Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence

Parametric Properties of Ideal Semantics*

Wolfgang Dvořák
Institute of Information Systems
Vienna Univ. of Tech., Austria
dvorak@dbai.tuwien.ac.at

Paul E. Dunne
Dept. of Computer Science
Univ. of Liverpool, UK
ped@csc.liv.ac.uk

Stefan Woltran
Institute of Information Systems
Vienna Univ. of Tech., Austria
woltran@dbai.tuwien.ac.at

Abstract

The concept of "ideal semantics" has been promoted as an alternative basis for skeptical reasoning within abstract argumentation settings. Informally, ideal acceptance not only requires an argument to be skeptically accepted in the traditional sense but further insists that the argument is in an admissible set all of whose arguments are also skeptically accepted. The original proposal was couched in terms of the so-called preferred semantics for abstract argumentation. We argue, in this paper, that the notion of "ideal acceptability" is applicable to arbitrary semantics and justify this claim by showing that standard properties of classical ideal semantics, e.g. unique status, continue to hold in any "reasonable" extension-based semantics. We categorise the relationship between the divers concepts of "ideal extension wrt semantics σ" that arise and we present a comprehensive analysis of algorithmic and complexity-theoretic issues.

1 Introduction

Argumentation has evolved as an important field in AI with abstract argumentation frameworks (AFs, for short) as introduced by Dung [1995] being its most studied formalism. Meanwhile, a wide range of semantics for AFs has been proposed; for an overview see [Baroni and Giacomin, 2009]. One of the most interesting recent approaches is the ideal semantics which have been introduced as an alternative basis for skeptical argumentation by Dung, Mancarella and Toni [2007]. These define the concept of *ideal sets* as *admissible* sets, S, with the property that S is contained in all *preferred extensions* (i.e. subset-maximal admissible sets) of an AF $\langle \mathcal{X}, \mathcal{A} \rangle$, with the (unique) ideal *extension* of $\langle \mathcal{X}, \mathcal{A} \rangle$ being its *maximal* (again, wrt $\subseteq$) ideal set.

Although Dung *et al.* [2007] present ideal semantics using preferred extensions as the basis, we argue that the approach is appropriately handled as a *parametric model* of acceptance, applicable to multiple status semantics in general, with the specific example of preferred semantics being one instance. In fact, another instance of this schema has been proposed by Caminada [2007], who used semi-stable [Caminada, 2006] instead of preferred extension for the base semantics, and called this approach *eager semantics*.

A natural question of interest thus concerns to what extent general properties of ideal semantics can be established, particularly in the light of known results for the preferred and eager case. Our aim in this paper is to focus on complexity theoretic properties. In particular, we are interested in the question whether the complexity of ideal reasoning can be derived from known complexity results for the respective base semantics in terms of credulous or skeptical acceptance. As a basis we shall use the characterisations proven by Dunne [2009], who showed that the main reasoning problems here are located between NP resp. coNP and Θ_2^P.[1] As we will see these results are due to the fact that the problems of credulous and skeptical acceptance for the base semantics are located on different levels of the polynomial hierarchy. For many other base semantics the complexity of ideal and skeptical reasoning is the same, as we show in this paper.

What we call ideal semantics, is also known as prudent reasoning in the world of non-monotonic reasoning (this should not to be confused with prudent semantics for AFs as proposed in [Coste-Marquis *et al.*, 2005]). Typically, prudent reasoning is considered as an additional reasoning mode besides credulous and skeptical reasoning, and in case of default logic can be defined as follows, see e.g. [Besnard and Schaub, 1998]: a formula φ follows from a default theory (W, D) by prudent reasoning if φ follows from W plus those defaults from D which are generating defaults for all extensions of (W, D). Another similar concept is the one of *free consequences* [Benferhat *et al.*, 1993] which are defined as formulas being entailed by the intersection of all maximal consistent subsets of a (possibly inconsistent) knowledge base.

This motivates our view of understanding ideal acceptance as a third reasoning mode for abstract argumentation, different from credulous and skeptical acceptance and, in principle, applicable to any semantics. We recall that here the difference to skeptical acceptance is due to the fact that the intersection of the extensions given by the chosen base semantics

*This work was supported by the Vienna Science and Technology Fund (WWTF) under grant ICT08-028 and by the Austrian Science Fund (FWF) under grant P20704-N18.

[1] Exact bounds seem to be very hard to establish; Dunne [2009] gives Θ_2^P–hardness proofs with respect to *randomised* reductions.

is not necessarily an admissible set. A well-known example is the AF $(\{a,b,c,d\},\{(a,b),(b,a),(a,c),(b,c),(c,d)\})$ where the preferred extensions are $\{a,d\}$ and $\{b,d\}$, thus d is skeptically accepted, but the only ideal set is $\emptyset$ (since $\{d\}$ is not admissible). For prudent reasoning in default logic, the difference to skeptical reasoning is a bit more subtle. As an example, consider, $\Delta = (\{\alpha \vee \beta \rightarrow \gamma\}, \{\frac{:\neg\alpha}{\beta}, \frac{:\neg\beta}{\alpha}\})$. Here, we also have two extensions, one containing α the other containing β. Hence, γ follows skeptically from Δ but not via prudent reasoning, since none of the defaults is generating in all extensions. Let us finally mention that such a situation is also possible in the argumentation setting as soon as arguments are not fully abstract anymore. As an example, consider the AF from above with additional attack (d,c). Then $\{d\}$ is the ideal extension coinciding with the intersection of skeptically accepted arguments. If one considers now that, for instance, a's claim is $\alpha \wedge \neg\beta$, b's claim is $\beta \wedge \neg\alpha$, and d's claim is $\alpha \vee \beta \rightarrow \gamma$, then γ follows skeptically from the AF, but γ is not derivable from the ideal extension.

Our main contributions can be summarised as follows:

- We show general results about properties ideal semantics for abstract argumentation satisfy. Most notably, we show that whenever the base semantics is based on conflict-free sets (which is true for any reasonable semantics) there exists a *unique* ideal extension. Moreover, it holds that for each base semantics satisfying the reinstatement property [Baroni and Giacomin, 2007], the ideal extension is also complete.

- We provide two alternative algorithms how to compute an ideal extension. One is a generalisation of concepts introduced in [Dunne, 2009] and relies on credulous acceptance for the base semantics. A novel characterisation uses the skeptically accepted arguments instead.

- We provide a thorough complexity analysis. We are able to provide general results for ideal reasoning which can be derived from complexity results for the base semantics. Moreover, we give novel exact bounds for the following realisations of ideal semantics: semi-stable (we note that complexity has not been analysed in [Caminada, 2007]), resolution-based grounded [Baroni *et al.*, 2011a], stage [Verheij, 1996], and naive semantics. For naive semantics, reasoning in the ideal extension remains tractable.

2 Preliminaries

Throughout $\langle \mathcal{X}, \mathcal{A} \rangle$ is a (finite) AF with argument set $\mathcal{X} = \{x_1, \ldots, x_n\}$ and attack relation $\mathcal{A} \subseteq \mathcal{X} \times \mathcal{X}$. For $S \subseteq \mathcal{X}$, let S^- be $\{x \mid \exists y \in S \text{ s.t. } \langle x,y \rangle \in \mathcal{A}\}$ and, similarly, S^+ be $\{x \mid \exists y \in S \text{ s.t. } \langle y,x \rangle \in \mathcal{A}\}$.

A *semantics* for AFs is any predicate $\sigma : 2^{\mathcal{X}} \rightarrow \langle \top, \bot \rangle$ prescribing criteria under which a set of arguments is considered collectively "justified". A semantics is *extension–based* if it satisfies for all $S \subseteq \mathcal{X}$: if $\sigma(S)$ then $\neg\sigma(T)$ for all $T \supset S$. We denote by $\mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle)$ the set of subsets within the AF $\langle \mathcal{X}, \mathcal{A} \rangle$, i.e. $\mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle) = \{S \subseteq \mathcal{X} : \sigma(S)\}$. An (extension-based) semantics, σ, is said to be *unique status* if for every AF $\langle \mathcal{X}, \mathcal{A} \rangle$, $|\mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle)| = 1$ holds; otherwise the semantics is *multiple status*. For σ a unique status semantics, $E_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle)$ denotes the unique set comprising $\mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle)$.

An argument, $x \in \mathcal{X}$ is *acceptable wrt* $S \subseteq \mathcal{X}$ if for any $y \in \mathcal{X}$ for which $\langle y,x \rangle \in \mathcal{A}$ there is some $z \in S$ such that $\langle z,y \rangle \in \mathcal{A}$. The *characteristic function*, $\mathcal{F} : 2^{\mathcal{X}} \rightarrow 2^{\mathcal{X}}$, reports the set of arguments that are acceptable to a given set. It is shown by Dung [1995] that the function $\mathcal{F}$ has a least fixed point which is called the *grounded extension*.

We require the following concepts for an AF $F = \langle \mathcal{X}, \mathcal{A} \rangle$:

$$
\begin{aligned}
\mathcal{E}_{cf}(F) &= \{S \subseteq \mathcal{X} \mid \forall x,y \in S,\ \langle x,y \rangle \notin \mathcal{A}\} \\
\mathcal{E}_{adm}(F) &= \{S \in \mathcal{E}_{cf}(F) \mid S \subseteq \mathcal{F}(S)\} \\
\mathcal{E}_{comp}(F) &= \{S \in \mathcal{E}_{adm}(F) \mid \mathcal{F}(S) \subseteq S\} \\
\mathcal{E}_{gr}(F) &= \mathcal{F}^k(\emptyset), \text{ for } k \text{ such that } \mathcal{F}^k(\emptyset) = \mathcal{F}^{k+1}(\emptyset) \\
\mathcal{E}_{naive}(F) &= \{S \in \mathcal{E}_{cf}(F) \mid S \subset T \Rightarrow T \notin \mathcal{E}_{cf}(F)\} \\
\mathcal{E}_{pr}(F) &= \{S \in \mathcal{E}_{adm}(F) \mid S \subset T \Rightarrow T \notin \mathcal{E}_{adm}(F)\} \\
\mathcal{E}_{sst}(F) &= \{S \in \mathcal{E}_{adm}(F) \mid S \cup S^+ \subset T \cup T^+ \Rightarrow T \notin \mathcal{E}_{adm}(F)\} \\
\mathcal{E}_{stage}(F) &= \{S \in \mathcal{E}_{cf}(F) \mid S \cup S^+ \subset T \cup T^+ \Rightarrow T \notin \mathcal{E}_{cf}(F)\}
\end{aligned}
$$

$\mathcal{E}_{cf}(F)$ are called the conflict-free sets of F, $\mathcal{E}_{adm}(F)$ and $\mathcal{E}_{comp}(F)$ are the admissible, and respectively, complete sets. The remaining cases define extension-based semantics, and with the exception of the grounded semantics gr, these are all multiple status. In fact, we have defined here naive extensions ($\mathcal{E}_{naive}(F)$), preferred extensions ($\mathcal{E}_{pr}(F)$), semi-stable extensions ($\mathcal{E}_{sst}(F)$), and stage extensions ($\mathcal{E}_{stage}(F)$).

Definition 1 *Let σ and θ be semantics. If for all* AF*s F, $\mathcal{E}_\sigma(F) \subseteq \mathcal{E}_\theta(F)$, we call σ a θ-*preserving *semantics.*

All of the above semantics are cf-preserving, but naive and stage semantics are, for instance, not adm-preserving. Moreover, sst is pr-preserving, and $stage$ is $naive$-preserving.

The *resolution-based* semantics [Baroni *et al.*, 2011a] are a parametric approach defined in the following way: given an AF $F = \langle \mathcal{X}, \mathcal{A} \rangle$, let $\mu(F)$ be the set of (unordered) pairs $\{\{x,y\} \mid \{\langle x,y \rangle, \langle y,x \rangle\} \subseteq \mathcal{A}\}$. A *(full) resolution*, β, of F contains exactly one of each of the attacks $\langle x,y \rangle$, $\langle y,x \rangle$ for each $\{x,y\} \in \mu(F)$. We denote the set of all full resolutions as $\gamma(F)$. If σ is an extension-based semantics the *resolution based* σ semantics, $\sigma*$, is given through $\mathcal{E}_{\sigma*}(\langle \mathcal{X}, \mathcal{A} \rangle) = \min \bigcup_{\beta \in \gamma(\langle \mathcal{X}, \mathcal{A} \rangle)} \{\mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \setminus \beta \rangle)\}$ where $\min$ is with respect to $\subseteq$. We note that $gr*$ – the resolution-based grounded semantics – is *multiple status*. As is known, $gr*$ satisfies many desirable properties, thus we will focus on this instantiation here. In fact, the semantics for which we will consider ideal extensions in this paper, are pr, sst, $stage$, $gr*$, and $naive$.

We assume the reader has knowledge about standard complexity concepts, as P, NP, Σ_2^P, Π_2^P, LOGSPACE (L), and oracles. Beyond these we need the following complexity classes: $\Theta_2^P = \mathrm{P}_{\parallel}^{\mathrm{NP}}$, where $\mathrm{P}_{\parallel}^{\mathcal{C}}$ is the class of decision problems that can be solved by a deterministic polynomial time algorithm which is allowed to make $O(n)$ non-adaptive calls to the $\mathcal{C}$-oracle; $\mathrm{FP}_{\parallel}^{\mathcal{C}}$ which is the corresponding class of function problems; D^P, the class of decision problems L that can be characterised as $L_1 \cap L_2$ for decision problems $L_1 \in$ NP and $L_2 \in$ coNP; D_2^P, the class of decision problems L that can be characterised as $L_1 \cap L_2$ for $L_1 \in \Sigma_2^P$ and $L_2 \in \Pi_2^P$.

Typical *computational problems* of interest are:

a. Credulous acceptance – $\mathrm{CA}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle, x)$ $(x \in \mathcal{X})$. Is there *any* $S \in \mathcal{E}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle)$ for which $x \in S$?

σ	CA_σ	SA_σ	VER_σ	NE_σ
$comp$	NP-c	P-c	in L	NP-c
pr	NP-c	Π_2^P-c	coNP-c	NP-c
sst	Σ_2^P-c	Π_2^P-c	coNP-c	NP-c
$stage$	Σ_2^P-c	Π_2^P-c	coNP-c	in L
$gr*$	NP-c	coNP-c	in P	in P
$naive$	in L	in L	in L	in L

Table 1: Complexity landscape; $\mathcal{C}$-c stands for $\mathcal{C}$-complete.

b. Skeptical acceptance – $\mathrm{SA}_\sigma(\langle\mathcal{X},\mathcal{A}\rangle, x)$ ($x \in \mathcal{X}$). Is $x \in S$ for *every* $S \in \mathcal{E}_\sigma(\langle\mathcal{X},\mathcal{A}\rangle)$?

c. Verification – $\mathrm{VER}_\sigma(\langle\mathcal{X},\mathcal{A}\rangle, S)$ ($S \subseteq \mathcal{X}$). Is it the case that $S \in \mathcal{E}_\sigma(\langle\mathcal{X},\mathcal{A}\rangle)$?

d. Non-emptiness – $\mathrm{NE}_\sigma(\langle\mathcal{X},\mathcal{A}\rangle)$. Is there any $S \in \mathcal{E}_\sigma(\langle\mathcal{X},\mathcal{A}\rangle)$ for which $S \neq \emptyset$?

The complexity of these problems is well explored [Baroni *et al.*, 2011a; Dimopoulos and Torres, 1996; Dunne and Bench-Capon, 2002; Dvořák and Woltran, 2010] (see Table 1).

3 Parameterised Ideal Semantics

We give a general definition of ideal semantics, abstracting from the approaches in [Caminada, 2007; Dung *et al.*, 2007].

Definition 2 *Let $\langle\mathcal{X},\mathcal{A}\rangle$ be an* AF *and σ a semantics that promises at least one extension*[2]. *The* ideal sets *wrt* base semantics *σ of $\langle\mathcal{X},\mathcal{A}\rangle$ are those that satisfy the constraints: (I1) $S \in \mathcal{E}_{adm}(\langle\mathcal{X},\mathcal{A}\rangle)$ and (I2) $S \subseteq \bigcap_{T \in \mathcal{E}_\sigma(\langle\mathcal{X},\mathcal{A}\rangle)} T$. We say that S is an ideal* extension *of $\langle\mathcal{X},\mathcal{A}\rangle$ wrt σ, if S is a $\subseteq$-maximal ideal set (of $\langle\mathcal{X},\mathcal{A}\rangle$) wrt σ. $\mathcal{E}_\sigma^{\mathrm{IDL}}$ denotes the collection of ideal* sets *wrt σ and $\mathcal{E}_\sigma^{\mathrm{IE}}$ denotes the set of ideal extensions wrt σ. σ^{IE} denotes the corresponding semantics.*

Proposition 1 *If a semantics σ is cf-preserving, then σ^{IE} is a unique status semantics.*

Proof: It suffices to show that if $S \in \mathcal{E}_\sigma^{\mathrm{IDL}}(\langle\mathcal{X},\mathcal{A}\rangle)$ and $T \in \mathcal{E}_\sigma^{\mathrm{IDL}}(\langle\mathcal{X},\mathcal{A}\rangle)$ then $S \cup T \in \mathcal{E}_\sigma^{\mathrm{IDL}}(\langle\mathcal{X},\mathcal{A}\rangle)$. First note that $S \cup T \in \mathcal{E}_{cf}(\langle\mathcal{X},\mathcal{A}\rangle)$ since, by the definition of $\mathcal{E}_\sigma^{\mathrm{IDL}}$, $S \subseteq \bigcap_{V \in \mathcal{E}_\sigma(\langle\mathcal{X},\mathcal{A}\rangle)} V$ and $T \subseteq \bigcap_{V \in \mathcal{E}_\sigma(\langle\mathcal{X},\mathcal{A}\rangle)} V$; hence $S \cup T \subseteq V$ for every $V \in \mathcal{E}_\sigma(\langle\mathcal{X},\mathcal{A}\rangle)$. As by assumption σ promises at least one extension $V \in \mathcal{E}_\sigma(\langle\mathcal{X},\mathcal{A}\rangle)$ we then have that $S \cup T$ is a subset of a conflict-free set and thus is itself conflict-free. It must, however, further hold that $S \cup T \in \mathcal{E}_{adm}(\langle\mathcal{X},\mathcal{A}\rangle)$: both S and T are in $\mathcal{E}_{adm}(\langle\mathcal{X},\mathcal{A}\rangle)$ (since they are ideal sets wrt σ), hence any $y \in (S \cup T)^-$ either belongs to S^- (and so is counterattacked by some $z \in S$) or is in T^- (and, in the same way, counterattacked by some argument in T). It follows that $S \cup T \in \mathcal{E}_{adm}(\langle\mathcal{X},\mathcal{A}\rangle)$ forming a subset of every set in $\mathcal{E}_\sigma(\langle\mathcal{X},\mathcal{A}\rangle)$, i.e. $S \cup T \in \mathcal{E}_\sigma^{\mathrm{IDL}}(\langle\mathcal{X},\mathcal{A}\rangle)$. □

Let us now discuss a few further aspects of Definition 2. Indeed, one might ask why the notion of base semantics is parameterised, but the required properties of ideal sets and ideal extensions are to some extent fixed in the following sense: (a) each ideal set is admissible; (b) the ideal extensions are defined wrt subset maximality.

[2]Hence we exclude stable semantics as a base semantics here.

Proposition 1 gives an answer to the second issue already. In fact, using this result one can observe that considering admissible sets $S \subseteq \bigcap_{T \in \mathcal{E}_\sigma(\langle\mathcal{X},\mathcal{A}\rangle)} T$ which are maximal with respect to the cardinality $|S|$, the range $S \cup S^+$ or the cardinality of the range $|S \cup S^+|$ would always yield the ideal extension (defined via $\subseteq$-maximality), as well.

Let us thus turn to issue (a). Two directions are possible: (1) weakening the restriction for being an ideal set; (2) strengthening the properties ideal sets have to satisfy. For (1), the only natural such relaxation would be to consider conflict-free (instead of admissible) sets as ideal ones. However, as long as we assume that any reasonable base semantics is cf-preserving, we get that the set of skeptical accepted arguments is already conflict-free. Hence we would end up with the problem of skeptical acceptance. Against (2), one can argue that a restriction to complete sets is the closest reasonable such sharpening. We can give the following result.

Proposition 2 *If σ satisfies the reinstatement property*[3] *then σ^{IE} is comp-preserving.*

Proof: Let $E \in \mathcal{E}_\sigma^{\mathrm{IE}}(\langle\mathcal{X},\mathcal{A}\rangle)$. By definition E is admissible and it remains to show that every $a \in \mathcal{X}$ defended by E belongs to E. Thus consider such an a. As E is part of every σ-extension, each σ-extension S defends a and therefore (as σ satisfies the reinstatement property) $a \in S$. Hence a is skeptically accepted and defended by E; thus $E \cup \{a\}$ is an ideal set. But as E is already a maximal ideal set, $a \in E$. □

Thus, if the base-semantics satisfies the reinstatement property then the ideal extension is already a complete set and thus considering complete sets (instead of admissible sets in Condition I1 of Definition 2) would not change anything. On the other hand, if the base-semantics σ does not satisfy the reinstatement property, the existence of a complete set S which is contained in all σ-extensions is not guaranteed. For example, consider $\sigma = naive$ and AF $F = (\{a,b\},\{(a,b)\})$. Then, $\mathcal{E}_{naive}(F) = \{\{a\},\{b\}\}$, but $\emptyset$ is not complete here.

4 Algorithms

We discuss two types of algorithms for computing the ideal extension wrt to a given base semantics σ. The first is a generalisation of an algorithm from [Dunne, 2009] and relies on credulous acceptance for σ. We will show that such an algorithm can be used for any base semantics σ which is pr- or $naive$-preserving. Thus, this algorithm applies also to $\sigma = sst$. Our second algorithm is closer to the original definition of ideal sets and thus makes use of skeptical acceptance in σ. It is applicable to any cf-preserving base semantics σ. Having these two algorithms at hand clearly is also of practical value. In fact, whenever both algorithms are applicable, then one can now select in view of the computational complexity of credulous acceptance and skeptical acceptance for the base semantics. We first need a few technical results.

Proposition 3 *If σ is pr-preserving, then for each* AF *$F = \langle\mathcal{X},\mathcal{A}\rangle$ the following relations hold:*

[3]The concept of *reinstatement* [Baroni and Giacomin, 2007] separates complete from admissible sets. A semantics σ satisfies reinstatement iff for every AF $\langle\mathcal{X},\mathcal{A}\rangle$ and $E \in \mathcal{E}_\sigma(\mathcal{X},\mathcal{A})$, we have $\forall x \in \mathcal{X}\colon (\forall y \in \mathcal{X}((y,x) \in \mathcal{A} \rightarrow \exists z \in E : (z,y) \in \mathcal{A})) \rightarrow x \in E$.

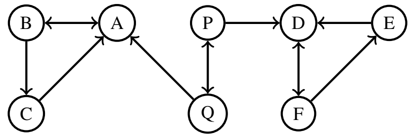

Figure 1: Resolution-based semantics vs. preferred semantics

(C1) $S \in \mathcal{E}_\sigma^{\text{IDL}}(F) \;\Leftrightarrow\; \begin{array}{l} S \in \mathcal{E}_{adm}(F)\ \ \& \\ \forall y \in S^- \ \neg\text{CA}_\sigma(F,y) \end{array}$

(C2) $x \in E_\sigma^{\text{IE}}(F) \Leftrightarrow \forall y \in \{x\}^- \left(\begin{array}{l} \neg\text{CA}_\sigma(F,y)\ \ \& \\ \{y\}^- \cap E_\sigma^{\text{IE}}(F) \neq \emptyset \end{array} \right)$

In the case of ideal sets wrt preferred extensions the characterisation of Proposition 3 has previously been shown in [Dung *et al.*, 2007; Dunne, 2009]. In addition, we now see that it holds also for the case when *semi-stable* semantics are employed as base semantics for ideal reasoning.

However, the above characterisation does not apply to semantics which are not based on admissibility, e.g. stage semantics. Thus our next step is to give a similar characterisation for $naive$-preserving semantics. The only subtle difference is due to the fact that naive semantics do not take the orientation of attacks into account. Thus, we have to add S^+ (resp. $\{x\}^+$) to the conditions from Proposition 3.

Proposition 4 *If σ is $naive$-preserving, then for each* AF $F = \langle \mathcal{X}, \mathcal{A} \rangle$ *the following relations hold:*

(C1') $S \in \mathcal{E}_\sigma^{\text{IDL}}(F) \;\Leftrightarrow\; \begin{array}{l} S \in \mathcal{E}_{adm}(F)\ \ \& \\ \forall y \in S^- \cup S^+ \ \neg\text{CA}_\sigma(F,y) \end{array}$

(C2') $x \in E_\sigma^{\text{IE}}(F) \Leftrightarrow \forall y \in \{x\}^- \cup \{x\}^+ \left(\begin{array}{l} \neg\text{CA}_\sigma(F)\ \ \& \\ \{y\}^- \cap E_\sigma^{\text{IE}}(F) \neq \emptyset \end{array} \right)$

The characterisations (C2) and (C2') from above results suggest how to compute the ideal extension w.r.t. a semantics σ, in case we have given a function that decides credulous acceptance for σ. Algorithm 1 describes this idea.

Algorithm 1 *Input:* AF $\langle \mathcal{X}, \mathcal{A} \rangle$*, function* CA_σ *deciding credulous acceptance.*
1. Determine the sets $\mathcal{X}_{OUT} = \{x \mid \neg\text{CA}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle, x)\}$ *and* $\mathcal{X}_{PSA} = \{x \in \mathcal{X} \setminus \mathcal{X}_{OUT} \mid \{x\}^- \cup \{x\}^+ \subseteq \mathcal{X}_{OUT}\}$.
2. Form the AF $\langle \mathcal{X}_{PSA} \cup \mathcal{X}_{OUT}, \mathcal{B} \rangle$ *in which* $\mathcal{B}$ *contains only attacks* $\langle x, y \rangle \in \mathcal{A}$ *with* $x \in \mathcal{X}_{PSA}$ *and* $y \in \mathcal{X}_{OUT}$ *or* $x \in \mathcal{X}_{OUT}$ *and* $y \in \mathcal{X}_{PSA}$*, i.e. the* AF *is bipartite.*
3. Return $\max\{S \subseteq \mathcal{X}_{PSA} \ : \ S \in \mathcal{E}_{adm}(\langle \mathcal{X}_{PSA}, \mathcal{X}_{OUT}, \mathcal{B} \rangle)\}$.

The correctness of this algorithm (for appropriate base semantics as outlined above) is an immediate consequence of Proposition 3 and Proposition 4. As well there exist simple algorithms for handling bipartite AFs [Dunne, 2007].

Theorem 1 *For any semantics σ which is pr-preserving or $naive$-preserving, Algorithm 1 constructs* E_σ^{IE}.

Although Algorithm 1 is already applicable to a wide range of base semantics, it does *not* work for resolution-based grounded semantics. Consider the example in Figure 1. We have $\mathcal{E}_{gr*} = \{\{P, F\}, \{Q, B\}\}$, thus $E_{GR*}^{\text{IE}} = \emptyset$. However, $\{B, F\}$ is admissible and none of its attackers credulously accepted wrt $gr*$. In fact, $\{B, F\}$ is the standard ideal extension (wrt preferred semantics), since the preferred extensions of the framework are $\{B, F, P\}$ and $\{B, F, Q\}$.

This motivates Algorithm 2 which follows the definition of ideal semantics more closely: it first computes the set of all *skeptically accepted* arguments and then iteratively computes the maximal admissible subset.

Algorithm 2 *Input:* AF $\langle \mathcal{X}, \mathcal{A} \rangle$*, function* SA_σ *deciding skeptical acceptance. 1. Compute* $\mathcal{X}_{\text{SA}} = \{x \mid \text{SA}_\sigma(\langle \mathcal{X}, \mathcal{A} \rangle, x)\}$. *2. Return* $E = \hat{\mathcal{F}}^{|\mathcal{X}_{\text{SA}}|}(\mathcal{X}_{\text{SA}})$*, with* $\hat{\mathcal{F}}(E) = \mathcal{F}(E) \cap E$.

We next show that this algorithm is correct for every reasonable base semantics, i.e. for base semantics that are cf-preserving. The following lemma captures the correctness of the fixed-point iteration in the algorithm.

Lemma 1 *Let* $(\mathcal{X}, \mathcal{A})$ *be an* AF *and* $S \in \mathcal{E}_{cf}(\langle \mathcal{X}, \mathcal{A} \rangle)$*. The* $\subseteq$*-maximal admissible set* $A \subseteq S$ *is* $\hat{\mathcal{F}}^{|S|}(S)$.

Theorem 2 *For any semantics σ which is cf-preserving, Algorithm 2 constructs* E_σ^{IE}.

Proof: Since σ is cf-preserving, the set $\mathcal{X}_{\text{SA}}$ is also conflict-free. Thus by Lemma 1, $E_\sigma^{\text{IE}} = \hat{\mathcal{F}}^{|\mathcal{X}_{\text{SA}}|}(\mathcal{X}_{\text{SA}})$. □

5 Instantiations

Next, we study ideal reasoning wrt concrete base semantics. In particular, we are interested how the ideal extension wrt to different base semantics relate to each other. In the case of complete sets as base semantics, our schema yields that the ideal extension wrt complete sets matches the grounded extension, i.e. for any AF F, $E_{comp}^{\text{IE}}(F) = E_{gr}(F)$.

The other base semantics, we are interested here are $gr*$, pr, sst, $naive$, and $stage$. Since all these semantics σ are cf-preserving, we know that σ^{IE} is a unique status semantics by Proposition 1. Moreover, for $\sigma \in \{pr, sst, gr*\}$, σ has the reinstatement property, and thus by Proposition 2, $E_\sigma^{\text{IE}}(\langle \mathcal{X}, \mathcal{A} \rangle)$ is a complete set for any AF $\langle \mathcal{X}, \mathcal{A} \rangle$. In general, this does not hold for $\sigma \in \{naive, stage\}$. However, the ideal extension wrt naive semantics is easily characterised.

Proposition 5 *For any* AF $F = \langle \mathcal{X}, \mathcal{A} \rangle$, $E_{naive}^{\text{IE}}(F) = \max\{A \mid A \in \mathcal{E}_{adm}(F), A \subseteq \mathcal{X}_{\text{SA}}\}$ *with* $\mathcal{X}_{\text{SA}} = \{x \in \mathcal{X} \mid (x,x) \notin \mathcal{A}, \{x\}^- \cup \{x\}^+ \subseteq \{y \mid (y,y) \in \mathcal{A}\}\}$.

We investigate now how ideal extensions of different base semantics are related to each other. Caminada [2007] has already shown that $E_{pr}^{\text{IE}}(F) \subseteq E_{sst}^{\text{IE}}(F)$ holds and that there exist AFs F, such that $E_{pr}^{\text{IE}}(F) \subset E_{sst}^{\text{IE}}(F)$. With the following results we give a full analysis of the $\subseteq$ relations between the ideal extensions wrt the base semantics considered here.

Theorem 3 *For any* AF F *the following* $\subseteq$*-relations hold:*

$$\begin{array}{ccccccc} E_{comp}^{IE}(F) & \subseteq & E_{gr*}^{IE}(F) & \subseteq & E_{pr}^{IE}(F) & \subseteq & E_{sst}^{IE}(F) \\ & & & & \cup\!\mid & & \\ & & & & E_{naive}^{IE}(F) & \subseteq & E_{stage}^{IE}(F) \end{array}$$

Furthermore, all these $\subseteq$*-relations are proper ones.*

6 Computational Complexity

In this section we provide two kind of complexity results for ideal reasoning. First we present generic complexity bounds for ideal reasoning problems, i.e. complexity bounds which depend on the complexity of reasoning problems for the base semantics. Then we use these results to draw the complexity landscape for the concrete base-semantics mentioned in Section 5. The computational problems we are interested here, are the following:

a. Credulous acceptance – $\text{CA}_\sigma^{\text{IDL}}(\langle \mathcal{X}, \mathcal{A}\rangle, x)$ ($x \in \mathcal{X}$). Is there *any* $S \in \mathcal{E}_\sigma^{\text{IDL}}(\langle \mathcal{X}, \mathcal{A}\rangle)$ for which $x \in S$?

b. Verification for ideal sets – $\text{VER}_\sigma^{\text{IDL}}(\langle \mathcal{X}, \mathcal{A}\rangle, S)$ ($S \subseteq \mathcal{X}$). Is it the case that $S \in \mathcal{E}_\sigma^{\text{IDL}}(\langle \mathcal{X}, \mathcal{A}\rangle)$?

c. Non-emptiness – $\text{NE}_\sigma^{\text{IDL}}(\langle \mathcal{X}, \mathcal{A}\rangle)$. Is there any $S \in \mathcal{E}_\sigma^{\text{IDL}}(\langle \mathcal{X}, \mathcal{A}\rangle)$ for which $S \neq \emptyset$?

b. Verification for ideal extension – $\text{VER}_\sigma^{\text{IE}}(\langle \mathcal{X}, \mathcal{A}\rangle, S)$ ($S \subseteq \mathcal{X}$). Is it the case that $S = E_\sigma^{\text{IE}}(\langle \mathcal{X}, \mathcal{A}\rangle)$?

e. Construction – $\text{CONS}_\sigma^{\text{IE}}(\langle \mathcal{X}, \mathcal{A}\rangle)$. Reports the ideal extension $E_\sigma^{\text{IE}}(\langle \mathcal{X}, \mathcal{A}\rangle$.

We first consider the construction problem using the algorithms given in Section 4.

Theorem 4 *For any cf-preserving semantics σ, E_σ^{IE} can be constructed in $\text{FP}_\parallel^{\mathcal{C}}$, if either $\text{SA}_\sigma \in \mathcal{C}$ or σ is pr-preserving (resp. $naive$-preserving) and $\text{CA}_\sigma \in \mathcal{C}$.*

We next give a strong connection between the problem $\text{CA}_\sigma^{\text{IDL}}$ and the function problem $\text{CONS}_\sigma^{\text{IE}}$, such that we can extend a hardness result for $\text{CA}_\sigma^{\text{IDL}}$ to $\text{CONS}_\sigma^{\text{IE}}$.

Theorem 5 *Let σ be a semantics where unconnected (via attack paths) arguments do not affect each others acceptance status. If $\text{CA}_\sigma^{\text{IDL}}$ is $\mathcal{C}$-hard then $\text{CONS}_\sigma^{\text{IE}}(\langle \mathcal{X}, \mathcal{A}\rangle)$ is $\text{FP}_\parallel^{\mathcal{C}}$-hard (under metric reductions).*

Note that upper bounds for the problem $\text{CONS}_\sigma^{\text{IE}}$ immediately lead to upper bounds for decision problems. The following theorem may provide strictly better upper bounds.

Theorem 6 *Let σ be a cf-preserving semantics and let $\mathcal{V}$ be the complexity of the problem VER_σ. Then, (i) $\text{CA}_\sigma^{\text{IDL}} \in \text{coNP}^{\mathcal{V}}$, (ii) $\text{VER}_\sigma^{\text{IDL}} \in \text{coNP}^{\mathcal{V}}$, (iii) $\text{NE}_\sigma^{\text{IDL}} \in \text{coNP}^{\mathcal{V}}$ and (iv) $\text{VER}_\sigma^{\text{IE}} \in \text{NP}^{\mathcal{V}} \wedge \text{coNP}^{\mathcal{V}}$.*

Proof: Given an AF $F = \langle \mathcal{X}, \mathcal{A}\rangle$, a base semantics σ, an argument $x \in \mathcal{X}$ and a set $S \subseteq \mathcal{X}$. Fact (i) follows from the following $\text{NP}^{\mathcal{V}}$ algorithm for disproving x to be in E_σ^{IE}:

1. Guess extensions $E_1, \ldots, E_n$ (for $n = |\mathcal{X}|$)
2. Verify extensions using the $\mathcal{V}$-oracle
3. Compute $S_{\text{SA}} := \bigcap_{i=1}^n E_i$
4. Compute the $\subseteq$-maximal admissible set $A \subseteq S_{\text{SA}}$.
5. Accept iff $x \notin A$

We get (ii) by a simple adaptation of the above algorithm: instead of testing $x \notin A$ one tests whether $S \not\subseteq A$. For (iii) we use another adaptation of the above algorithm, now one tests for $A = \emptyset$, which proves that $E_\sigma^{\text{IE}} = \emptyset$. (iv): To verify $E_\sigma^{\text{IE}} = S$, one tests whether $S \in \mathcal{E}_\sigma^{\text{IDL}}(F)$ using the $\text{coNP}^{\mathcal{V}}$-algorithm and disproves the ideal acceptance for all $a \in \mathcal{X} \setminus S$ via the above $\text{NP}^{\mathcal{V}}$-algorithm. □

In the following we give generic lower bounds, i.e. generic hardness results, for the problems $\text{VER}_\sigma^{\text{IDL}}$ and $\text{CA}_\sigma^{\text{IE}}$ depending on the complexity of the problem CA_σ.

To this end, we introduce the following property:

Definition 3 *A pair $(\langle \mathcal{X}, \mathcal{A}\rangle, x)$ of an AF $\langle \mathcal{X}, \mathcal{A}\rangle$ and $x \in \mathcal{X}$ satisfies the* mutual attack property *wrt a semantics σ iff $\text{CA}_\sigma(\langle \mathcal{X}, \mathcal{A}\rangle, x) \Rightarrow \text{CA}_\sigma(\langle \mathcal{X} \cup \{y\}, \mathcal{A} \cup \{\langle x, y\rangle, \langle y, x\rangle\}\rangle, x)$ (where $y \notin \mathcal{X}$). A semantics σ has the* mutual attack property *iff each $(\langle \mathcal{X}, \mathcal{A}\rangle, x)$ has the* mutual attack property *wrt σ.*

Semantics adm, $comp$, pr, $gr*$, and $naive$ all have the mutual attack property. But one can construct AFs which do not satisfy the mutual attack property for sst and $stage$.

Proposition 6 *If σ is pr-preserving or naive-preserving, $\text{CA}_\sigma \in \mathcal{C}$ for some complexity class $\mathcal{C}$ which is closed under $\cup$, and there exist $\mathcal{C}$–hard instances of the CA_σ problem that satisfy the mutual attack property then (a) $\text{VER}_\sigma^{\text{IDL}}$ is co$\mathcal{C}$–complete and (b) $\text{CA}_\sigma^{\text{IDL}}$ is co$\mathcal{C}$–hard.*

We finally provide novel exact bounds for computational problems of ideal reasoning wrt to base semantics $stage$, sst, $gr*$, and $naive$, partly exploiting the generic results.

Theorem 7 *Complexity results as depicted in Table 2 hold (the first two lines contain known results).*

Proof: For $\sigma \in \{sst, stage, gr*\}$, membership for the decision problems follows by Theorem 6 and the complexity of VER_σ. Membership for $\text{CONS}_\sigma^{\text{IE}}$ is a consequence of Theorem 4. Hardness for the construction problems will follow by Theorem 5 after having shown hardness for $\text{CA}_\sigma^{\text{IDL}}$.

By Theorem 4, we can construct E_{naive}^{IE} in P; thus all the decision problems for $naive$ semantics are in P. $\text{VER}_{naive}^{\text{IDL}} \in$ L follows by Proposition 6. The P-hardness of $\text{CA}_{naive}^{\text{IDL}}$, $\text{NE}_{naive}^{\text{IDL}}$ and $\text{VER}_{naive}^{\text{IE}}$ can be shown via a reduction from the P-hard problem of deciding whether some atom follows from a propositional definite Horn theory.

The hardness results for $\sigma \in \{sst, stage\}$ use the following reduction from [Dvořák and Woltran, 2010], which maps the Π_2^P–hard problem $QBF_2^\forall$ to SA_σ and co-CA_σ: Given a QBF formula $\Phi = \forall Y \exists Z \bigwedge_{c \in C} c$, construct $K_\Phi = (\mathcal{X}, \mathcal{A})$, where $\mathcal{X} = \{\Phi, \Psi, b\} \cup C \cup Y \cup \bar{Y} \cup Y' \cup \bar{Y}' \cup Z \cup \bar{Z}$ and $\mathcal{A} = \{\langle c, \Phi\rangle \mid c \in C\} \cup \{\langle \Phi, \Psi\rangle, \langle \Psi, \Phi\rangle, \langle \Phi, b\rangle, \langle b, b\rangle\} \cup \{\langle x, \bar{x}\rangle, \langle \bar{x}, x\rangle \mid x \in Y \cup Z\} \cup \{\langle y, y'\rangle, \langle \bar{y}, \bar{y}'\rangle, \langle y', y'\rangle, \langle \bar{y}', \bar{y}'\rangle \mid y \in Y\} \cup \{\langle l, c\rangle \mid$ literal l occurs in $c \in C\}$. In [Dvořák and Woltran, 2010], it is shown that Ψ is credulously accepted iff Φ is not valid. Since the pairs (K_Φ, Ψ) are Σ_2^P–hard instance for CA_σ satisfying the mutual attack property, we obtain by Proposition 6 that both $\text{VER}_\sigma^{\text{IDL}}$ and $\text{CA}_\sigma^{\text{IDL}}$ are Π_2^P–hard.

To show that $\text{NE}_\sigma^{\text{IDL}}$ is Π_2^P-hard, we use a variant of this construction by adding an argument y and attacks $(y, \Psi), (\Psi, y)$ to K_Φ. Using a restricted class of Π_2^P-hard $QBFs$ $\forall Y \exists Z \bigwedge_{c \in C} c$ (in which $\bigwedge_{c \in C} c$ has a model M, with $M \cap Z = \emptyset$ and a model M', with $Z \subseteq M'$), one can show that the only argument that can be ideal accepted is y. Hence $\text{NE}_\sigma^{\text{IDL}}$ is true iff y is ideal accepted iff Φ is valid. Thus $\text{NE}_\sigma^{\text{IDL}}$ is Π_2^P-hard.

The results for $gr*$ can be established by similar techniques (using a different reduction from UNSAT). Since $gr*$ preserves neither pr nor $naive$, a direct use of Proposition 6, is not possible, however. □

σ	$\text{VER}_{\sigma}^{\text{IDL}}$	$\text{CA}_{\sigma}^{\text{IDL}}$	$\text{NE}_{\sigma}^{\text{IDL}}$	$\text{VER}_{\sigma}^{\text{IE}}$	$\text{CONS}_{\sigma}^{\text{IE}}$
comp	P-c	P-c	in L	P-c	in FP
pr	coNP-c	in Θ_2^P	in Θ_2^P	in Θ_2^P	$\text{FP}_{\parallel}^{\text{NP}}$-c
sst	Π_2^P-c	Π_2^P-c	Π_2^P-c	D_2^P-c	$\text{FP}_{\parallel}^{\Sigma_2^P}$-c
stage	Π_2^P-c	Π_2^P-c	Π_2^P-c	D_2^P-c	$\text{FP}_{\parallel}^{\Sigma_2^P}$-c
gr∗	coNP-c	coNP-c	coNP-c	D^P-c	$\text{FP}_{\parallel}^{\text{NP}}$-c
naive	in L	P-c	P-c	P-c	in FP

Table 2: Complexity of ideal reasoning

7 Discussion

While in their original paper, Dung *et al.* [2007] have proposed their approach of ideal sets and extensions as a dedicated semantics for abstract AFs (being in the line with the logic programming semantics [Alferes *et al.*, 1993] which motivated their work), we believe that the idea behind ideal semantics is more in the spirit of an alternative, third, notion of acceptance besides credulous and skeptical acceptance. A view which is backed up by the introduction of eager semantics [Caminada, 2007] which is defined as ideal semantics but replacing preferred by semi-stable extensions as a basis.

In this paper we analysed common properties for different formings of ideal semantics and we investigated the instantiation wrt semi-stable, stage, naive, and resolution-based grounded semantics in detail. Our complexity analysis showed that ideal reasoning in these instantiations is of the same complexity as skeptical reasoning, which is in contrast to the standard ideal semantics, where reasoning is known to be in Θ_2^P [Dunne, 2009] and thus easier than the Π_2^P-hard problem of skeptical reasoning (for preferred semantics); see also Table 2. We identified as a main reason for this result the fact that credulous reasoning in preferred semantics is only in NP, while for the semantics we have considered here, credulous acceptance is on the same level of complexity as skeptical reasoning. We therefore also provided two different generic algorithms for computing the ideal extension, one based on the credulous acceptance problem, the other based on the skeptical acceptance problem of the base semantics.

As a final remark, we thus note that the view promoted in our treatment of ideal reasoning is by no means limited to semantics in Dung's AFs but can also be considered, in principle, with respect to such developments as value-based frameworks [Bench-Capon, 2003], extended AFs [Modgil, 2009], or frameworks with recursive attacks [Baroni *et al.*, 2011b].

References

[Alferes *et al.*, 1993] J. J. Alferes, P. M. Dung, and L. M. Pereira. Scenario semantics of extended logic programs. In *Proc. LPNMR 1993*, pages 334–348, 1993.

[Baroni and Giacomin, 2007] P. Baroni and M. Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artif. Intell.*, 171(10-15):675–700, 2007.

[Baroni and Giacomin, 2009] P. Baroni and M. Giacomin. Semantics of abstract argument systems. In *Argumentation in Artificial Intelligence*, pages 25–44. Springer, 2009.

[Baroni *et al.*, 2011a] P. Baroni, P. E. Dunne, and M. Giacomin. On the resolution-based family of abstract argumentation semantics and its grounded instance. *Artif. Intell.*, 175(3-4):791–813, 2011.

[Baroni *et al.*, 2011b] P. Baroni, F. Cerutti, M. Giacomin, and G. Guida. AFRA: Argumentation framework with recursive attacks. *Int. J. Approx. Reasoning*, 52(1):19–37, 2011.

[Bench-Capon, 2003] T. J. M. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *JLC*, 13(3):429–448, 2003.

[Benferhat *et al.*, 1993] S. Benferhat, D. Dubois, and H. Prade. Argumentative inference in uncertain and inconsistent knowledge bases. In *Proc. UAI 1993*, pages 411–419.

[Besnard and Schaub, 1998] P. Besnard and T. Schaub. Signed systems for paraconsistent reasoning. *J. Autom. Reasoning*, 20(1):191–213, 1998.

[Caminada, 2006] M. Caminada. Semi-stable semantics. In *Proc. COMMA 2006*, pages 121–130.

[Caminada, 2007] M. Caminada. Comparing two unique extension semantics for formal argumentation: ideal and eager. In *Proc. BNAIC 2007*, pages 81–87.

[Coste-Marquis *et al.*, 2005] S. Coste-Marquis, C. Devred, and P. Marquis. Prudent semantics for argumentation frameworks. In *Proc. ICTAI 2005*, pages 568–572.

[Dimopoulos and Torres, 1996] Y. Dimopoulos and A. Torres. Graph theoretical structures in logic programs and default theories. *TCS*, 170(1-2):209–244, 1996.

[Dung *et al.*, 2007] P. M. Dung, P. Mancarella, and F. Toni. Computing ideal sceptical argumentation. *Artif. Intell.*, 171(10-15):642–674, 2007.

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

[Dunne and Bench-Capon, 2002] P. E. Dunne and T. J. M. Bench-Capon. Coherence in finite argument systems. *Artif. Intell.*, 141(1/2):187–203, 2002.

[Dunne, 2007] P. E. Dunne. Computational properties of argument systems satisfying graph-theoretic constraints. *Artif. Intell.*, 171(10-15):701–729, 2007.

[Dunne, 2009] P. E. Dunne. The computational complexity of ideal semantics. *Artif. Intell.*, 173(18):1559–1591, 2009.

[Dvořák and Woltran, 2010] W. Dvořák and S. Woltran. Complexity of semi-stable and stage semantics in argumentation frameworks. *Inf. Process. Lett.*, 110(11):425–430, 2010.

[Modgil, 2009] S. Modgil. Reasoning about preferences in argumentation frameworks. *Artif. Intell.*, 173(9-10):901–934, 2009.

[Verheij, 1996] B. Verheij. Two approaches to dialectical argumentation: admissible sets and argumentation stages. In *Proc. NAIC 1996*, pages 357–368.