The Computational Complexity of Ideal Semantics I: Abstract Argumentation Frameworks

Paul E. Dunne
Department of Computer Science, The University of Liverpool, U.K.

Abstract

We analyse the computational complexity of the recently proposed *ideal semantics* within abstract argumentation frameworks. 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 admissibility semantics. In particular the task of *finding* the unique maximal 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. Finally we present a number of technical results which offer strong indications that typical problems in ideal argumentation are complete for the class $\mathrm{P}^{\mathrm{NP}}_{||}$: languages decidable by polynomial time algorithms allowed to make non-adaptive queries to an NP oracle.

1 Introduction

The extension-based semantics defined by *ideal extensions* were introduced by Dung, Mancarella and Toni [11, 12] as an alternative sceptical basis for defining collections of justified arguments in the frameworks promoted by Dung [10] and Bondarenko *et al.* [3].

Our principal concern in this article is in classifying the computational complexity of a number of natural problems related to ideal semantics in abstract argumentation frameworks such problems including 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 maximal ideal extension?, i.e. without, necessarily, being an ideal extension itself.

b2. Is S, itself, an ideal extension?

b3. Is S the *maximal* ideal extension?

c. Is the maximal ideal extension empty?

d. Does the maximal ideal extension coincide with the set of all *sceptically accepted* arguments?

e. Given an argumentation framework, construct its maximal ideal extension.

We obtain bounds for these problems ranging NP, CO-NP, and D^p-hardness through to an exact $FP_{||}^{NP}$–completeness classification for the construction problem defined in (e). In the remainder of this paper, background definitions are given in Section 2 together with formal definitions of the problems introduced in (a)–(e) above. In Section 3, two technical lemmata are given which characterise properties of ideal extensions (Lemma 1) and of *arguments* belonging to the *maximal* ideal extension (Lemma 2). The complexity of decision questions is considered in Section 4 while Section 5 provides details of efficient solution approaches for the special case of *bipartite* argumentation frameworks, a class whose properties have previously been studied in [13]. Our main technical result is presented in Section 6 wherein an exact classification for the complexity of *finding* the maximal ideal extension is given. One consequence of this result is that (under the usual complexity-theoretic assumptions) *constructing* the maximal ideal extension of a given framework is, in general, easier than *deciding* if one of its arguments is sceptically accepted.

The results of Section 4 leave a gap between lower (hardness) bounds and upper bounds for a number of the decision questions. In Section 7 we present strong evidence that problems (a), (b1), (b3) and (c) are not contained in any complexity class strictly below $P_{||}^{NP}$: specifically that all of these problems are $P_{||}^{NP}$–hard via *randomized reductions* which are correct with probability approaching 1.

2 Abstract Argumentation Frameworks

The following concepts were introduced in Dung [10].

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*

a. $s \in S$ *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.*

b. $x \in \mathcal{X}$ *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*.*

c. S *is* conflict-free *if no argument in* S *is attacked by any other argument in* S*.*

d. A conflict-free set S *is* admissible *if every* $y \in S$ *is acceptable w.r.t* S*.*

e. S *is a* preferred extension *if it is a maximal (with respect to* $\subseteq$*) admissible set.*

f. S *is a* stable extension *if* S *is conflict free and every* $y \notin S$ *is attacked by* S*.*

g. $\mathcal{H}$ *is* coherent *if every preferred extension in* $\mathcal{H}$ *is also a stable extension.*

h. S *is an* ideal extension *([11, 12]) of* $\mathcal{H}$ *if* S *is admissible* and *a subset of every preferred extension of* $\mathcal{H}$*.*

i. $\mathcal{H}$ *is* cohesive *if its maximal ideal extension coincides with the intersection of all preferred extensions of* $\mathcal{H}$*, i.e. if the set*

$$\bigcap_{S \subseteq \mathcal{X} \;:\; S \text{ is preferred}} S$$

is admissible.

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

$$\begin{array}{lll} S^- & =_{\text{def}} & \{\, p \;:\; \exists\, q \in S \;\; \textit{such that } \langle p, q \rangle \in \mathcal{A} \} \\ S^+ & =_{\text{def}} & \{\, p \;:\; \exists\, q \in S \;\; \textit{such that } \langle q, p \rangle \in \mathcal{A} \} \end{array}$$

An argument x *is* credulously accepted *if there is* some *preferred extension containing it;* x *is* sceptically accepted *if it is a member of* every *preferred extension.*

Decision Problem	Instance	Question
CA	$\mathcal{H}(\mathcal{X},\mathcal{A})$, $x \in \mathcal{X}$	Is x credulously accepted in $\mathcal{H}$?
SA	$\mathcal{H}(\mathcal{X},\mathcal{A})$, $x \in \mathcal{X}$	Is x sceptically accepted in $\mathcal{H}$?

Table 1: Decision Problems in Argumentation Frameworks

The concepts of credulous and sceptical acceptance motivate the decision problems of Table 1 that have been considered in [9, 14].

These questions are formulated in terms of *single* arguments, it will be useful to consider analogous concepts with respect to *sets*. Thus $\text{CA}_{\{\}}$ denotes the decision problem whose instances are an AF $\langle \mathcal{X},\mathcal{A}\rangle$ together with a subset S of $\mathcal{X}$: the instance being accepted if there is a preferred extension T for which $S \subseteq T$. Similarly, $\text{SA}_{\{\}}$ accepts instances for which S is a subset of *every* preferred extension.

The results of [9] establish that CA is NP–complete, while [14] proves SA to be Π_2^p–complete.

We consider a number of decision problems relating to properties of ideal extensions in argumentation frameworks as described in Table 2.

Problem Name	Instance	Question
IE	$\mathcal{H}(\mathcal{X},\mathcal{A})$; $S \subseteq \mathcal{X}$	Is S an ideal extension?
IA	$\mathcal{H}(\mathcal{X},\mathcal{A})$; $x \in \mathcal{X}$	Is x in the maximal ideal extension?
$\text{MIE}_{\emptyset}$	$\mathcal{H}(\mathcal{X},\mathcal{A})$	Is the maximal ideal extension empty?
MIE	$\mathcal{H}(\mathcal{X},\mathcal{A})$; $S \subseteq \mathcal{X}$	Is S the *maximal* ideal extension?
CS	$\mathcal{H}(\mathcal{X},\mathcal{A})$	Is $\mathcal{H}(\mathcal{X},\mathcal{A})$ *cohesive*?

Table 2: Decision questions for Ideal Semantics

We also consider *search* (so-called *function problems*) where the aim is not simply to verfify that a given set has a specific property but to *construct* an example. In particular we examine the function problem FMIE in which, given an AF $\mathcal{H}(\mathcal{X},\mathcal{A})$, it is required to return the maximal 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 $L_2 \in \text{co-NP}$. 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 [17, 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 Δ_2^p and $\mathrm{F}\Delta_2^p$ (sometimes denoted P^{NP} and $\mathrm{FP}^{\mathrm{NP}}$) consist of those decision problems (respectively function or *search* 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).[1] An important (presumed) subset of Δ_2^p 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 denoted Θ_2^p and considered in Wagner [20, 21], Jenner and Toran [15]: we use the notation $\mathrm{FP}_{||}^{\mathrm{NP}}$ adopted for the function computation class in [15], i.e. $\mathrm{FP}_{||}^{\mathrm{NP}} \equiv_{\mathrm{def}} \mathrm{F}\Theta_2^p$. Under the standard complexity-theoretic assumptions, it is conjectured that,

$$\mathrm{P} \subset \left\{ \begin{array}{c} \mathrm{NP} \\ \mathrm{CO\text{-}NP} \end{array} \right\} \subset \mathrm{D}^p \subset \Theta_2^p \subset \Delta_2^p \subset \left\{ \begin{array}{c} \Sigma_2^p \\ \Pi_2^p \end{array} \right\}$$

We prove the following complexity classifications.

a. IE is CO-NP–complete.

b. IA is CO-NP–hard.

c. $\mathrm{MIE}_\emptyset$ is NP–hard.

d. MIE is D^p–hard.

e. CS is Σ_2^p–complete.

f. FMIE is $\mathrm{FP}_{||}^{\mathrm{NP}}$–complete.

g. All the problems (a)–(f) are polynomial time solvable for *bipartite* frameworks.

3 Characteristic properties of ideal extensions

The upper bound proofs exploit firstly a characterisation of ideal extensions in terms of credulous acceptability presented in Lemma 1. We also present,

[1] We refer the reader to e.g. [17, pp. 415–423] for further background concerning these classes.

in Lemma 2, a necessary and sufficient condition for a given *argument* to be a member of the maximal ideal extension.

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

I1. 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. $\forall\ p \in S^-\ \ \neg\textsc{ca}(\mathcal{H}, p)$.*

Proof:
($\Rightarrow$) Suppose that $S \subseteq \mathcal{X}$ is an ideal extension of $\mathcal{H}$. It is immediate from the definition of ideal that S is admissible so that (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 extension.

($\Leftarrow$) 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 extension. 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 the maximal ideal extension of $\mathcal{H}(\mathcal{X},\mathcal{A})$. 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. $\forall\ y \in \{x\}^-\ \neg\textsc{ca}(\mathcal{H}, y)$.

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 maximal ideal extension of $\mathcal{H}$. Since $\mathcal{M}$ is an ideal extension, 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 extension, 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 extension 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 extension. The set $\mathcal{M}$ is, however, already *maximal* so that $\mathcal{M} \cup \{x\} = \mathcal{M}$, i.e. $x \in \mathcal{M}$ as required. □

4 Decision questions in ideal argumentation

Theorem 1 IE *is co*-NP*–complete.*

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

$$\textsc{adm}(\mathcal{H}, S)\ \wedge \bigwedge_{q \in S^-} \neg\textsc{ca}(\mathcal{H}, q)$$

where $\text{ADM}(\mathcal{H}, S)$ is the predicate returning $\top$ if and only if S is an admissible set of $\mathcal{H}$. Correctness follows from Lemma 1 and since $\text{CA}(\mathcal{H}, q)$ is decidable in NP, its complement is decidable by a co-NP algorithm.[2]

To prove IE is co-NP–hard we reduce from CNF-UNSAT (without loss of generality, restricted to instances which are 3-CNF). Given a 3-CNF formula

$$\Phi(x_1, \ldots, x_n) \quad = \quad \bigwedge_{i=1}^{m} C_i \quad = \quad \bigwedge_{i=1}^{m} (z_{i,1} \vee z_{i,2} \vee z_{i,3})$$

as an instance of UNSAT we form an instance $\langle \mathcal{F}_\Phi, S \rangle$ of IE as follows. First construct the AF $\mathcal{H}_\Phi$ (described in Appendix A.1) from the CNF Φ. In this, via [9, 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 $\Phi(Z_n)$ is unsatisfiable. The AF, $\mathcal{F}_\Phi$, is formed from $\mathcal{H}_\Phi$ by adding an argument Ψ together with attacks

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

The instance of IE is completed by setting $S = \{\Psi\}$.

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

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

$$\neg\text{CA}(\mathcal{F}_\Phi, \Phi) \quad \Leftrightarrow \quad \neg\text{CA}(\mathcal{H}, \Phi) \quad \Leftrightarrow \quad \text{UNSAT}(\Phi)$$

□

Corollary 1 IA *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. 1, defines a positive instance of IA if and only if $\Phi(Z_n)$ is unsatisfiable. □

Corollary 2 $\text{MIE}_\emptyset$ *is* NP*–hard*

Proof: The AF $\mathcal{F}_\Phi$ defined in Thm. 1 has an empty maximal ideal extension if and only if $\Phi(Z_n)$ is satisfiable. □

[2] The form $\bigwedge_{q \in S^-} \neg\text{CA}(\mathcal{H}, q)$ is equivalent to $\forall \, T \ \text{ADM}(\mathcal{H}, T) \Rightarrow (T \cap S^- = \emptyset)$ so that it is not necessary to use $|S|$ distinct co-NP tests: a form which would preclude membership proper in the class co-NP.

Theorem 2 MIE *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. 1 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 MIE 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 maximal ideal extension, M, of $\mathcal{F}_{\langle \Phi_1, \Phi_2 \rangle}$: $M = \emptyset$ (both Φ_1 and Φ_2 are satisfiable); $M = \{\Psi_1, \Psi_2\}$ (neither formula is satisfiable); $M = \{\Psi_1\}$ (Φ_1 is unsatisfiable and Φ_2 is satisfiable; $M = \{\Psi_2\}$ (Φ_1 is satisfiable and Φ_2 is unsatisfiable). Only the final case corresponds with the set given in the constructed instance. □

Theorem 3 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

$$\text{ADM}\left(\mathcal{H}(\mathcal{X}, \mathcal{A}), \bigcap_{S \subseteq \mathcal{X} \,:\, S \text{ is preferred}} S \right)$$

which can be tested by checking

$$\exists\ S\ \text{IE}(\mathcal{H}, S) \quad \wedge \bigwedge_{x \in \mathcal{X} \setminus S} \neg\text{SA}(\mathcal{H}, x) \tag{1}$$

That is, there is a subset (S) of $\mathcal{X}$ which defines an ideal extension of $\mathcal{H}$ and for which no argument outside S is in every preferred extension. From Thm. 1, IE is in co-NP; in addition since SA $\in \Pi_2^p$ its complement $\neg$SA is in Σ_2^p hence (1) gives a Σ_2^p test for CS.[3]

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 [14] from the complementary problem – QSAT_2^Π, details of which are presented in Appendix A.2.

[3] 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 $\neg\text{SA}(\mathcal{H}, x)$ is replaced by verifying that U_x is a preferred extension of $\mathcal{H}$ (co-NP) and that $x \notin U_x$.

Given an instance $\Phi(Y_n, Z_n)$ of QSAT$_2^\Sigma$, consider the AF $\mathcal{G}_\Phi$, defined from this as described in Appendix A.2. Noting that the maximal ideal extension of $\mathcal{G}_\Phi$ is the empty set (which is always an ideal extension), 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$_2^\Sigma$.

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

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

where R_β denotes the subset of $\{\ z_i, \neg z_i\ \ : 1\ \leq\ i\ \leq\ n\}$ 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$_2^\Sigma$. 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.*

Proof: From [14, Corollary 18], the AF $\mathcal{G}_\Phi$ used in the proof of Thm. 3 is coherent if and only the argument Φ is sceptically accepted. Thus $\mathcal{G}_\Phi$ is coherent if and only if it is *not* cohesive. □

5 Frameworks with Efficient Algorithms

Theorem 4 *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

$$\mathcal{M} \quad = \bigcap_{S \subseteq \mathcal{Y} \cup \mathcal{Z}\ : S \text{ is a preferred extension of } \mathcal{B}} S$$

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:

$$\begin{array}{lcl} \mathcal{Y}_{\mathrm{SA}} & = & \mathcal{Y} \cap \mathcal{M} \\ \mathcal{Y}_{\mathrm{CA}} & = & \{ \ y \in \mathcal{Y} \ : \ \textsc{ca}(\mathcal{B}, y) \ \} \setminus \mathcal{M} \\ \mathcal{Y}_{\mathrm{OUT}} & = & \{ \ y \in \mathcal{Y} \ : \ \neg\textsc{ca}(\mathcal{B}, y)\} \\ \mathcal{Z}_{\mathrm{SA}} & = & \mathcal{Z} \cap \mathcal{M} \\ \mathcal{Z}_{\mathrm{CA}} & = & \{ \ z \in \mathcal{Z} \ : \ \textsc{ca}(\mathcal{B}, z) \ \} \setminus \mathcal{M} \\ \mathcal{Z}_{\mathrm{OUT}} & = & \{ \ z \in \mathcal{Z} \ : \ \neg\textsc{ca}(\mathcal{B}, z)\} \end{array}$$

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

$$\begin{array}{ccc} \mathcal{Y}^{-}_{\mathrm{SA}} \subseteq \mathcal{Z}_{\mathrm{OUT}} & ; & \mathcal{Y}^{+}_{\mathrm{SA}} \subseteq \mathcal{Z}_{\mathrm{OUT}} \\ \mathcal{Z}^{-}_{\mathrm{SA}} \subseteq \mathcal{Y}_{\mathrm{OUT}} & ; & \mathcal{Z}^{+}_{\mathrm{SA}} \subseteq \mathcal{Y}_{\mathrm{OUT}} \end{array}$$

In addition,

$$\begin{array}{l} \forall \ y \in \mathcal{Y}_{\mathrm{CA}} \ \exists \ z \in \mathcal{Z}_{\mathrm{CA}} \ \ \langle z, y\rangle \in \mathcal{A} \\ \forall \ z \in \mathcal{Z}_{\mathrm{CA}} \ \exists \ y \in \mathcal{Y}_{\mathrm{CA}} \ \ \langle y, z\rangle \in \mathcal{A} \end{array}$$

To see this, suppose without loss of generality, that $\mathcal{Z}_{\mathrm{CA}}$ does not attack $y \in \mathcal{Y}_{\mathrm{CA}}$: then since $\neg$ *attacks*$(\mathcal{Z}_{\mathrm{SA}} \cup \mathcal{Z}_{\mathrm{CA}}, y)$ 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 maximal ideal extension of $\mathcal{B}(\mathcal{Y}, \mathcal{Z}, \mathcal{A})$ may be constructed in polynomial time.*

Proof: From Thm. 4 the maximal ideal extension of $\mathcal{B}$ corresponds with the set of all sceptically accepted arguments of $\mathcal{B}$. Applying the methods described in [13, Thm. 6] this set can be identified in polynomial time. □

We note that as a consequence of Thm. 4 and Corollary 4 in the case of bipartite AFs, the decision problems IE, MIE and IA are all in P and CS is trivial.

6 Finding the Maximal Ideal Extension

Theorem 5 FMIE *is* $\text{FP}_{||}^{\text{NP}}$*-complete.*

Proof: We first present the argument that FMIE 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 $\chi(\Xi) = c_1 c_2 c_3 \cdots c_r \in [0, 2^r - 1]$ 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 maximal ideal extension of this framework, from Thm. 2, 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 FMIE is $\text{FP}_{||}^{\text{NP}}$–hard.

To see that FMIE $\in \text{FP}_{||}^{\text{NP}}$ 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. 4),

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

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 $\mathcal{X}_{\text{OUT}}$.[4]

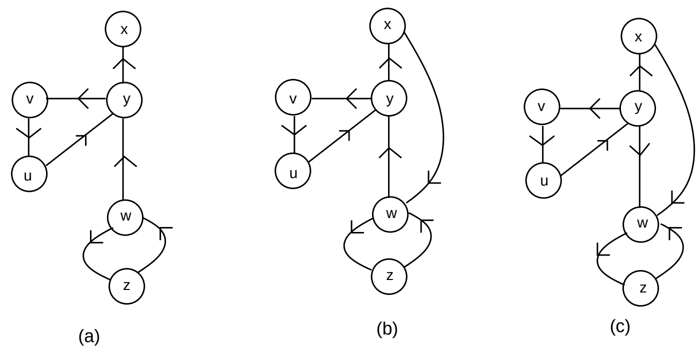

Figure 1:
a) $\mathcal{X}_{\mathrm{SA}} = \emptyset \subset \{x, v\} = \mathcal{X}_{\mathrm{PSA}}$; $\mathcal{X}_{\mathrm{OUT}} = \{y, u\}$; $\mathcal{X}_{\mathrm{CA}} = \{w, z\}$.
b) $\mathcal{X}_{\mathrm{PSA}} = \{v\}$; $\mathcal{X}_{\mathrm{OUT}} = \{y, u\}$; $\mathcal{X}_{\mathrm{CA}} = \{x, w, z\}$.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

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

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

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

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

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

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

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 maximal 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 [13, Thm. 6(a)] this set can be found in polynomial time and thus FMIE $\in$ $\text{FP}_{||}^{\text{NP}}$ as claimed. □

The technique employed to established $\text{FP}_{||}^{\text{NP}}$–hardness in proving Thm. 5 can be used to demonstrate $\text{P}_{||}^{\text{NP}}$–hardness for a number of (admittedly rather artificial) decision problems concerning properties of the maximal ideal extension. For example,

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

Proof: Membership is immediate from the construction of Thm. 5. Hardness follows from the fact – [20, 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. 5. □

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 maximal 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 $\forall x \in \mathcal{X}$ CA$(\mathcal{H}, x)$ 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 $\mathcal{X}_{\text{PSA}} = \{ x \ : \ \text{SA}(\mathcal{H}, x)\}$ and $\mathcal{X}_{\text{PSA}}$ is the maximal ideal extension, i.e. $\mathcal{H}$ is cohesive. □

Combining the results we obtain the picture of the relative complexities of checking whether an argument is credulously/ideally/sceptically acceptable shown in Table 3.

Acceptability Semantics	Lower bound	Upper Bound
CA	NP–hard	NP
IA	CO-NP–hard	$\text{P}_{\|\|}^{\text{NP}}$
SA	Π_2^p–hard	Π_2^p

Table 3: Relative Complexity of Testing Acceptability.

Similarly, Table 4 considers checking whether a given *set* of arguments collectively satisfies the requirements of a given semantics or is a maximal such set.

Semantics	Decision Problem	Lower bound	Upper Bound
Credulous	ADM	P	P
Ideal	IE	CO-NP–hard	CO-NP
Sceptical	$\text{SA}_{\{\}}$	Π_2^p–hard	Π_2^p
Credulous	PREF-EXT	CO-NP–hard	CO-NP
Ideal	MIE	D^p–hard	$\text{P}_{\|\|}^{\text{NP}}$
Sceptical	MAS	D_2^p–hard	D_2^p

Table 4: Deciding set and maximality properties

We note in passing that the problem MAS 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 $L_2 \in \Pi_2^p$.

Table 5 completes this overview giving the complexity of constructing maximal sets:[5]

Semantics	Lower bound	Upper Bound
Ideal	$\text{FP}_{\|\|}^{\text{NP}}$–hard	$\text{FP}_{\|\|}^{\text{NP}}$
Sceptical	$\text{FP}_{\|\|}^{\Sigma_2^p}$–hard	$\text{FP}_{\|\|}^{\Sigma_2^p}$

Table 5: Complexity of constructing maximal extensions

In this table $\text{FP}_{||}^{\Sigma_2^p}$ is the analogous class to $\text{FP}_{||}^{\text{NP}}$ in which polynomially many queries may be made (in parallel) to a Σ_2^p oracle. In total the classifications given by the three tables above reinforce the the case that deciding

[5]We do not formulate this problem for credulous semantics since, in contrast to ideal and sceptical semantics, $\mathcal{H}$ will, in general, not have a *unique* maximal admissible set, so that the corresponding function is *multi-valued*.

credulous admissibility semantics is easier than ideal semantics which, in turn, is easier than sceptical admissibility semantics.

7 Reducing the complexity gaps

In [4], Chang and Kadin introduce the concepts of a language having the properties $\textsc{op}_2$ and $\textsc{op}_\omega$ where $\textsc{op}$ is one of the Boolean operators $\{\textsc{and}, \textsc{or}\}$. Formally,

Definition 2 *([4, pp. 175–76] Let L be a language, i.e. a set of finite words over an alphabet. The languages, $\textsc{and}_k(L)$ and $\textsc{or}_k(L)$ $(k \geq 1)$ are*

$$\begin{array}{lll} \textsc{and}_k(L) & =_{\text{def}} & \{\langle w_1, w_2, \ldots, w_k\rangle \ : \ \forall\ 1 \leq i \leq k\ w_i \in L\} \\ \textsc{or}_k(L) & =_{\text{def}} & \{\langle w_1, w_2, \ldots, w_k\rangle \ : \ \exists\ 1 \leq i \leq k\ w_i \in L\} \end{array}$$

The languages $\textsc{and}_\omega(L)$ and $\textsc{or}_\omega(L)$ are,

$$\textsc{and}_\omega(L) \ =_{\text{def}} \ \bigcup_{k\geq 1} \textsc{and}_k(L) \quad ; \quad \textsc{or}_\omega(L) \ =_{\text{def}} \ \bigcup_{k\geq 1} \textsc{or}_k(L)$$

A language, L, is said to have property $\textsc{op}_k$ (resp. $\textsc{op}_\omega$) if $\textsc{op}_k(L) \ \leq^p_m \ L$ (resp. $\textsc{op}_\omega(L) \ \leq^p_m \ L$).

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

Fact 1 *([4, Thm. 9, p. 182])*
A language L is $\textsc{p}^{\text{NP}}_{||}$–complete (via $\leq^p_m$ reducibility) if and only if all of the following hold.

F1. $L \in \textsc{p}^{\text{NP}}_{||}$.

F2. L *is* NP*–hard* and L *is co-*NP*–hard.*

F3. L *has property* $\textsc{and}_2$.

F4. L *has property* $\textsc{or}_\omega$.

As a consequence of Fact 1, we have,

Theorem 6

a. If IA *is* NP*–hard then* IA *is* $\textsc{p}^{\text{NP}}_{||}$*–complete.*

b. If IA $\in$ *co-*NP *then* MIE *is* $\textsc{d}^p$*–complete.*

c. If IA $\in$ *co-*NP *then* $\text{MIE}_\emptyset$ *is* NP*–complete.*

d. If MIE *has property* OR_ω *then* MIE *is* $\text{P}_{||}^{\text{NP}}$*–complete.*

e. If $\text{MIE}_\emptyset$ *is co-*NP*–hard then* $\text{MIE}_\emptyset$ *is* $\text{P}_{||}^{\text{NP}}$*–complete.*

Proof: (Outline)

a. With the assumption that IA is NP–hard, IA would satisfy conditions (F1) and (F2) of Fact 1. To complete the argument it suffices to show that IA already has property AND_2 and property OR_ω. For the first of these consider any instance $\langle\langle\mathcal{H}_1, x\rangle, \langle\mathcal{H}_2, y\rangle\rangle$ of $\text{AND}_2(\text{IA})$. Form the 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 IA if and only if $\text{IA}(\langle\mathcal{H}_1, x\rangle) \wedge \text{IA}(\langle\mathcal{H}_2, y\rangle)$. To see that IA has property OR_ω 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 $\text{OR}_\omega(\text{IA})$. Form an AF, $\mathcal{H}$, from these m frameworks, adding two new arguments, $\{y, z\}$. The instance of IA 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 IA if and only if $\vee_{i=1}^{m}\ \text{IA}(\langle\mathcal{H}_i, x_i\rangle)$.

b. It has already been shown that MIE is D^p–hard. Consider the languages,

$$
\begin{array}{lll}
L_1 & =_{\text{def}} & \{\langle\mathcal{H}, S\rangle\ :\ \forall\ x \in S,\ \text{IA}(\mathcal{H}, x)\ \} \\
L_2 & =_{\text{def}} & \{\langle\mathcal{H}, S\rangle\ :\ \forall\ x \notin S,\ \neg\text{IA}(\mathcal{H}, x)\ \}
\end{array}
$$

We have $L_1 \in$ co-NP (by the assumption IA is in co-NP and by the straighforward generalisation of (a) that shows IA has property AND_ω). In addition, $L_2 \in$ NP (from the premise IA $\in$ co-NP and the fact that $\neg$IA has property AND_ω since IA has property OR_ω). With these choices of L_1 and L_2, $\langle\mathcal{H}, S\rangle$ is accepted as an instance MIE if and only if $\langle\mathcal{H}, S\rangle \in L_1 \cap L_2$ so that MIE $\in \text{D}^p$.

c. Easy consequence of (b).

d. It has already been shown that MIE satisfies (F1) and (F2) of Fact 1. In addition, MIE has property AND_ω (thus, trivially, also 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 $\text{AND}_\omega(\text{MIE})$ 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 MIE if and only if $\wedge_{i=1}^{m}\ \text{MIE}(\mathcal{H}_i, S_i)$. It follows that were MIE to have property OR_ω, then MIE would be $\text{P}_{||}^{\text{NP}}$-complete via Fact 1.

e. It suffices to show $\text{MIE}_\emptyset$ has OR_ω. Given $\langle \mathcal{H}_1, \mathcal{H}_2, \ldots, \mathcal{H}_m \rangle$ form the AF, $\mathcal{H}$, from these together with additional arguments $\{y_1, y_2, \ldots, y_m, z\}$. Introduce attacks $\mathcal{A}_i = \{\langle x, y_i \rangle \ : \ \forall\ x \in \mathcal{H}_i\}$ $(1 \leq i \leq m)$ together with $\{\langle y_i, z \rangle \ : \ 1 \leq i \leq m\}$. From Lemma 2, $\mathcal{H}$ is accepted as an instance of $\text{MIE}_\emptyset$ if and only if $\vee_{i=1}^{m}\ \text{MIE}_\emptyset(\mathcal{H}_i)$.

□

We may interpret Thm. 6 as focusing the issue of obtaining exact classifications in terms of IA. If IA $\in$ co-NP (so that, with the usual assumption of NP$\neq$co-NP, IA would not be NP–hard) then we obtain exact classifications of the complexity of $\{\text{IA}, \text{MIE}, \text{MIE}_\emptyset\}$ as $\{\text{co-NP}, \text{D}^p, \text{NP}\}$–complete. On the other hand, an alternative hypothesis, in the event of IA $\notin$ co-NP, is that suggested by Thm. 6 (a): that IA is $\text{P}_{||}^{\text{NP}}$–complete, a result which would follow by demonstrating IA to be NP–hard.

In fact, there is strong evidence that IA $\notin$ co-NP and, using one suite of techniques is more likely to be complete within $\text{P}_{||}^{\text{NP}}$. Our formal justification of these claims rests on a number of technical analyses using results of Chang *et al.* [5], which in turn develop ideas of [1, 2, 18]. Two key concepts in our further analyses of IA 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 instantatiation?

Determining the exact complexity of USAT remains an open problem. It is known that USAT $\in \text{D}^p$ and while Blass and Gurevich [2] show it to be co-NP–hard[6], USAT has only be shown to be complete for D^p using a *randomized* reduction technique of Valiant and Vazirani [18]. Two concepts of such reductions are studied in Chang *et al.* [5] specifically with respect to USAT via the following general definition.

Definition 3 *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*

[6] The reader should note that [16, p. 93] has a typographical slip whereby Blass and Gurevich's result is described as proving USAT to be NP–hard.

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* $\langle 0,1\rangle^{q(|x|}$

$$\begin{array}{lcl} x \in L_1 & \Rightarrow & Prob[f(x,z) \in L_2] \ \geq \ \delta \\ x \notin L_1 & \Rightarrow & Prob[f(x,z) \notin L_2] \ = \ 1 \end{array}$$

We have the following properties of USAT and randomized reductions:

Fact 2

a. SAT $\leq_m^{rp}$ USAT *with probability* $1/(4n)$*. ([18, 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* $1-2^{-n}$*. ([5, Fact 1, p. 361][7])*

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

Theorem 7 USAT $\leq_m^p$ IA.

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. 1, 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.[8] 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 IA is $\langle\mathcal{K}(\mathcal{X},\mathcal{A}),x\rangle$ and the resulting AF is illustrated in Fig. 2.

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 maximal ideal 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

[7] The bound actually stated in [5] is for arbitrary exponentially decreasing functions, i.e. not just 2^{-n}.

[8] 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 $\{\langle C_j,d_j\rangle,\langle d_j,e_j\rangle,\langle e_j,C_j\rangle\}$.

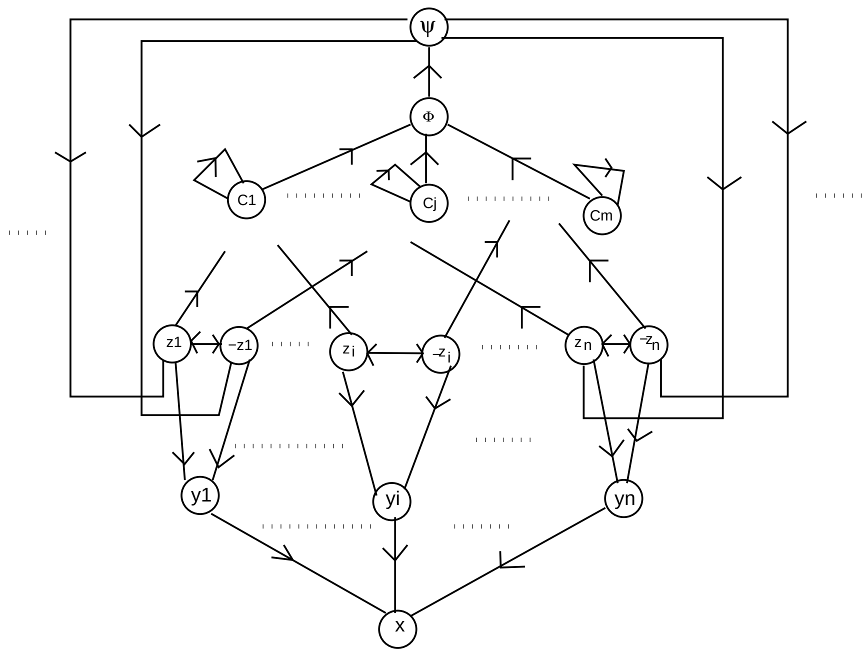

Figure 2: The Argumentation Framework $\mathcal{K}_\Phi$

accepted, in particular, each of the arguments z_i and $\neg z_i$ fail to be accepted. It easily follows that $x \notin \mathcal{M}_\mathcal{K}$ 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 $\alpha = \langle a_1, a_2, \ldots, a_n \rangle$ 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 IA.

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 maximal ideal extension of $\mathcal{K}(\mathcal{X}, \mathcal{A})$. □

Combining Thms. 6 and 7 with Facts 1 and 2 gives the following corollaries.

Corollary 7 USAT $\leq_m^p$ $\neg$MIE$_\emptyset$

Proof: The AF $\mathcal{K}_\Phi$ of Thm. 7 has a *non-empty* maximal ideal extension, $\mathcal{M}_\mathcal{K}$, if and only if $x \in \mathcal{M}_\mathcal{K}$. □

Corollary 8 IA *is complete for* $\mathrm{P}_{||}^{\mathrm{NP}}$ *under* $\leq_m^{rp}$ *reductions with probability* $1 - 2^{-n}$.

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

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

(as observed in [5, 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 OR$_\omega$(SAT-UNSAT).

Now, via [18], SAT-UNSAT $\leq_m^{rp}$ USAT with probability $1/(4n)$ so that, combining these randomized reductions,

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

Now applying the (deterministic) reduction of Thm. 7 shows

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

As demonstrated in the proof of Thm. 6(a), IA has property OR$_\omega$ so that via Fact 2(b) we obtain,

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

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

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

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

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

Proof: Trivial consequence of Corollary 9 since MIE$_{\emptyset}$ is a restricted special case of MIE. ▫

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

Corollary 11 *If* IA $\leq_m^p$ USAT *(note* deterministic *reduction) then the Polynomial Hierarchy (*PH*) collapses to* Σ_3^p, *i.e.*

$$\text{IA} \leq_m^p \text{USAT} \quad \Rightarrow \quad \bigcup_{k\geq 3} \Sigma_k^p \cup \bigcup_{k\geq 3} \Pi_k^p \quad \subseteq \quad \Sigma_3^p$$

Proof: Suppose it is the case that IA $\leq_m^p$ USAT. We then have

$$\begin{array}{lll l} \text{OR}_\omega(\text{USAT}) & \leq_m^p & \text{OR}_\omega(\text{IA}) & \text{by Thm. 7} \\ & \leq_m^p & \text{IA} & \text{since IA has property OR}_\omega \\ & \leq_m^p & \text{USAT} & \text{by premise} \end{array}$$

So that USAT would have property OR$_\omega$: [5, Thm. 5, p. 364] demonstrates that this leads to the collapse stated. ▫

Now, noting that $\leq_m^p$ reductions can be interpreted as "$\leq_m^{rp}$ reductions with probability 1", we can reconsider the lower bounds of Tables 3 and 4 using hardness via $\leq_m^{rp}$ reductions (with "high" probability) instead of hardness via deterministic $\leq_m^p$ reducibility, as shown in Table 6.

Decision Problem	Complexity	$\leq_m^{rp}$ probability
CA	NP–complete	1
IA	$\mathrm{P}_{\|\|}^{\mathrm{NP}}$–complete	$1 - 2^{-n}$
SA	Π_2^p–complete	1
ADM	P	–
IE	co-NP–complete	1
$\mathrm{SA}_{\{\}}$	Π_2^p–complete	1
PREF-EXT	co-NP–complete	1
$\mathrm{PREF\text{-}EXT}_{\emptyset}$	co-NP–complete	1
MIE	$\mathrm{P}_{\|\|}^{\mathrm{NP}}$–complete	$1 - 2^{-n}$
$\mathrm{MIE}_{\emptyset}$	$\mathrm{P}_{\|\|}^{\mathrm{NP}}$–complete	$1 - 2^{-n}$
MAS	D_2^p–complete	1

Table 6: Complexity of ideal semantics relative to randomized reductions

8 Conclusions and Further Work

We have considered the computational complexity of decision and search problems arising in the Ideal semantics for abstract argumentation frameworks introduced in [11, 12]. It has been shown that all but one of these[9] can be resolved within $\mathrm{P}_{||}^{\mathrm{NP}}$ or its functional analogue $\mathrm{FP}_{||}^{\mathrm{NP}}$: classes believed to lie strictly below the second level of the polynomial hierarchy. 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 maximal ideal extension, and if the maximal ideal extension is empty, are not contained within any complexity class falling strictly within $\mathrm{P}_{||}^{\mathrm{NP}}$: all of these problems being $\mathrm{P}_{||}^{\mathrm{NP}}$–hard with respect to $\leq_m^{rp}$ reductions of probability $1-2^{-n}$. Although this complexity class compares unfavourably with the NP and co-NP–complete status of related questions under the credulous preferred semantics of [10], it represents an improvement on the Π_2^p–completeness level of similar issues within the sceptical preferred semantics.

Given that a precondition of ideal acceptance is that the argument is sceptically accepted this reduction in complexity may appear surprising. 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 extension: as well as being sceptically accepted, the set must be admissible. This condition plays a significant role in the complexity shift. An impor-

[9] The exception being the problem CS.

tant reason why testing sceptical acceptance of a gven argument x fails to belong to co-NP (assuming co-NP $\neq \Pi_2^p$) 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 [19] 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 [14]. In contrast, as demonstrated in the characterisation of ideal extensions 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 co-NP.

The reason why *finding* the maximal 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. Some further properties of this partition are reviewed in Appendix B: among these it is shown that deciding if $S = \mathcal{X}_{\mathrm{OUT}}$ given $\langle \mathcal{H}(\mathcal{X}, \mathcal{A}), S \rangle$ is D^p–complete.

This paper has focussed on the graph-theoretic abstract argumentation framework model from [10]. A natural continuation, and the subject of current work, is to consider the divers instantiations of *assumption-based* argumentation frameworks (ABF) [3]. Complexity-theoretic analyses of this model with respect to credulous and sceptical semantics have been presented in a series of papers by Dimopoulos, Nebel, and Toni [6, 7, 8]. In these, the computational complexity of specific decision questions is shown to be linked with that of deciding $\Delta \models \varphi$ where Δ is a given collection of formulae and $\models$ is a derivability relation whose precise semantics are dependent on the logical theory described by the ABF, e.g. [3] describe how ABFs may be formulated to capture a variety of non-classical logics. While one might reasonably expect a number of the techniques described above to translate to ABF instantiations, there are non-trivial issues for cases where $\Delta \models \varphi$ is NP–complete, e.g. the default logic instantiation of ABFs. A significant problem in such cases concerns the mechanisms explored in Section 7 in order to amplify the co-NP–hardness (via $\leq_m^p$ reductions) of IA to $\mathrm{P}_{||}^{\mathrm{NP}}$. For example, the reduction from USAT and the concept of languages with property OR_ω: whether a "natural" analogue of USAT for the second level of the polynomial hierarchy can be formulated is far from clear.

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

[3] 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.

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

[5] 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.

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

[7] 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.

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

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

[10] 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.

[11] 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.

[12] P. M. Dung, P. Mancarella, and F. Toni. Computing ideal sceptical argumentation. *Artificial Intelligence*, 2007. (in press), doi:10.1016/j.artint.2007.05.003.

[13] P. E. Dunne. Computational properties of argument systems satisfying graph-theoretic constraints. *Artificial Intelligence*, 2007. (in press), doi:10.1016/j.artint.2007.03.006.

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

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

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

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

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

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

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

[21] 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 [9, 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

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

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

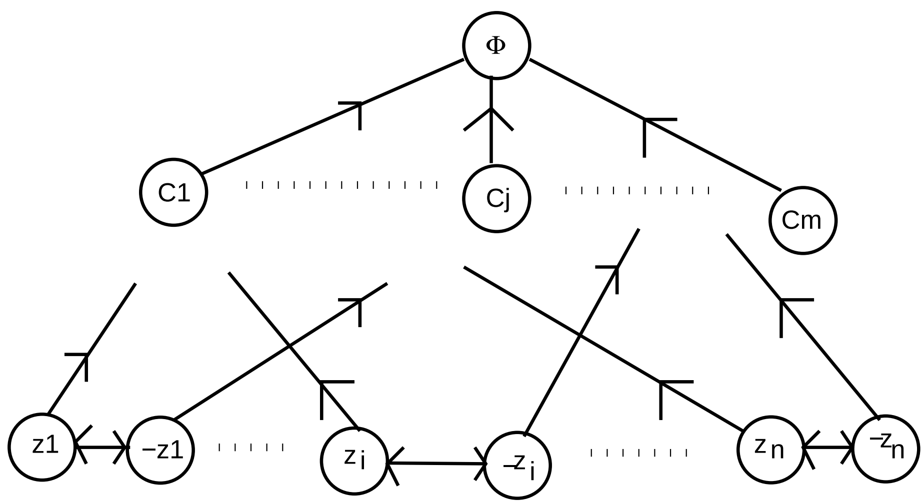

Figure 3: The Argumentation Framework $\mathcal{H}_\Phi$

Fact 3 *(Dimopoulous and Torres [9]) 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 [14] uses a reduction from QSAT_2^Π 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

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

The resulting AF is shown in Fig. 4.

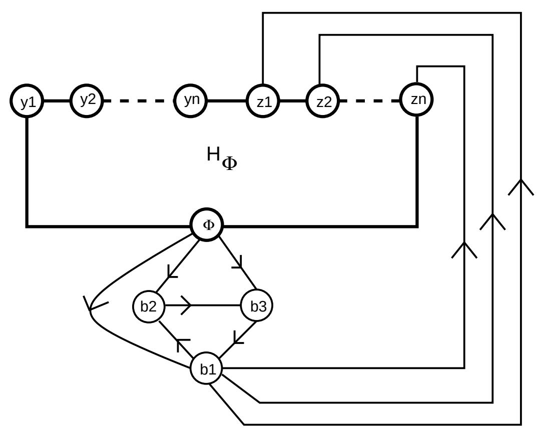

Figure 4: The Argumentation Framework $\mathcal{G}_\Phi$.

Fact 4 *(Dunne and Bench-Capon [14])*

a. $\Phi(Y_n, Z_n)$ *is accepted as an instance of* QSAT_2^Π *if and only if* $\text{SA}(\mathcal{G}_\Phi, \Phi)$.

b. $\Phi(Y_n, Z_n)$ *is accepted as an instance of* QSAT_2^Π *if and only if* $\mathcal{G}_\Phi$ *is coherent.*

Appendix B: The partition $\langle \mathcal{X}_{\text{OUT}}, \mathcal{X}_{\text{PSA}}, \mathcal{X}_{\text{CA}} \rangle$

The proof of Thm. 5 defines a partition of the arguments of an AF for any $\mathcal{H}(\mathcal{X}, \mathcal{A})$, i.e. the sets

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

This partition has the property

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

In the sequel Q will denote one of the three cases $\{\text{OUT}, \text{PSA}, \text{CA}\}$ so that we can consider the decision problems of Table 7:

Problem Name	Instance	Question
$\mathcal{X}^{\text{Q}}$**-set** ($\mathcal{X}^{\text{Q}}_{\{\}}$)	$\mathcal{H}(\mathcal{X}, \mathcal{A})$; $S \subseteq \mathcal{X}$	Is $S \subseteq \mathcal{X}_{\text{Q}}$?
Maximal $\mathcal{X}^{\text{Q}}$**-set** ($\text{M}\mathcal{X}^{\text{Q}}_{\{\}}$)	$\mathcal{H}(\mathcal{X}, \mathcal{A})$; $S \subseteq \mathcal{X}$	Is $S = \mathcal{X}_{\text{Q}}$?

Table 7: Decision problems related to partition of $\mathcal{X}$

For the maximality variants, the search forms are also of some interest.

Theorem 8

a. $\text{M}\mathcal{X}^{\text{PSA}}_{\{\}}$ *is* D^p*–hard.*

b. $\text{M}\mathcal{X}^{\text{CA}}_{\{\}}$ *is* D^p*–hard.*

c. $\text{M}\mathcal{X}^{\text{OUT}}_{\{\}}$ *is* D^p*-complete.*

Proof:

a. To see that $\text{M}\mathcal{X}^{\text{PSA}}_{\{\}}$ is D^p-hard, consider any instance – $\langle \Phi_1, \Phi_2\rangle$ – of SAT-UNSAT. Forming the AF, $\mathcal{F}_{\langle \Phi_1, \Phi_2\rangle}$ consisting of a copy of the framework $\mathcal{F}_{\Phi_1}$ and $\mathcal{F}_{\Phi_2}$ described in the proof of Thm. 1, and choosing $S = \{\Psi_2\}$ it is easily seen that $\langle \Phi_1, \Phi_2\rangle$ is accepted as an instance of SAT-UNSAT if and only if $\langle \mathcal{F}_{\langle \Phi_1, \Phi_2\rangle}, \{\Psi_2\}\rangle$ is accepted as an instance of $\text{M}\mathcal{X}^{\text{PSA}}_{\{\}}$.

b. Given an instance, $\langle\Phi_1,\Phi_2\rangle$ of SAT-UNSAT construct an AF, $\mathcal{F}_{\langle\Phi_1,\Phi_2\rangle}$ consisting of the AF $\mathcal{H}_{\Phi_1}$ and $\mathcal{F}_{\Phi_2}$ but with the latter having the attack $\langle\Psi_2,\Phi_2\rangle$ removed. To complete the instance of $\text{M}\mathcal{X}_{\{\}}^{\text{CA}}$ set S to contain all arguments of $\mathcal{H}_{\Phi_1}$. With these, Φ_1 is satisfiable if and only if $S\subseteq\mathcal{X}_{\text{CA}}$ and Φ_2 is unsatisfiable if and only if $\mathcal{X}_{\text{CA}}\subseteq S$ (we note that the set $\mathcal{X}_{\text{PSA}}=\emptyset$ irrespective of the satisfiability properties of $\langle\Phi_1,\Phi_2\rangle$). In total, $\langle\Phi_1,\Phi_2\rangle$ defines a positive instance of SAT-UNSAT if and only if $\langle\mathcal{F}_{\langle\Phi_1,\Phi_2\rangle},S\rangle$ defines a positive instance of $\text{M}\mathcal{X}_{\{\}}^{\text{CA}}$.

c. That $\text{M}\mathcal{X}_{\{\}}^{\text{OUT}}\in\text{D}^p$ follows by choosing L_1 as

$$\{\langle\mathcal{H}(\mathcal{X},\mathcal{A}),S\rangle\ :\ \forall\,T\subseteq\mathcal{X}\ :\ \text{ADM}(\mathcal{H},T)\Rightarrow T\cap S=\emptyset\}$$

and L_2 as the set of pairs $\langle\mathcal{H}(\mathcal{X},\mathcal{A}),S\rangle$ for which,

$$\exists\ \langle T_1,T_2,\ldots,T_{|\mathcal{X}\setminus S|}\rangle\ :\ \bigwedge_{x_i\in\mathcal{X}\setminus S}\text{ADM}(\mathcal{H},T_i)\wedge(x_i\in T_i)$$

So that L_1 captures all $S\subseteq\mathcal{X}_{\text{OUT}}$ and is in co-NP. Similarly L_2 describes sets, S, for which every member of $\mathcal{X}\setminus S$ is credulously accepted, i.e. subsets of $\mathcal{X}_{\text{PSA}}\cup\mathcal{X}_{\text{CA}}$ with $L_2\in\text{NP}$. We deduce that $\text{M}\mathcal{X}_{\{\}}^{\text{OUT}}=L_1\cap L_2$.

To show that $\text{M}\mathcal{X}_{\{\}}^{\text{OUT}}$ is D^p–hard, we again use a reduction from SAT-UNSAT. Given an instance $\langle\Phi_1,\Phi_2\rangle$ of SAT-UNSAT for an AF $\mathcal{F}_{\langle\Phi_1,\Phi_2\rangle}$ consisting of a copy of the AF, $\mathcal{H}_{\Phi_1}$ and a copy of the AF $\mathcal{F}_{\Phi_2}$ of Thm.1 modified by removing the attack $\langle\Psi,\Phi_2\rangle$. The instance of $\text{M}\mathcal{X}_{\{\}}^{\text{OUT}}$ is completed by setting S to contain every argument of $\mathcal{F}_{\Phi_2}$. Let $\mathcal{X}$ denote the set of all arguments in the resulting system. We claim that $\langle\Phi_1,\Phi_2\rangle$ is accepted as an instance of SAT-UNSAT if and only if $\langle\mathcal{F}_{\langle\Phi_1,\Phi_2\rangle},S\rangle$ is accepted as an instance of $\text{M}\mathcal{X}_{\{\}}^{\text{OUT}}$. Suppose first that Φ_1 is satisfiable and Φ_2 is unsatisfiable. From the fact that Φ_2 is unsatisfiable it is certainly the case that no argument in $\mathcal{F}_{\Phi_2}$ can be credulously accepted, i.e. $S\subseteq\mathcal{X}_{\text{OUT}}$. Similarly, since Φ_1 is satisfiable, its argument Φ_1 is credulously accepted (so that $\Phi_1\not\in\mathcal{X}_{\text{OUT}}$) and it is easy to see that every other argument in $\mathcal{H}_{\Phi_1}$ is credulously accepted. We deduce that $S=\mathcal{X}_{\text{OUT}}$ as required. On the other hand, suppose that $\langle\mathcal{F}_{\langle\Phi_1,\Phi_2\rangle},S\rangle$ is accepted as an instance of $\text{M}\mathcal{X}_{\{\}}^{\text{OUT}}$. From the fact that $\Phi_1\not\in S$ and $S=\mathcal{X}_{\text{OUT}}$, it must be the case the case that Φ_1 is credulously accepted, i.e. Φ_1 is satisfiable. Similarly, no argument in S is credulously accepted and, in particular, the argument Φ_2 of $\mathcal{F}_{\Phi_2}$: this can only be the case, however, when Φ_2 is unsatisfiable. We deduce that $\langle \Phi_1, \Phi_2 \rangle$ is accepted as an instance of SAT-UNSAT if and only if $\langle \mathcal{F}_{\langle \Phi_1, \Phi_2 \rangle}, S \rangle$ is accepted as an instance of $\mathrm{M}\mathcal{X}^{\mathrm{OUT}}_{\{\}}$.

□