NP}}_{||}$–hard languages from [7] leading to Fact 9, is through a simulation of oracle computations and does *not* explicitly depend on the oracle itself being from $\textsc{np}$. We thus obtain [20]

Fact 20 *For all* $k \geq 1$*, a language* L *is* $\textsc{p}^{\Sigma^p_k}_{||}$*–complete if all of the following hold.*

F1. $L \in \textsc{p}^{\Sigma^p_k}_{||}$.
F2. L *is* Σ^p_k*–hard* and L *is* Π^p_k*–hard.*

[20] We have chosen to state this generalisation in terms of classes Σ^p_k and Π^p_k within the polynomial hierarchy – which are the cases of interest for our later development – rather than arbitrary complexity classes $\mathcal{C}$. We further note that the "only if" part is not needed in the subsequent treatment.

F3. L *has property* AND_2.
F4. L *has property* OR_ω.

In order to amplify the lower bounds to Corollary 18 to $\mathrm{P}_{||}^{\Sigma_2^p}$–hardness, Fact 20 suggests an approach, analogous to the devices discussed in the commentary following the proof of Thm. 10, namely

S1. Prove that $\mathrm{CA}_{\mathrm{IDL}}^{\mathrm{DL}}$ is Σ_2^p–hard (noting that we already know it to be Π_2^p–hard).

S2. Prove that $\mathrm{CA}_{\mathrm{IDL}}^{\mathrm{DL}}$ has property AND_2.

S3. Prove that $\mathrm{CA}_{\mathrm{IDL}}^{\mathrm{DL}}$ has property OR_ω.

We deal with S2 and S3 in Thms. 21 and 22.

Theorem 21 $\mathrm{CA}_{\mathrm{IDL}}^{\mathrm{DL}}$ *has property* AND_2.

Proof: We restrict attention to instances $\langle\langle T, A, ^- \rangle, M\alpha\rangle$ for $M\alpha \in A$. In this case it suffices to show that given instances $\langle \mathcal{T}_1, M\alpha_1\rangle$ and $\langle \mathcal{T}_2, M\alpha_2\rangle$ of $\mathrm{CA}_{\mathrm{IDL}}^{\mathrm{DL}}$ we can form an instance $\langle \mathcal{U}, M\beta\rangle$ for which

$$\mathrm{CA}_{\mathrm{IDL}}^{\mathrm{DL}}(\mathcal{U}, M\beta) \text{ if and only if } \mathrm{CA}_{\mathrm{IDL}}^{\mathrm{DL}}(\mathcal{T}_1, M\alpha_1) \ \wedge\ \mathrm{CA}_{\mathrm{IDL}}^{\mathrm{DL}}(\mathcal{T}_2, M\alpha_2)$$

Let $\langle\langle T_1, A_1, ^- \rangle, M\alpha_1\rangle$ and $\langle\langle T_2, A_2, ^- \rangle, M\alpha_2\rangle$ be instances of $\mathrm{CA}_{\mathrm{IDL}}^{\mathrm{DL}}$ where, without loss of generality, the underlying languages L_1 and L_2 are disjoint.[21] Let $\langle W_1, D_1\rangle$ and $\langle W_2, D_2\rangle$ be the default theories from which $\mathcal{T}_1$ and $\mathcal{T}_2$ are defined. Define $\mathcal{U}$ to be the ABF built from the default theory $\langle W_1 \cup W_2, D_1 \cup D_2 \cup D\rangle$ where D is a new set of defaults over new literals $\langle y_1, y_2, z\rangle$ and

$$D \ = \ \left\{ \frac{\top \ :\ M\alpha_1}{\neg y_1},\ \frac{\top \ :\ M\alpha_2}{\neg y_2},\ \frac{\top \ :\ Mz}{z},\ \frac{\top \ :\ M(y_1 \vee y_2)}{\neg z} \right\}$$

$\langle U, B, ^- \rangle$ with $U = T_1 \cup T_2$, $B = A_1 \cup A_2 \cup \{Mz,\ M(y_1 \vee y_2\}$. We claim that

$$\mathrm{CA}_{\mathrm{IDL}}^{\mathrm{DL}}(\mathcal{U}, Mz) \ \Leftrightarrow\ \mathrm{CA}_{\mathrm{IDL}}^{\mathrm{DL}}(\mathcal{T}_1, M\alpha_1) \ \wedge\ \mathrm{CA}_{\mathrm{IDL}}^{\mathrm{DL}}(\mathcal{T}_2, M\alpha_2)$$

Suppose that $\langle \mathcal{U}, Mz\rangle$ defines a positive instance of $\mathrm{CA}_{\mathrm{IDL}}^{\mathrm{DL}}$ and let $\Theta_{\mathcal{U}}$ be the ideal extension of $\mathcal{U}$. From $att(\{M(y_1 \vee y_2)\}, \{Mz\})$ we obtain $att(\Theta_{\mathcal{U}}, \{M(y_1 \vee y_2)\})$, i.e. $\Theta_{\mathcal{U}} \models \neg(y_1 \vee y_2) \ \equiv \neg y_1 \wedge \neg y_2$. It follows, therefore, that $\{M\alpha_1, M\alpha_2\} \subset \Theta_{\mathcal{U}}$. The ideal extension of $\mathcal{U}$ is, however, $\Theta_1 \cup \Theta_2 \cup \{Mz\}$ (where Θ_i is the ideal extension of $\mathcal{T}_i$, so that from $A_1 \cap A_2 = \emptyset$ it follws $M\alpha_1 \in \Theta_1$ and $M\alpha_2 \in \Theta_2$ as claimed.

[21] This can always be guaranteed by renaming the literal terms in each.

For the converse direction, assume that $M\alpha_1 \in \Theta_1$ and $M\alpha_2 \in \Theta_2$ (so that $\text{CA}^{\text{DL}}_{\text{IDL}}(\mathcal{T}_1, M\alpha_1) \wedge \text{CA}^{\text{DL}}_{\text{IDL}}(\mathcal{T}_2, M\alpha_2)$ and consider the $\Theta_1 \cup \Theta_2 \cup \{Mz\} \subset B$. This set is admissible: $\neg att(\Theta_1 \cup \Theta_2 \cup \{Mz\}, \Theta_1 \cup \Theta_2 \cup \{Mz\})$; if $\Gamma \subseteq B$ is such that $att(\Gamma, \Theta_1 \cup \Theta_2 \cup \{Mz\})$ then one of $att(\Theta_1, \Gamma)$, $att(\Theta_2, \Gamma)$ or $att(\Gamma, \{Mz\})$ must hold. In the last of these $\Gamma \models \neg z$ so that $M(y_1 \vee y_2) \in \Gamma$. Now from $\{M\alpha_1, M\alpha_2\} \subseteq \Theta_1 \cup \Theta_2$ we obtain the counterattack $\{M\alpha_1, M\alpha_2\} \models \neg y_1 \wedge \neg y_2$. It is easily seen that $\neg\text{CA}^{\text{DL}}_{\text{ADM}}(\mathcal{U}, \{M(y_1 \vee y_2)\})$: $M\alpha_1 \in \Theta_1$ and $M\alpha_2 \in \Theta_2$ so that for every preferred extension Δ_1 of $\mathcal{T}_1$ and every preferred extension Δ_2 of $\mathcal{T}_2$, $\{M\alpha_1, M\alpha_2\} \subseteq \delta_1 \cup \Delta_2 \models \neg(y_1 \vee y_2)$. We deduce that $\Theta_1 \cup \Theta_2 \cup \{Mz\}$ is an admissible set none of whose attackers is admissible, so that via Lemma 3, this is an ideal set hence, $\text{CA}^{\text{DL}}_{\text{IDL}}(\mathcal{U}, Mz)$. □

Theorem 22 $\text{CA}^{\text{DL}}_{\text{IDL}}$ *has property* OR_ω.

Proof: We use a similar construction to that of Thm 21: given

$$\{\langle \mathcal{T}_1, M\alpha_1 \rangle, \ldots, \langle \mathcal{T}_k, M\alpha_k \rangle\}$$

a set of k instances of $\text{CA}^{\text{DL}}_{\text{IDL}}$ with $\mathcal{T}_i = \langle T_i, A_i, ^{-} \rangle$ defined from default theories $\langle W_i, D_i \rangle$, we form the default theory (W, D) with

$$W \;=\; \bigcup_{i=1}^{k} W_i \;\; ; \;\; D \;=\; D' \cup \bigcup_{i=1}^{k} D_i$$

where D' is a new set of defaults built using additional literals over $\{y, z\}$ given by

$$D' \;\; \left\{ \frac{\top \; : \; My}{\neg z}, \; \frac{\top \; : \; M\alpha_1}{\neg y}, \; \ldots \; , \frac{\top \; : \; M\alpha_k}{\neg y} \right\}$$

We denote by $\mathcal{U}$ the ABF capturing this default theory and fix the instance of $\text{CA}^{\text{DL}}_{\text{IDL}}$ to be $\langle \mathcal{U}, Mz \rangle$. With this instance it is not hard to show that

$$\bigvee_{i=1}^{k} \text{CA}^{\text{DL}}_{\text{IDL}}(\mathcal{T}_i, M\alpha_i) \;\; \Leftrightarrow \;\; \text{CA}^{\text{DL}}_{\text{IDL}}(\mathcal{U}, Mz)$$

□

From the constructions of Thms. 21 and 22, in order to improve the lower bounds obtained to $\text{P}^{\Sigma_2^p}_{||}$–hard, we need to show that $\text{CA}^{\text{DL}}_{\text{IDL}}$ is Σ_2^p–hard. Faced with the related issue (proving NP–hardness of CA_{IDL}) in AF settings, we used a reduction from USAT combined with properties of a randomized reduction from SAT to USAT, i.e. the approach described in Fact 12, Thm. 13. There are two methods we might attempt to use in adapting these techniques to proving $\textsc{ca}^{\mathrm{DL}}_{\mathrm{IDL}}$ to be Σ_2^p–hard. Define the quantifier, $\exists!$ to hold whenever a witnessing solution is *unique*, e.g. USAT is then the language whose positive instances are those for which $\exists!\ \alpha\ :\ \varphi(\alpha)$. Consider the following two formulations of "unique satisfiability" for the second level of PH,

$$\begin{array}{ll} \textsc{usat}_2^{\exists,1} & \exists\ \alpha_Y\ \exists!\ \beta_Z\ \varphi(\alpha_Y,\beta_Z)\equiv\top \\ \textsc{usat}_2^{\exists,2} & \exists!\ \alpha_Y\ \forall\ \beta_Z\ \varphi(\alpha_Y,\beta_Z)\equiv\top \end{array}$$

Given the result of Marx [31, Thm. 5], showing the first of these to be Σ_2^p–complete, one might attempt to prove $\textsc{usat}_2^{\exists,1}\ \leq_m^p\ \textsc{ca}^{\mathrm{DL}}_{\mathrm{IDL}}$. Unfortunately, attempts to translate the form of $\textsc{usat}_2^{\exists,1}$ into a simulating default theory turn out to be problematic: [22] the device used earlier to map multiple satisfying assignments (or unsatisfiability) to non-membership of an argument in the ideal extension, fails with the form $\exists\alpha_Y\exists!\beta_Z\varphi(\alpha_Y,\beta_Z)$ since it is unable to deal with two (or more) instantiations of Y_n *all of which* reduce φ to a CNF having a unique satisfying instantiation.

We first show that if we consider the form described in the second variant – which we will now denote by $\textsc{usat}_2^{\exists}$, i.e. the language of formulae $\varphi(Y_n,Z_n)$ for which $\exists!\alpha_Y\forall\beta_Z\ \varphi(\alpha_Y,\beta_Z)$ – then $\textsc{usat}_2^{\exists}\ \leq_m^p\ \textsc{ca}^{\mathrm{DL}}_{\mathrm{IDL}}$. We then address the question of Σ_2^p–hardness for $\textsc{usat}_2^{\exists}$.

It should be noted that we do *not* assume instances of $\textsc{usat}_2^{\exists}$ to be in a normal form.

Theorem 23 $\textsc{usat}_2^{\exists}\leq_m^p\ \textsc{ca}^{\mathrm{DL}}_{\mathrm{IDL}}$

Proof: Let $\varphi(Y_n,Z_n)$ be an instance of $\textsc{usat}_2^{\exists}$. Consider the following instance – $\langle\mathcal{T}_\varphi,Mw\rangle$ – of $\textsc{ca}^{\mathrm{DL}}_{\mathrm{IDL}}$, in which $\mathcal{T}_\varphi=\langle T_\varphi,A_\varphi,^{-}\rangle$. First form the default theory $\langle W_\varphi,D_\varphi\rangle$ over the literals

$$\{y_i,\ \neg y_i,z_i,\ \neg z_i,\ x_i,\ \neg x_i\ :\ 1\leq i\leq n\}\cup\{w,\neg w,\varphi,\neg\varphi\}$$

with $W_\varphi=\emptyset$ and $D_\varphi\ =\ D_1\cup D_2\cup D_3$ where,

$$D_1\ =\ \left\{\frac{\top:My_i}{y_i},\ \frac{\top:M\neg y_i}{\neg y_i},\ \frac{\top:My_i}{\neg x_i},\ \frac{\top:M\neg y_i}{\neg x_i},\ \frac{\top:Mx_i}{\neg w},\ \frac{\top:Mx_i}{\neg x_i}\ :\ 1\leq i\leq n\right\}$$

$$D_2\ =\ \left\{\frac{\top:M(\neg\varphi(Y_n,Z_n))}{y_i},\ \frac{\top:M(\neg\varphi(Y_n,Z_n))}{\neg y_i}\ :\ 1\leq i\leq n\right\}$$

[22] This is, perhaps, unsurprising: a close inspection of Marx' proof indicates that were it possible directly to prove $\textsc{usat}_2^{\exists,1}\leq_m^p\ \textsc{ca}^{\mathrm{DL}}_{\mathrm{IDL}}$ then it is likely that one could directly derive $\textsc{qsat}_2^{\Sigma}\leq_m^p\ \textsc{ca}^{\mathrm{DL}}_{\mathrm{IDL}}$ thereby obviating any need to consider a generalisation of USAT for the second level of the Polynomial Hierarchy.

$$D_3 \;=\; \left\{ \frac{\top : Mw}{w},\; \frac{\top : M\varphi}{\varphi},\; \frac{\top : M(\neg\varphi(Y_n, Z_n))}{\varphi(Y_n, Z_n)}, \frac{\top : M(\neg\varphi(Y_n, Z_n))}{\neg\varphi} \right\}$$

Note that we distinguish $\varphi(Y_n, Z_n)$ the *propositional formula* presented as an instance of $\mathrm{USAT}_2^\exists$ from φ a formal *literal* in the constructed default theory (W_φ, D_φ). The ABF $\mathcal{T}_\varphi$ then has $T_\varphi = \emptyset$ and

$$A_\varphi \;=\; \{My_i,\; M\neg y_i,\; Mx_i \;:\; 1 \leq i \leq n\} \cup \{M\varphi,\; Mw,\; M(\neg\varphi(Y_n, Z_n))\}$$

with $\overline{M\alpha} \;=\; \neg\alpha$ for each $M\alpha \in A_\varphi$.

The instance of $\mathrm{CA}_{\mathrm{IDL}}^{\mathrm{DL}}$ formed is $\langle \mathcal{T}_\varphi, Mw\rangle$.

We claim that $\varphi(Y_n, Z_n)$ is accepted as an instance of $\mathrm{USAT}_2^\exists$, i.e. $\exists! \alpha_Y\, \forall \beta_Z\, \varphi(\alpha_Y, \beta_Z)$, if and only if $\langle \mathcal{T}_\varphi, Mw\rangle$ is accepted as an instance of $\mathrm{CA}_{\mathrm{IDL}}^{\mathrm{DL}}$.

Let $\underline{\alpha} \;=\; \langle a_1, \ldots, a_n\rangle \in \langle \bot, \top\rangle^n$ be the unique instantiation of Y_n that witnesses $\varphi(Y_n, Z_n)$ as a positive instance of $\mathrm{USAT}_2^\exists$. Consider the subset $\Gamma_{\underline{\alpha}}$ of A_φ given by

$$\Gamma_{\underline{\alpha}} \;=\; \{My_i \;:\; a_i = \top\} \,\cup\, \{M\neg y_i \;:\; a_i = \bot\} \cup \{M\varphi, Mw\}$$

We first note that $\Gamma_{\underline{\alpha}}$ defines an admissible subset of A_φ. It is it clear that $\neg att(\Gamma_{\underline{\alpha}}, \Gamma_{\underline{\alpha}})$. Consider any $\Delta \subseteq A_\varphi$ for which $att(\Delta, \Gamma_{\underline{\alpha}})$. If $My_i \in \Delta$ (resp. $M\neg y_i \in \Delta$) for some $M\neg y_i$ (resp. My_i) in $\Gamma_{\underline{\alpha}}$ then such attacks are countered since $\{M\neg y_i\} \models \neg y_i$ and $\{My_i\} \models y_i$. If $Mx_i \in \Delta$ (so that $\Delta \models \neg w$) then we either have $My_i \in \Gamma_{\underline{\alpha}}$ or $M\neg y_i \in \Gamma_{\underline{\alpha}}$ so that $\Gamma_{\underline{\alpha}} \models \neg x_i$. Finally, if $M(\neg\varphi(Y_n, Z_n)) \in \Delta$ so that $\Delta \models \neg\varphi$ since $\underline{\alpha}$ is such that $\varphi(\underline{\alpha}, Z_n) \equiv \top$ we obtain

$$\Gamma_{\underline{\alpha}} \;\models\; \varphi(Y_n, Z_n) \;\;=\;\; \overline{M(\neg\varphi(Y_n, Z_n))}$$

To complete the first part of the proof it remains only to show that no subset, Δ of A_φ for which $att(\Delta, \Gamma_{\underline{\alpha}})$ is admissible, whence via Lemma 3(a), it follows that $\langle \mathcal{T}_\varphi, Mw\rangle$ is accepted as an instance of $\mathrm{CA}_{\mathrm{IDL}}^{\mathrm{DL}}$. Thus, consider any Δ for which $att(\Delta, \Gamma_{\underline{\alpha}})$. From the earlier discussion showing that $\Gamma_{\underline{\alpha}}$ is admissible we have the following possibilities.

D1. $\Delta \cap \{M(\neg\varphi(Y_n, Z_n)), Mx_1, \ldots, Mx_n\} \neq \emptyset$
In this case $att(\Delta, \Delta)$: $\Delta \models \overline{M(\neg\varphi(Y_n, Z_n))}$ (from D_3) or $\Delta \models \overline{Mx_i}$ (from D_1) so that Δ cannnot be admissible.

D2. $M\neg y_i \in \Delta$ for some $My_i \in \Gamma_{\underline{\alpha}}$.
In order for Δ to be admissible, we must have $\Delta \models \overline{M(\neg\varphi(Y_n, Z_n))}$. This, however, implies that there is an instantiation $\underline{\beta} = \beta_1, \ldots, \beta_n$ of Y_n for which $\varphi(\underline{\beta}, Z_n) \equiv \top$ so contradicting the premise that $\underline{\alpha}$ is the *unique* such instantiation of Y_n.

We deduce that if $\varphi(Y_n, Z_n)$ is a positive instance of $\text{USAT}_2^{\exists}$ then $\langle \mathcal{T}_\varphi, Mw\rangle$ is a positive instance of $\text{CA}_{\text{IDL}}^{\text{DL}}$.

On the other hand suppose that, Θ_φ the ideal extension of $\mathcal{T}_\varphi$ contains the assumption Mw. Then Θ_φ contains exactly one assumption from $\{My_i, M\neg y_i\}$ for each $1 \leq i \leq n$ (at least one in order that $\Theta_\varphi \models \neg x_i$ is needed since $\{Mx_i\} \models \neg w$; at most one since $\{My_i, M\neg y_i\}$ cannot be belong to Θ_φ). Without loss of generality suppose that $\{My_1, \ldots, My_n\} \subset \Theta_\varphi$. Since $\{M(\neg\varphi(Y_n, Z_n))\} \models \neg y_i$ we deduce that $\Theta_\varphi \models \overline{M(\neg\varphi(Y_n, Z_n))}$, i.e. that the instantiation α_Y in which $y_i := \top$ for each i is such that $\varphi(\alpha_Y, Z_n) \equiv \top$. Furthermore, from the fact that no subset Γ of A_φ having $att(\Gamma, \Theta_\varphi)$ can be admissible, we deduce that α_Y is the unique instantiation of Y_n with $\varphi(\alpha_Y, Z_n) \equiv \top$. Hence if $\langle \mathcal{T}_\varphi, Mw\rangle$ is accepted as instance of $\text{CA}_{\text{IDL}}^{\text{DL}}$ then $\varphi(Y_n, Z_n)$ is accepted as an instance of $\text{USAT}_2^{\exists}$. ▫

Improvements to the lower bounds of Corollary 18 are established as a consequence of the following result.

Theorem 24 $\text{QSAT}_2^{\Sigma} \leq_m^{rp} \text{USAT}_2^{\exists}$ *with probability* $1/4n$.

Proof: The detailed argument is presented in Appendix C below. ▫

Corollary 20

a. $\text{CA}_{\text{IDL}}^{\text{DL}}$ *is* $\text{P}_{||}^{\Sigma_2^p}$*–complete via* $\leq_m^{rp}$ *with probability* $1 - 2^{-n}$.

b. $\text{VER}_{\text{IE}}^{\text{DL},\emptyset}$ *is* $\text{P}_{||}^{\Sigma_2^p}$*–complete via* $\leq_m^{rp}$ *with probability* $1 - 2^{-n}$.

c. $\text{VER}_{\text{IE}}^{\text{DL}}$ *is* $\text{P}_{||}^{\Sigma_2^p}$*–complete via* $\leq_m^{rp}$ *with probability* $1 - 2^{-n}$.

Proof: Noting that membership in $\text{P}_{||}^{\Sigma_2^p}$ for each case has already been shown and recalling the earlier arguments of Corollary 18, it suffices to prove only (a).

Consider the $\text{P}_{||}^{\Sigma_2^p}$–complete problem [23] $\text{OR}_\omega(\text{QSAT}_2^{\Sigma}\text{–QSAT}_2^{\Pi})$ instances of which comprise n pairs of CNF formulae – $\langle \varphi_i(Y_m, Z_m), \psi_i(V_m, W_m)\rangle$ – over disjoint sets of variables (so the instance involves $4nm$ distinct variables in total). Such an instance being accepted if there is at least one pair $\langle \varphi_i(Y_m, Z_m), \psi_i(V_m, W_m)\rangle$ for which $\varphi_i(Y_m, Z_m)$ is accepted as an instance of QSAT_2^{Σ}, i.e. $\exists\alpha_Y\ \neg\varphi_i(\alpha_Y, Z_m) \equiv \top$ *and* $\psi_i(Y_m, Z_m)$ is accepted as an instance of QSAT_2^{Π}, i.e. $\forall\alpha_Y \exists\beta_Z\ \psi(\alpha_Y, \beta_Z)$.

[23] That the given problem is indeed $\text{P}_{||}^{\Sigma_2^p}$–complete is an easy generalisation of the methods used to establish $\text{OR}_\omega(\text{SAT–UNSAT})$ is $\text{P}_{||}^{\text{NP}}$–complete.

Choosing uniformly at random one of the n pairs in such an instance immediately yields

$$\text{OR}_\omega(\text{QSAT}_2^\Sigma\text{–QSAT}_2^\Pi) \quad \leq_m^{rp} \quad \text{QSAT}_2^\Sigma\text{–QSAT}_2^\Pi \quad \text{with probability } 1/n$$

From Thm. 24,

$$\text{QSAT}_2^\Sigma\text{–QSAT}_2^\Pi \quad \leq_m^{rp} \quad \text{USAT}_2^\exists\text{–QSAT}_2^\Pi \quad \text{with probability } 1/4n$$

so that combining Corollary 18(a) and Thm. 23 gives

$$\text{OR}_\omega(\text{QSAT}_2^\Sigma\text{–QSAT}_2^\Pi) \quad \leq_m^{rp} \quad \text{AND}_2(\text{CA}_{\text{IDL}}^{\text{DL}}) \quad \text{with probability } 1/4n^2$$

Hence, via Thm. 21,

$$\text{OR}_\omega(\text{QSAT}_2^\Sigma\text{–QSAT}_2^\Pi) \quad \leq_m^{rp} \quad \text{CA}_{\text{IDL}}^{\text{DL}} \quad \text{with probability } 1/4n^2$$

Finally, applying Fact 12 and Thm. 22 we obtain

$$\text{OR}_\omega(\text{QSAT}_2^\Sigma\text{–QSAT}_2^\Pi) \quad \leq_m^{rp} \quad \text{CA}_{\text{IDL}}^{\text{DL}} \quad \text{with probability } 1 - 2^{-n}$$

from which (a) is immediate. □

4 Conclusions and Further Work

We have considered the computational complexity of decision and search problems arising in the ideal semantics of [19,20], addressing both the AF model of Dung [17] and *flat* frameworks within the ABF approach of Bondarenko *et al.* [5].

It has been shown that for settings in which credulous reasoning can be carried out in a complexity class $\mathcal{C}$, the principal computational problems of interest can be resolved within $\text{P}_{||}^{\mathcal{C}}$ or its functional analogue $\text{FP}_{||}^{\mathcal{C}}$: classes believed to lie strictly below $\text{coNP}^{\mathcal{C}}$ the complexity of *sceptical* reasoning in such environments. We have, in addition, presented compelling evidence that deciding if an argument is acceptable under the ideal semantics, if a set of arguments defines the ideal extension, and if the ideal extension is empty, are not contained within any complexity class falling strictly within $\text{P}_{||}^{\mathcal{C}}$: all of these problems being $\text{P}_{||}^{\mathcal{C}}$–hard with respect to $\leq_m^{rp}$ reductions of probability $1-2^{-n}$. Although this complexity class compares unfavourably with the $\mathcal{C}$ and co$\mathcal{C}$–complete status of related questions under the credulous preferred semantics, it represents an improvement on the $\text{coNP}^{\mathcal{C}}$–completeness level of similar issues within the sceptical preferred semantics.

Given that sceptical acceptance is a precondition of membership in an ideal set this reduction in complexity may appear surprising. If we consider the AF cases, the apparent discrepancy is, however, accounted for by examining the second condition that a *set* of arguments must satisfy in order to form an ideal set: as well as being sceptically accepted, the set must be admissible. This condition plays a significant role in the complexity shift. An important reason why testing sceptical acceptance of a gven argument x fails to belong to coNP (assuming coNP $\neq \Pi_2^p$) is that the condition "no attacker of x is credulously accepted" while *necessary* for sceptical acceptance of x is not *sufficient*: a fact which seems first to have been observed by Vreeswijk and Prakken [35] in their analysis of sound and complete proof procedures for credulous acceptance. Although this condition *is* sufficient in *coherent* frameworks, deciding if $\mathcal{H}$ is coherent is already Π_2^p–complete [23]. In contrast, as demonstrated in the characterisation of ideal sets given in Lemma 1, an *admissible* set, S, is also sceptically accepted if and only if no argument in S^- – i.e. attacker of S – is credulously accepted: we thus have a condition which can be tested in coNP. With an analogous characterisation of ideal sets also holding in flat assumption based fraemworks – Lemma 3 – a similar reduction in complexity is obtained.

The reason why *finding* the ideal extension (and consequently decision questions predicated on its properties, e.g. cardinality, membership, etc.) can be performed more efficiently than testing sceptical acceptance stems from the fact this set can be readily computed given the *bipartite* framework, $\mathcal{B}(\mathcal{X}_{\mathrm{PSA}}, \mathcal{X}_{\mathrm{OUT}}, \mathcal{F})$ associated with $\mathcal{H}(\mathcal{X}, \mathcal{A})$. Construction of this framework only requires determining the set, $\mathcal{X}_{\mathrm{OUT}}$, of arguments which are not credulously accepted, so that *explicit* consideration of sceptical acceptance is never required. Although a direct representation of the structures as a bipartite graph is not employed, the analogous partition of assumptions in flat ABFs affords a simlar device by which explicit testing of sceptical acceptance is avoided.

This paper has focused on the graph-theoretic abstract argumentation framework model from [17] and instantiations of *flat* assumption based frameworks, e.g. those realising LP and DL theories. Of the questions left open in ideal semantics, possibly, the most challenging concerns the computational complexity of problems in non-flat ABFs, in particular those realising AEL. Of these it would be of some interest to demonstrate, as we conjecture is indeed the case, that verifying a set of assumptions as an ideal set is Π_4^p–hard in AEL settings, thereby giving an exact bound. A final collection of issues concern the performance of Algorithm 1 as a practical mechanism for constructing the ideal extension. Thus Dung, Mancarella, and Toni [20] describe dialectic approaches for identifying the ideal extension using a variation of a procedure described in [18]. To what extent the methods of Algorithm 1 can be used to complement or offer an effective alternative is a question of some interest.

References

[1] L. M. Adleman and K. Manders. Reducibility, randomness and intractibility. In *Proc. 9th ACM Symposium on Theory of Computing*, pages 151–163, 1979.

[2] S. Arnborg, J. Lagergren, and D. Seese. Easy problems for tree-decomposable graphs. *Jnl. of Algorithms*, 12:308–340, 1991.

[3] T. J. M. Bench-Capon and P. E. Dunne. Argumentation in artificial intelligence. *Artificial Intelligence*, 171:619–641, 2007.

[4] A. Blass and Y. Gurevich. On the unique satisfiability problem. *Information and Control*, 55:80–82, 1982.

[5] A. Bondarenko, P. M. Dung, R. A. Kowalski, and F. Toni. An abstract, argumentation-theoretic approach to default reasoning. *Artificial Intelligence*, 93:63–101, 1997.

[6] S. Brass, J. Dix, B. Freitag, and U. Zukowski. Transformation-based bottom-up computation of the well-founded model. In *Theory and Proactice of Logic Programming*, 2001.

[7] R. Chang and J. Kadin. On computing Boolean connectives of characteristic functions. *Math. Syst. Theory*, 28:173–198, 1995.

[8] R. Chang, J. Kadin, and P. Rohatgi. On unique satisfiability and the threshold behavior of randomised reductions. *Jnl. of Comp. and Syst. Sci.*, pages 359–373, 1995.

[9] S. Costantini, O. D'Antona, and A. Provetti. On the equivalence and range of applicability of graph-based representations of logic programs. *Inf. Proc. Letters*, 84:241–249, 2002.

[10] B. Courcelle. The monadic second-order logic of graphs. I. recognizable sets of finite graphs. *Information and Computation*, 85(1):12–75, 1990.

[11] B. Courcelle. The monadic second-order logic of graphs III: tree-decompositions, minor and complexity issues. *Informatique Théorique et Applications*, 26:257–286, 1992.

[12] Y. Dimopoulos, B. Nebel, and F. Toni. Preferred arguments are harder to compute than stable extensions. In D. Thomas, editor, *Proc. of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99-Vol1)*, pages 36–43, San Francisco, 1999. Morgan Kaufmann Publishers.

[13] Y. Dimopoulos, B. Nebel, and F. Toni. Finding admissible and preferred arguments can be very hard. In A. G. Cohn, F. Giunchiglia, and B. Selman, editors, *KR2000: Principles of Knowledge Representation and Reasoning*, pages 53–61, San Francisco, 2000. Morgan Kaufmann.

[14] Y. Dimopoulos, B. Nebel, and F. Toni. On the computational complexity of assumption-based argumentation for default reasoning. *Artificial Intelligence*, 141:55–78, 2002.

[15] Y. Dimopoulos and A. Torres. Graph theoretical structures in logic programs and default theories. *Theoretical Computer Science*, 170:209–244, 1996.

[16] R. G. Downey and M. R. Fellows. Fixed parameter tractability and completeness I: basic results. *SIAM Jnl. on Computing*, 24:873–921, 1995.

[17] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and N-person games. *Artificial Intelligence*, 77:321–357, 1995.

[18] P. M. Dung, R. A. Kowalski, and F. Toni. Dialectic proof procedures for assumption-based, admissible argumentation. *Artificial Intelligence*, 170:114–159, 2006.

[19] P. M. Dung, P. Mancarella, and F. Toni. A dialectical procedure for sceptical assumption-based argumentation. In P. E. Dunne and T. J. M. Bench-Capon, editors, *Proc. 1st Int. Conf. on Computational Models of Argument*, volume 144 of *FAIA*, pages 145–156. IOS Press, 2006.

[20] P. M. Dung, P. Mancarella, and F. Toni. Computing ideal sceptical argumentation. *Artificial Intelligence*, 171:642–674, 2007.

[21] P. E. Dunne. Computational properties of argument systems satisfying graph-theoretic constraints. *Artificial Intelligence*, 171:701–729, 2007.

[22] P. E. Dunne. The computational complexity of ideal semantics I: abstract argumentation frameworks. In *Proc. 2nd Int. Conf. on Computational Models of Argument*, volume 172 of *FAIA*, pages 147–158. IOS Press, 2008.

[23] P. E. Dunne and T. J. M. Bench-Capon. Coherence in finite argument systems. *Artificial Intelligence*, 141:187–203, 2002.

[24] P. E. Dunne and T. J. M. Bench-Capon. Two party immediate response disputes: properties and efficiency. *Artificial Intelligence*, 149:221–250, 2003.

[25] K. Eshghi and R. A. Kowalski. Abduction compared with negation as failure. In *Proc. 6th Int. Conf. on Logic Programming*, pages 234–254, 1989.

[26] G. Gottlob. Complexity results for nonmonotonic logcis. *Journal of Logic and Computation*, 2(3):397–425, 1992.

[27] B. Jenner and J. Toran. Computing functions with parallel queries to NP. *Theoretical Computer Science*, 141:175–193, 1995.

[28] D. S. Johnson. A catalog of complexity classes. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science. Volume A: Algorithms and Complexity*, pages 67–161. Elsevier Science, 1998.

[29] K. Konolige. On the relation between default and autoepistemic logic. *Artificial Intelligence*, 35:343–382, 1988.

[30] T. Linke. Graph theoretical characterization and computation of answer sets. In *Proc. IJCAI-2001*, pages 641–648, 2001.

[31] D. Marx. Complexity of unique list colorability. Technical report, Dept. of Comp. Science and Inf. Theory, Budapest Univ. of Tech. and Econ., December 2007. `http://www.cs.bme.hu/~dmarx/papers/marx-unique.pdf`.

[32] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

[33] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[34] L. G. Valiant and V. V. Vazirani. NP is as easy as detecting unique solutions. *Theoretical Computer Science*, 47:85–93, 1986.

[35] G. Vreeswijk and H. Prakken. Credulous and sceptical argument games for preferred semantics. In *Proc. of JELIA'2000, The 7th European Workshop on Logic for Artificial Intelligence.*, pages 224–238, Berlin, 2000. Springer LNAI 1919, Springer Verlag.

[36] K. Wagner. Bounded query computations. In *Proc. 3rd Conf. on Structure in Complexity Theory*, pages 260–277, 1988.

[37] K. Wagner. Bounded query classes. *SIAM Jnl. Comput.*, 19:833–846, 1990.

Appendix A

A.1 The argumentation framework $\mathcal{H}_\Phi$

The form we describe is virtually identical to that first presented by Dimopoulos and Torres [15, Thm. 5.1, p. 227] where it is used to establish NP–hardness of CA via a reduction from 3-SAT.

Given a CNF formula $\Phi(Z_n) = \wedge_{j=1}^{m} C_j$ with each C_j a disjunction of literals from $\{z_1, \ldots, z_n, \neg z_1, \ldots, \neg z_n\}$, the AF, $\mathcal{H}_\Phi(\mathcal{X}, \mathcal{A})$ has

$$\begin{aligned}
\mathcal{X} &= \{\Phi, C_1, \ldots, C_m\} \cup \{z_i,\ \neg z_i\ :\ 1 \le i \le n\} \\
\mathcal{A} &= \{\langle C_j, \Phi\rangle\ :\ 1 \le j \le m\} \cup \{\langle z_i, \neg z_i\rangle,\ \langle \neg z_i, z_i\rangle\ :\ 1 \le i \le n\} \cup \\
&\quad \{\langle z_i, C_j\rangle\ :\ z_i \text{ occurs in } C_j\} \cup \{\langle \neg z_i, C_j\rangle\ :\ \neg z_i \text{ occurs in } C_j\}
\end{aligned}$$

Fig. 3 illustrates $\mathcal{H}_\Phi$.

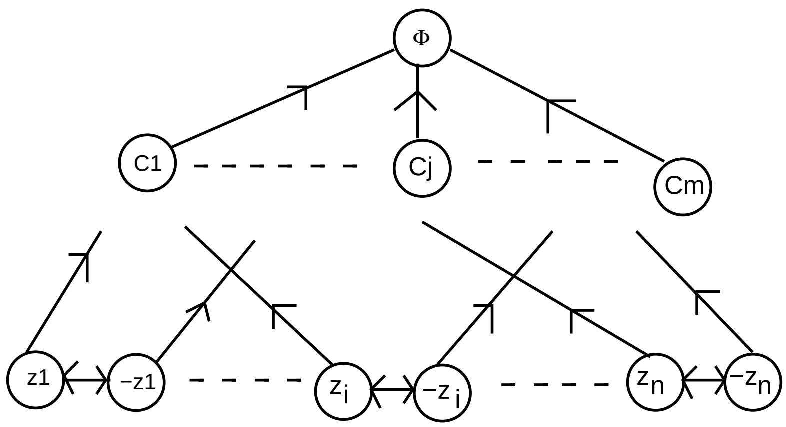

Fig. 3. The Argumentation Framework $\mathcal{H}_\Phi$

Fact 25 *(Dimopoulous and Torres [15]) Let $\Phi(Z_n)$ be an instance of 3-SAT, i.e. a 3-CNF formula. Then $\Phi(Z_n)$ is satisfiable if and only if* CA$(\mathcal{H}_\Phi(\mathcal{X}, \mathcal{A}), \Phi)$.

.

A.2 The argumentation framework $\mathcal{G}_\Phi$

The proof that SA is Π_2^p–complete from [23] uses a reduction from QSAT$_2^\Pi$ instances of which may, without loss of generality, be restricted to 3-CNF formulae, $\Phi(Y_n, Z_n)$, accepted if $\forall\, \alpha_Y \exists\, \beta_Z\ \Phi(\alpha_Y, \beta_Z)$, i.e. for every instantiation of the propositional variables Y_n (α_Y) there is some instantiation of Z_n (β_Z) for which $\langle \alpha_Y, \beta_Z \rangle$ satisfies Φ.

The AF $\mathcal{G}_\Phi(\mathcal{W}, \mathcal{B})$ is formed from $\mathcal{H}_\Phi(\mathcal{X}, \mathcal{A})$, i.e. $\mathcal{X} \subset \mathcal{W}$ and $\mathcal{A} \subset \mathcal{B}$, so that

$$
\begin{array}{rcl}
\mathcal{W} & = & \{\Phi, C_1, \ldots, C_m\} \cup \{y_i, \neg y_i, z_i, \neg z_i \ :\ 1 \le i \le n\} \cup \{b_1, b_2, b_3\} \\
\mathcal{B} & = & \{\langle C_j, \Phi \rangle \ :\ 1 \le j \le m\} \cup \\
 & & \{\langle y_i, \neg y_i \rangle, \langle \neg y_i, y_i \rangle, \langle z_i, \neg z_i \rangle, \langle \neg z_i, z_i \rangle \ :\ 1 \le i \le n\} \cup \\
 & & \{\langle y_i, C_j \rangle \ :\ y_i \text{ occurs in } C_j\} \cup \{\langle \neg y_i, C_j \rangle \ :\ \neg y_i \text{ occurs in } C_j\} \cup \\
 & & \{\langle z_i, C_j \rangle \ :\ z_i \text{ occurs in } C_j\} \cup \{\langle \neg z_i, C_j \rangle \ :\ \neg z_i \text{ occurs in } C_j\} \cup \\
 & & \{\langle \Phi, b_1 \rangle, \langle \Phi, b_2 \rangle, \langle \Phi, b_3 \rangle, \langle b_1, b_2 \rangle, \langle b_2, b_3 \rangle, \langle b_3, b_1 \rangle\} \cup \\
 & & \{\langle b_1, z_i \rangle, \langle b_1, \neg z_i \rangle \ :\ 1 \le i \le n\}
\end{array}
$$

The resulting AF is shown in Fig. 4.

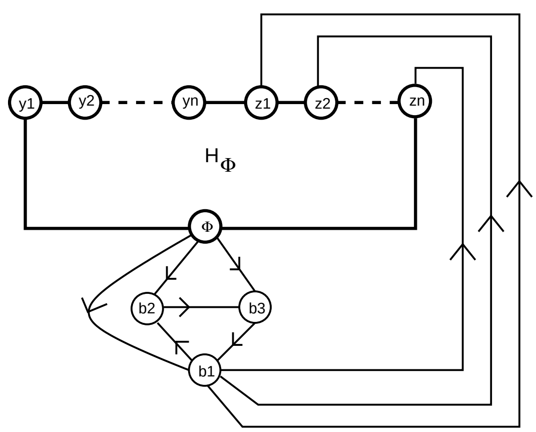

Fig. 4. The Argumentation Framework $\mathcal{G}_\Phi$.

Fact 26 *(Dunne and Bench-Capon [23])*

a. $\Phi(Y_n, Z_n)$ *is accepted as an instance of* QSAT$_2^\Pi$ *if and only if* SA$(\mathcal{G}_\Phi, \Phi)$.
b. $\Phi(Y_n, Z_n)$ *is accepted as an instance of* QSAT$_2^\Pi$ *if and only if* $\mathcal{G}_\Phi$ *is coherent.*

We recall there is a possible objection to defining $\langle \mathcal{H}, x\rangle \in \text{SA}_{ST}$ as accepted even when $\mathcal{G}$ has *no* stable extension whatsoever, i.e. $\mathcal{H} \not\in \text{EXISTS}_{ST}$.[24] In order to deal with this objection one might require as a precondition of $\langle \mathcal{H}, x\rangle \in \text{SA}_{ST}$ that $\mathcal{H} \in \text{EXISTS}_{ST}$. In this appendix we present a proof that this variant, which we will denote by $\text{SA}_{ST}^{\geq 1}$ is D^p–complete.

Theorem 27 *The decision problem* $\text{SA}_{ST}^{\geq 1}$ *accepting instances* $\langle \mathcal{H}, x\rangle$ *for which* $\mathcal{H}$ *has at least one stable extension* and x *is a member of every such extension is* D^p*–complete.*

Proof: Membership in D^p follows by observing that accepted instances are exactly those belonging to,

$$L_1 \cap L_2 \;=\; \text{CA}_{ST} \cap \{\langle\langle \mathcal{X}, \mathcal{A}\rangle, x\rangle \;:\; \forall\, S \subseteq \mathcal{X}\; (\text{VER}_{ST}(\langle \mathcal{X}, \mathcal{A}\rangle, S) \;\Rightarrow\; (x \in S)\}$$

As $L_1 \in \text{NP}$ and $L_2 \in \text{coNP}$ the upper bound is immediate.

For the matching D^p–hardness lower bound we again employ the standard translation from 3-CNF instances described in Appendix A, to instances $\langle \varphi_1, \varphi_2\rangle$ of the canonical D^p–hard problem SAT-UNSAT. Given an instance $\langle \varphi_1, \varphi_2\rangle$ of SAT-UNSAT, form the instance $\langle \mathcal{K}, \psi_2\rangle$ of $\text{SA}_{ST}^{\geq 1}$ in which the AF, $\mathcal{K}$ has arguments $\mathcal{X}_{\varphi_1} \cup \mathcal{X}_{\varphi_2} \cup \{\alpha, \psi_1, \psi_2\}$, i.e. the arguments of the AFs $\mathcal{H}_{\varphi_1}$ and $\mathcal{H}_{\varphi_2}$ together with new arguments $\{\psi_1, \psi_2, \alpha\}$. In addition to those attacks already present in $\mathcal{H}_{\varphi_j}$, $\mathcal{K}$ contains attacks

$\{\langle \varphi_i, \psi_i\rangle, \langle \psi_i, \varphi_i\rangle\}$ $\;1 \leq i \leq 2$

$\langle \psi_1, y_i\rangle$ for each literal y_i of φ_1

$\langle \psi_2, y_i\rangle$ for each literal y_i of φ_2

$\{\langle \varphi_1, \alpha\rangle, \langle \alpha, \psi_2\rangle\}$

This AF is illustrated in Fig. 5.

The instance $\langle \mathcal{K}, \psi_2\rangle$ of $\text{SA}_{ST}^{\geq 1}$ satisfies both EXISTS_{ST} *and* has ψ_2 a member of every stable extension if and only if $\varphi_1 \in$ 3-SAT *and* $\varphi_2 \in$ 3-UNSAT.

[24] Resulting in AFs for which $\langle\langle \mathcal{X}, \mathcal{A}\rangle, x\rangle \in \text{SA}_{ST}$ and $\langle\langle \mathcal{X}, \mathcal{A}\rangle, x\rangle \not\in \text{CA}_{ST}$ for *every* $x \in \mathcal{X}$, i.e. *every* argument is sceptically accepted but *none* credulously so.

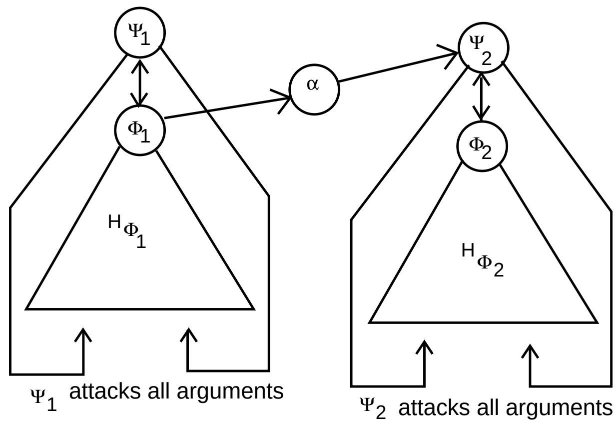

Fig. 5. The Argumentation Framework $\mathcal{K}$

To see this, first note that $\mathcal{K}$ has no stable extension if and only if φ_1 is unsatisfiable: $\text{CA}_{\text{PR}}(\mathcal{K}, \varphi_1)$ if and only if φ_1 is satisfiable, so that if this is not case then α is unattacked in every preferred extension. Thus the instance $\langle \mathcal{K}, \psi_2 \rangle$ can only be accepted if φ_1 is satisfiable. In addition, should φ_1 be satisfiable, then ψ_2 belongs to every stable extension if and only if φ_2 is *unsatisfiable*. That is, if φ_2 has a satisfying instantiation either $\mathcal{K}$ has no stable extension at all (φ_1 being unsatisfiable) or if both φ_1 and φ_2 are satisfiable then $\{\varphi_1, \varphi_2\}$ together with the literals selected by the witnessing satisfying assignments form a stable extension that does not contain ψ_2. □

Appendix C: Proof of Theorem 24

We recall that Thm. 24 asserts

$$\text{QSAT}_2^{\Sigma} \ \leq_m^{rp} \ \text{USAT}_2^{\exists} \text{ with probability } 1/4n.$$

The proof of this is based on the fact that Valiant and Vazirani's randomized reduction from SAT to USAT [34] exploits a *combinatorial* property of *arbitrary* subsets S of the n-dimensional vector space $\langle 0, 1 \rangle^n$ formed via the operations $\{\oplus, \wedge\}$ [25] and randomly chosen elements from this, i.e. it does not explicitly depend on *satisfiability per se*.

[25] That is the vector space $\mathbf{GF}[2]^n$: the operation $\oplus$ being Boolean "exclusive–or", i.e. $x \oplus y = 1$ if and only if $x \neq y$.

In our subsequent discussion, we view instantiations to a set of n propositional variables as n-tuples from $\langle 0,1\rangle^n$. Given

$$\underline{x} = \langle x_1, x_2, \ldots, x_n\rangle \in \langle 0,1\rangle^n$$
$$\underline{y} = \langle y_1, y_2, \ldots, y_n\rangle \in \langle 0,1\rangle^n$$

the *inner product* w.r.t. $\{\oplus, \wedge\}$ of $\underline{x}$ and $\underline{y}$ (denoted $\underline{x} \cdot \underline{y}$) is the value in $\langle 0,1\rangle$ given by,

$$\underline{x} \cdot \underline{y} \;=\; \bigoplus_{i=1}^{n} (x_i \wedge y_i)$$

The reduction from SAT to USAT in [34] builds on the following result.

Fact 28 *([34, Thm. 2.4, p. 89]) Let $S \subseteq \langle 0,1\rangle^n$ and w_1, w_2, …, w_n be chosen uniformly at random from $\langle 0,1\rangle^n$. For each $1 \leq i \leq n$, define S_i to be the set*

$$S_i \;=\; \left\{ v \in S \;:\; \bigwedge_{j=1}^{i} (v \cdot w_j \;=\; 0) \right\}$$

Furthermore, let $P_n(S)$ be the probability that, for some $i \leq n$, $|S_i| = 1$. Then $P_n(S) \geq 1/4$.

The main device needed is a mechanism for manipulating the structure of formulae in order to exploit Fact 28. This is achieved in the following development of [34, Lemma 2.1, p. 88], where the notion of a set of *candidates* for an instance $\varphi(Y_n, Z_n)$ of QSAT$_2^\Sigma$ – denoted $C(\varphi)$ – is defined as

$$C(\varphi) \;=\; \{\; \alpha \in \langle 0,1\rangle^n \;:\; \forall \beta \in \langle 0,1\rangle^n,\; \varphi(\alpha, \beta) = 1\}$$

Note the set of candidates is well-defined irrespective of the exact form taken by φ, i.e. it is not required that φ be either in CNF or DNF: we restrict attention, however, to formulae defined over the logical basis $\{\wedge, \vee, \neg\}$.

Lemma 5 *Let $\varphi(Y_n, Z_n)$ be a formula defining an instance of* QSAT$_2^\Sigma$ *and let $w_1,\ w_2, \ldots,\ w_k$ be elements of $\langle 0,1\rangle^n$.*

a. There is a propositional formula, $\psi_k(Y_n, Z_n)$ (that may be constructed in linear time) and is such that $C(\psi_k) \subseteq C(\varphi)$ and

$$\forall\, \alpha \in C(\psi_k) \quad \bigwedge_{i=1}^{k} (\alpha \cdot w_i = 0)$$

b. Given ψ_k one may construct in polynomial time a formula over the basis $\{\wedge, \vee, \neg\}$, $\chi_k(Y_n \cup U_m, Z_n)$ using variables $Y_n \cup Z_n \cup U_m$ for some value of m such that defining

$$C(\chi_k) \;=\; \{\; \gamma \in \langle 0,1\rangle^{n+m} \;:\; \forall\, \beta\; \chi(\gamma,\beta) = 1\}$$

Then there exists some $\langle \gamma_{n+1}, \ldots, \gamma_{n+m}\rangle \in \langle 0,1\rangle^m$ for which

$$\langle \alpha_1, \alpha_2, \ldots, \alpha_n, \gamma_{n+1}, \gamma_{n+2}, \ldots, \gamma_{n+m}\rangle \in C(\chi_k)$$

if and only if $\langle \alpha_1, \alpha_2, \ldots, \alpha_n\rangle \in C(\psi_k)$

Proof: Given $\varphi(Y_n, Z_n)$ as defined in the Lemma statement and $\langle w_1, \ldots, w_k\rangle$ from $\langle 0,1\rangle^n$, define the formula $\psi_k(Y_n, Z_n)$ to be $\varphi(Y_n, Z_n) \wedge \omega_k(Y_n)$ where,

$$\begin{aligned}
\omega_1(Y_n) &= \left(1 \;\oplus \bigoplus_{j\,:\,w_{1,j}=1} y_j\right) \quad \text{for } i = 1 \\
\omega_i(Y_n) &= \omega_{i-1} \;\wedge\; \left(1 \;\oplus \bigoplus_{j\,:\,w_{i,j}=1} y_j\right) \quad \text{for } i > 1
\end{aligned}$$

Notice that $C(\psi_k) \subseteq C(\varphi)$ since,

$$\begin{aligned}
C(\psi_k) &= \{\alpha \;:\; \psi_k(\alpha, Y_n) \equiv 1\} \\
&= \{\alpha \;:\; \varphi(\alpha, Y_n) \equiv 1\} \cap \{\alpha \;:\; \omega_k(\alpha) = 1\} \\
&\subseteq \{\alpha \;:\; \varphi(\alpha, Y_n) \equiv 1\} \\
&= C(\varphi)
\end{aligned}$$

Furthermore, any such α must satisfy $\omega_k(\alpha) = 1$, so that for each $1 \le i \le k$, $(1 \;\oplus\; \oplus_{j\,:\,w_{i,j}=1}\; y_j)$, it follows that

$$\left(\bigoplus_{j\,:\,w_{i,j}=1} y_j\right)(\alpha) \;=\; 0$$

i.e. $w_i \;\cdot\; y = 0$. This establishes part (a).

For part (b), with $\varphi(Y_n, Z_n)$ assumed already to be over the logical basis $\{\wedge, \vee, \neg\}$ to convert ψ_k to a formula χ of this form it suffices to observe that only the sub-formula ω_k (involving the opeartion $\oplus$) is not in the required form. Furthermore ω_k consists of a conjunction of terms Q_i each having the form

$$1 \;\oplus\; y_{i_1} \;\oplus\; y_{i_2} \;\oplus\; \cdots \;\oplus\; y_{i_m}$$

for some subset $\{y_{i_1}, \ldots, y_{i_m}\}$ of Y_n. Without loss of generality it suffices to show that $1 \oplus y_1 \oplus y_2 \oplus \cdots \oplus y_m$ may be efficiently translated to a formula, $\chi(Y_m \cup U_{m-1})$ using only operations from $\{\wedge, \vee, \neg\}$.

Note that if $m = 0$ no translation is needed; and if $m = 1$ then $1 \oplus y_1 \equiv \neg y_1$, so we may assume $m \geq 2$. Introducing new variables $\{u_1, u_2, \ldots, u_{m-1}\}$, $\chi(Y_n \cup U_{m-1})$ is formed from

$$(u_1 \leftrightarrow (y_1 \oplus y_2)) \wedge (u_2 \leftrightarrow (u_1 \oplus y_3)) \wedge \ldots \wedge (u_{m-1} \leftrightarrow (u_{m-2} \oplus y_m)) \wedge (u_{m-1} \oplus 1)$$

Since

$$\begin{aligned}(x \leftrightarrow (y \oplus z)) &\equiv (1 \oplus x \oplus y \oplus z)\\ &\equiv (\neg x \oplus y \oplus z)\\ &\equiv (\neg x \vee y \vee z)(\neg x \vee \neg y \vee \neg z)(x \vee y \vee \neg z)(x \vee \neg y \vee z)\end{aligned}$$

this conversion can be performed in polynomial time without introducing further new variables. Letting $\chi_k(Y_n \cup U_m, Z_n)$ be the formula resulting by translating $\psi_k(Y_n, Z_n) = \varphi(Y_n, Z_n) \wedge \omega_k(Y_n)$ in this way, it remains only to note that any candidate

$$\langle \alpha_1, \alpha_2, \ldots, \alpha_n, \gamma_{n+1}, \gamma_{n+2}, \ldots, \gamma_{n+m} \rangle$$

of $\chi_k(Y_n \cup U_m, Z_n)$ maps to a unique candidate ($\langle \alpha_1, \alpha_2, \ldots, \alpha_n \rangle$) of $\psi_k(Y_n, Z_n)$ and that for $\langle \alpha_1, \alpha_2, \ldots, \alpha_n \rangle \in C(\psi_k)$ there is some choice of $\langle \gamma_{n+1}, \gamma_{n+2}, \ldots, \gamma_{n+m} \rangle$ for which

$$\langle \alpha_1, \alpha_2, \ldots, \alpha_n, \gamma_{n+1}, \gamma_{n+2}, \ldots, \gamma_{n+m} \rangle \in C(\chi_k)$$

$\square$

We observe that the argument and construction of Lemma 5 is effectively identical to that of [34, Lemma 2.1]: the latter deals with satisfying instantiations of $\varphi(X_n)$ instead of the notion of candidates of $\varphi(Y_n, Z_n)$.

We now have sufficient machinery in place to provide the

Proof: (of Thm. 24) Given $\varphi(Y_n, Z_n)$ an instance of QSAT_2^{Σ} choose (uniformly at random) a value k in $\{1, 2, \ldots, n\}$ and then (also uniformly at random) k n-tuples w_1, w_2, …, w_k from $\langle 0, 1 \rangle^n$. The instance of $\mathrm{USAT}_2^{\exists}$ constructed is the formula $\chi_k(Y_n \cup U_m, Z_n)$ of Lemma 5.

To see that $\varphi(Y_n, Z_n)$ is accepted as an instance of QSAT_2^Σ if and only if $\chi_k(Y_n \cup U_m, Z_n)$ is accepted as an instance of $\text{USAT}_2^\exists$, first observe that if $\varphi(Y_n, Z_n)$ fails to define a positive instance of QSAT_2^Σ, i.e. $C(\varphi) = \emptyset$, then χ_k can never define a positive instance of $\text{USAT}_2^\exists$: in this case $C(\chi_k) = \emptyset$.

On the other hand, if $|C(\varphi)| = t > 0$, then from Fact 28, given n randomly chosen elements, $\langle w_1, w_2, \ldots, w_n\rangle$ from $\langle 0, 1\rangle^n$, we know that with probability at least $1/4$ there is a choice of $k \in \{1, 2, 3, \ldots, n\}$ such that

$$\left|\left\{ v \in C(\varphi) \ : \ \bigwedge_{i=1}^{k} \ (v \ \cdot \ w_i \ = \ 0)\right\}\right| \ = \ 1$$

Hence the correct choice of k (in forming χ_k) is made with probability (at least) $1/n$, from which it follows that with probability at least $1/4n$ exactly one such candidate will survive as a member of $C(\chi_k)$. In consequence χ_k will be accepted as an instance of $\text{USAT}_2^\exists$. $\square$