ECAI 2010
H. Coelho et al. (Eds.)
IOS Press, 2010

doi:10.3233/978-1-60750-606-5-119

Computation in Extended Argumentation Frameworks

Paul E. Dunne and **Sanjay Modgil** and **Trevor Bench-Capon**[1]

Abstract. Extended Argumentation Frameworks (EAFs) are a recently proposed formalism that develop abstract argumentation frameworks (AFs) by allowing attacks between *arguments* to be attacked themselves: hence EAFs add a relationship $\mathcal{D} \subseteq \mathcal{X} \times \mathcal{A}$ to the arguments ($\mathcal{X}$) and attacks ($\mathcal{A} \subseteq \mathcal{X} \times \mathcal{X}$) in an AF's basic directed graph structure $\langle \mathcal{X}, \mathcal{A} \rangle$. This development provides a natural way to represent and reason about preferences between arguments. Studies of EAFs have thus far focussed on acceptability semantics, proof-theoretic processes, and applications. However, no detailed treatment of their practicality in computational settings has been undertaken. In this paper we address this lacuna, considering algorithmic and complexity properties specific to EAFs. We show that (as for standard AFs) the problem of determining if an argument is acceptable w.r.t. a subset of $\mathcal{X}$ is polynomial time decidable and, thus, determining the grounded extension and verifying admissibility are efficiently solvable. We, further, consider the status of a number of decision questions specific to the EAF formalism in the sense that these have no counterparts within AFs.

1 Introduction

Dung's abstract model of argumentation [11] has become firmly established as a basis for research on computational aspects of argumentation (see, e.g. the survey of Bench-Capon and Dunne [8]). The central element in Dung's approach is an *argumentation framework* (AF) comprising a pair $\langle \mathcal{X}, \mathcal{A} \rangle$ wherein $\mathcal{X}$ represents a set of abstract atomic *arguments* and $\mathcal{A} \subseteq \mathcal{X} \times \mathcal{X}$ defines the so-called *attack relation*. This relation describes a view of two arguments being "incompatible" in the sense that if $\langle x, y \rangle \in \mathcal{A}$ then the argument x *attacks* the argument y. A major emphasis of subsequent study has been in defining divers formalisms describing intuitive ideas of "collections of acceptable arguments". Typically such proposals are formulated in terms of some predicate $\sigma : 2^{\mathcal{X}} \to \langle \top, \bot \rangle$ so that acceptable collections are those $S \subseteq \mathcal{X}$ for which $\sigma(S)$ holds. In addition to the canonical forms presented in [11] – grounded, admissible, preferred and stable sets – a number of alternatives have been put forward: a detailed review of such argumentation semantics may be found in the survey of Baroni and Giacomin [5].

Recently, Modgil [16] has proposed developing the structure $\langle \mathcal{X}, \mathcal{A} \rangle$ of AFs by incorporating arguments that eliminate attacks. The resulting *Extended Argumentation Frameworks* (EAFs) are defined as triples $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$ whereby $\mathcal{D} \subseteq \mathcal{X} \times \mathcal{A}$. Modgil [16] builds on the standard Dung acceptability semantics to provide analogous concepts in EAFs and demonstrates that EAFs admit a unifying semantic treatment of earlier proposals – notably Preference-based argumentation [1] and Value-based frameworks (VAFs) from [6] – in which mechanisms for disregarding attacks in $\mathcal{A}$ on the basis of the relative strengths of arguments had been proposed: the former through explicit specification of preferences for an argument over an attacker, the latter through the concepts of the value promoted by an argument and ordering of such values w.r.t. an *audience*. Furthermore, by incorporating *arguments* that eliminate attacks, EAFs additionally accommodate argumentation based reasoning about (possibly contradictory) preferences, values and audiences, which are external to the frameworks of [1] and [6]. EAFs thus provide for principled integration of meta and object level argumentation based reasoning, as has been shown in application domains such as agent reasoning [14], normative reasoning [18], and case law [9].

Subsequently, in [15], Modgil presented EAF argument game proof procedures and argument labelling schemes for selected EAF semantics. The principal concerns of [16, 15] have thus been to formalise acceptability concepts in EAFs; questions of algorithmic and complexity properties are not considered. However, such questions have been treated in depth – see the recent survey of Dunne and Wooldridge [13] for an overview – not only within Dung's original frameworks but also within developments such as VAFs, the resolution-based model of Baroni and Giacomin from [4] in [3], and the weighted argument system model of [12]. While it is trivially the case that complexity lower bounds established for the various semantics in AFs continue to be lower bounds in EAFs (by the simple expedient of fixing $\mathcal{D} = \emptyset$), the extent to which *upper* bounds are preserved is far from clear. In total, the distinctive features and range of application of EAFs, and the volume of established studies on algorithms and complexity in related AF contexts, strongly motivate examining similar questions within EAFs. The contribution of this paper is to provide a preliminary treatment of these issues.

In Section 2 we reprise the basic concepts from [11], and the components and analogous structures in EAFs as given in [16]. A significant distinction between AFs and EAFs is found in how the core concept "x is *acceptable* w.r.t. a set S" is treated: for reasons discussed at length within [16], this has a rather more intricate character in EAFs than its AF counterpart. As such, while the definition of "acceptability" (in AFs) leads in a straightforward manner to polynomial time decision procedures, it is less clear whether similar efficient methods follow for acceptability in EAFs. In Section 3 we show that it is, indeed, the case that acceptability in EAFs is decidable in polynomial time. Immediate consequences of this result are that EAF analogues of deciding if a set is admissible as well as construction of an EAF's grounded extension are polynomial time computable, so mirroring the status of these in AFs. In Section 4 we consider a number of decision problems that arise specifically in EAFs in the sense that these reflect behavioural characteristics not featuring in standard AFs. For these problems we establish that deciding whether given EAFs exhibit the characteristics in question is, typically, NP–hard. We conclude Section 4 by examining a problem whose definition while specific to EAFs, has natural analogues not only within classical AFs but

[1] Department of Computer Science, The University of Liverpool, U.K. email: {ped,tbc}@csc.liv.ac.uk, sanjaymodgil@yahoo.co.uk

also within variants such as value-based frameworks: specifically the problem of constructing a (maximal) conflict free set of arguments containing a given subset of arguments. In Section 5 we briefly discuss our results in the context of the EAF+ model proposed by Baroni *et al.* [2]. We conclude and discuss further work in Section 6.

2 Preliminary Background and Notation

The following concepts were introduced in Dung [11]:

Definition 1 *An* argumentation framework *(*AF*) is a pair $\langle \mathcal{X}, \mathcal{A} \rangle$, in which $\mathcal{X}$ is a finite set of* arguments *and $\mathcal{A} \subseteq \mathcal{X} \times \mathcal{X}$ is the* attack relationship. *A pair $\langle x, y \rangle \in \mathcal{A}$ is referred to as 'y* is attacked by *x' or 'x* attacks *y'.*

$x \in \mathcal{X}$ *is* acceptable with respect to $S \subseteq \mathcal{X}$ *if for every $y \in \mathcal{X}$ that attacks x there is some $z \in S$ that attacks y. The* characteristic function*, $\mathcal{F} : 2^{\mathcal{X}} \rightarrow 2^{\mathcal{X}}$ is the mapping which, given $S \subseteq \mathcal{X}$, reports the set of $y \in \mathcal{X}$ for which y is acceptable w.r.t. S.*

$S \subseteq \mathcal{X}$ *is* conflict free *if no argument in S is attacked by any other argument in S. A conflict free set S is:*

- *an* admissible *set if every $y \in S$ is acceptable w.r.t. S;*
- *a* preferred extension *if it is a maximal (with respect to $\subseteq$) admissible set;*
- *a* stable extension *if every $y \notin S$ is attacked by some $x \in S$;*
- *a* complete *extension if $x \in S$ if and only if x is acceptable w.r.t. S (i.e., a fixed point of $\mathcal{F}$);*
- *a* grounded *extension if it is a minimal (w.r.t. $\subseteq$) complete extension. It is shown in [11] that every $\langle \mathcal{X}, \mathcal{A} \rangle$ has a unique grounded extension which is obtained as the least fixed-point of $\mathcal{F}$.*

The preference based approach of [1] and value-based frameworks of [6] allow the notion of "conflict free" set to be modified by describing conditions under which attacks in $\mathcal{A}$ are discounted. Modgil [16]'s *Extended Argumentation Frameworks* generalise and extend (in that they accommodate argumentation *about* preferences, values, audiences etc.) such approaches:

Definition 2 *An* Extended Argumentation Framework *(*EAF*) is a triple $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$ where $\langle \mathcal{X}, \mathcal{A} \rangle$ is an* AF *and $\mathcal{D} \subseteq \mathcal{X} \times \mathcal{A}$ describes the property of "an argument x attacking an* attack *$\langle y, z \rangle$". The following condition is imposed: if $\{\langle x, \langle y, z \rangle \rangle,\ \langle x', \langle z, y \rangle \rangle\} \subseteq \mathcal{D}$ then $\{\langle x, x' \rangle, \langle x', x \rangle\} \subseteq \mathcal{A}$.*

Given $S \subseteq \mathcal{X}$, an attack $\langle x, y \rangle \in \mathcal{A}$ succeeds w.r.t. *S (written $x \rightarrow^S y$) if there is no $z \in S$ for which $\langle z, \langle x, y \rangle \rangle \in \mathcal{D}$.*

*$S \subseteq \mathcal{X}$ is (*EAF*)* conflict free *if for all x, $y \in S$, if $\langle x, y \rangle \in \mathcal{A}$ then $\langle y, x \rangle \notin \mathcal{A}$ and there is some $z \in S$ for which $\langle z, \langle x, y \rangle \rangle \in \mathcal{D}$. Note that mutually attacking pairs of arguments are disqualified from* both *being members of the same conflict free set.*

Thus in EAFs, while $S \subseteq \mathcal{X}$ may fail to be conflict free (in the literal sense of Defn. 1), it becomes so after eliminating attacks between members of S. For example, if $S = \{x, y\}$, $\langle x, y \rangle \in \mathcal{A}$, then S is not (EAF) conflict free, but $S' = \{x, y, z\}$ is (EAF) conflict free if $\langle z, \langle x, y \rangle \rangle \in \mathcal{D}$.

The analogue of "x is acceptable w.r.t. S" in EAFs requires the idea of a *reinstatement set*.

Definition 3 *Given $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$, let $S \subseteq \mathcal{X}$ and $v \rightarrow^S w$. Then $R_S = \{y_1 \rightarrow^S z_1, \ldots, y_n \rightarrow^S z_n\}$ defines a* reinstatement set *for $v \rightarrow^S w$ if R_S satisfies all of the following:*

R1. $v \rightarrow^S w \in R_S$.
R2. $y_i \in S$ for each $1 \leq i \leq n$.
R3. For every $y \rightarrow^S z \in R_S$ and every $\langle z', \langle y, z \rangle \rangle \in \mathcal{D}$ there is some $y' \rightarrow^S z' \in R_S$.

*The argument $x \in \mathcal{X}$ is said to be (*EAF*)* acceptable w.r.t. *$S \subseteq \mathcal{X}$ if whenever $z \rightarrow^S x$ there exists $y \rightarrow^S z$* and *a reinstatement set R_S for $y \rightarrow^S z$.*

Via the formalism of reinstatement sets we obtain direct analogues of admissible, complete, preferred and stable extensions:

Definition 4 *Given $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$, and the* EAF *conflict free set $S \subseteq \mathcal{X}$: S is (*EAF*) admissible if every $x \in S$ is* EAF *acceptable w.r.t. S; S is an* EAF preferred *extension if it is a maximal admissible set; S is an* EAF *stable extension if for every $y \notin S$ there is some $x \in S$ such that $x \rightarrow^S y$; S is an* EAF complete *extension if $x \in S$ if and only if x is* EAF *acceptable w.r.t. S.*

For $S \subseteq \mathcal{X}$ which is EAF conflict free*, the* EAF *characteristic function $\mathcal{F}_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle}(S)$ reports the set of arguments which are* EAF *acceptable w.r.t. S.*

In contrast to Defn. 1, the domain of the EAF characteristic function is limited to *conflict free* sets and furthermore (even within finitary frameworks) $\mathcal{F}_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle}$ is *not* guaranteed to have a least fixed point. However it is the case ([16, Propn. 6]) that iterating $\mathcal{F}_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle}$ starting from the empty set provides a fixed point, i.e. defining $\mathcal{F}^0_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle} = \emptyset$, $\mathcal{F}^{i+1}_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle} = \mathcal{F}_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle}(\mathcal{F}^i_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle})$, for any finite EAF not only is $\mathcal{F}^i_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle}$ EAF conflict free but also there is a finite value k for which $\mathcal{F}^{k+1}_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle} = \mathcal{F}^k_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle}$. These properties lead to the EAF *grounded extension* being defined as

$$GE(\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle) \;=\; \bigcup_{k=0}^{\infty} \mathcal{F}^k_{\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle}$$

The formulation of EAF conflict free, EAF acceptable w.r.t. a set S, and that of $GE(\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle)$, occasions some significant differences between AFs and EAFs. In the former setting we have:

Fact 1

a. Let $\langle \mathcal{X}, \mathcal{A} \rangle$ be an AF*, $S \subseteq \mathcal{X}$ be an admissible set and $x \in \mathcal{X}$ be acceptable w.r.t. S. The set $S \cup \{x\}$ is an admissible set of $\langle \mathcal{X}, \mathcal{A} \rangle$.*
b. Given $\langle \mathcal{X}, \mathcal{A} \rangle$ let $GE(\langle \mathcal{X}, \mathcal{A} \rangle)$ be its grounded extension and $\mathcal{E}_{pr}(\langle \mathcal{X}, \mathcal{A} \rangle)$ the set of preferred extensions.

$$GE(\langle \mathcal{X}, \mathcal{A} \rangle) \;\subseteq\; \bigcap_{S \in \mathcal{E}_{pr}(\langle \mathcal{X}, \mathcal{A} \rangle)} S$$

c. Let $\langle \mathcal{X}, \mathcal{A} \rangle$ be an AF*, S, T subsets of $\mathcal{X}$ with $S \subseteq T$ and $x \in \mathcal{X}$. If x is acceptable w.r.t. S then x is acceptable w.r.t. T.*
d. If $S \subseteq \mathcal{X}$ is a stable extension of $\langle \mathcal{X}, \mathcal{A} \rangle$ then S is also a preferred extension of $\langle \mathcal{X}, \mathcal{A} \rangle$.
e. If $S \subseteq \mathcal{X}$ is not conflict free then every superset of S also fails to be conflict free.

In EAFs *only* Fact 1(a) holds in general: that is, given an EAF $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$, S an EAF admissible set, and arguments x, y both of which are EAF acceptable w.r.t. S, the sets $S_x = S \cup \{x\}$ and $S_y = S \cup \{y\}$ are both EAF admissible; furthermore, x is EAF acceptable w.r.t. S_y *and* y is EAF acceptable w.r.t. S_x. Regarding Fact 1(b), [16] offers examples of EAF preferred extensions that do not have the EAF grounded extension as a subset. Similarly, in contrast to Fact 1(c), cases are illustrated with $S \subset T \subseteq \mathcal{X}$ and $x \in \mathcal{X}$

for which x is EAF acceptable w.r.t. S but is *not* EAF acceptable w.r.t. T. Also, if S is EAF stable and $x \rightarrow^S y$ ($x \in S$, $y \notin S$), then if no other attacks on y are present, $S \cup \{y\}$ will be EAF admissible if $\langle y, \langle x, y\rangle\rangle \in \mathcal{D}$, so that S is not *maximal* (and so preferred). Hence, Fact 1(d) does not hold within EAFs [2]. Finally, as discussed earlier, Fact 1(e) does not hold within EAFs.

The formal definition of EAF acceptability motivates consideration of the following decision problem for EAFs:

Acceptability (ACC)
Instance: EAF $\langle \mathcal{X}, \mathcal{A}, \mathcal{D}\rangle$; $S \subseteq \mathcal{X}$, $x \in \mathcal{X}$.
Question: Is x EAF acceptable w.r.t. S?

In addition, the analogues of Fact 1 (b)–(e) failing to hold in general, motivate consideration of the following decision problems specific to EAFs:

Grounded-Scepticism (GS)
Instance: EAF $\langle \mathcal{X}, \mathcal{A}, \mathcal{D}\rangle$.
Question: Is $GE(\langle \mathcal{X}, \mathcal{A}, \mathcal{D}\rangle) \subseteq S$ for all EAF preferred extensions S of $\langle \mathcal{X}, \mathcal{A}, \mathcal{D}\rangle$?

Monotonicity Failure (MF)
Instance: EAF $\langle \mathcal{X}, \mathcal{A}, \mathcal{D}\rangle$.
Question: Are there sets S and T with $S \subset T \subseteq \mathcal{X}$ and an argument $x \in \mathcal{X}$ for which x is EAF acceptable w.r.t. S but x is *not* EAF acceptable w.r.t. T?

Semi-Coherence[3] (SC)
Instance: EAF $\langle \mathcal{X}, \mathcal{A}, \mathcal{D}\rangle$.
Question: Is it the case that every EAF stable extension of $\langle \mathcal{X}, \mathcal{A}, \mathcal{D}\rangle$ is also an EAF preferred extension of $\langle \mathcal{X}, \mathcal{A}, \mathcal{D}\rangle$?

Conflict free extension (CFE)
Instance: EAF $\langle \mathcal{X}, \mathcal{A}, \mathcal{D}\rangle$; $S \subseteq \mathcal{X}$
Question: Is there a subset T of $\mathcal{X}$ for which $S \subseteq T$ *and* T is EAF conflict free?

We consider these questions in the next sections of this paper.

3 Acceptability in EAFs

Our main result in this section establishes that, just as acceptability w.r.t. a set is polynomial time decidable in AFs, so too EAF acceptability w.r.t. a set is decidable efficiently. Consider Algorithm 1 which takes a subset S and argument x of an EAF $\langle \mathcal{X}, \mathcal{A}, \mathcal{D}\rangle$ as its parameters. Before dealing with its correctness, some informal discussion of the rationale underlying the algorithm may be helpful. The argument x will fail to be EAF acceptable w.r.t. S unless it is possible to counterattack (using arguments within S) each attack $\langle y, x\rangle$. Having removed those attacks which do not succeed w.r.t. S (in l. 2) such attacks are singled out (by colouring these **BLUE** in l. 3). In order to confirm that x is EAF acceptable w.r.t. S we need to identify (for each $y \rightarrow^S x$) *some* $z \in S$ for which $z \rightarrow^S y$ *and* a reinstatement set for $z \rightarrow^S y$ can be found. If such a collection exists then it will be a subset of the attacks coloured **RED** (notice that every $u \rightarrow^S v$ coloured **RED** must have $u \in S$ from l. 6). The final loop (ll. 12–20) eliminates from $\mathcal{D}$ and $\mathcal{A}$ attacks which are "irrelevant" to determining the final status of x w.r.t. S.

We illustrate this in Fig. 1

Algorithm 1 Deciding EAF Acceptability

1: **Input:** $\langle \mathcal{X}, \mathcal{A}, \mathcal{D}\rangle$, $S \subseteq \mathcal{X}$; $x \in \mathcal{X}$;
2: $\mathcal{A} := \mathcal{A} \setminus \{\langle y, z\rangle : \neg(y \rightarrow^S z)\}$;
3: Colour each *attack* $y \rightarrow^S x$ **BLUE**;
4: **repeat**
5:   **for** $y \in \mathcal{X}$ s.t. $y \rightarrow^S x$ or $\langle y, \langle u, v\rangle\rangle$ is **BLUE do**
6:     Colour $z \rightarrow^S y$ **RED** for each $z \in S$ s.t. $z \rightarrow^S y$;
7:   **end for**
8:   **for** $z \rightarrow^S y$ coloured **RED do**
9:     Colour each attack $\langle v, \langle z, y\rangle\rangle \in \mathcal{D}$ **BLUE**
10:   **end for**
11: **until** No change in attack colours
12: **repeat**
13:   **if** $\exists y \in \mathcal{X}$ s.t. $\langle y, \langle v, w\rangle\rangle$ is **BLUE** *and* there is no $u \rightarrow^S y$ coloured **RED then**
14:     $\mathcal{A} := \mathcal{A} \setminus \{\langle v, w\rangle\}$;
15:     $\mathcal{D} := \mathcal{D} \setminus \{\langle y, \langle v, w\rangle\rangle\}$;
16:   **end if**
17:   **if** $\exists z \in S$ s.t. ($\langle z, y\rangle$ is **RED** with $\langle y, \langle u, v\rangle\rangle$ **BLUE**) *and* there is no $\langle p, \langle z, y\rangle\rangle \in \mathcal{D}$ **then**
18:     $\mathcal{D} := \mathcal{D} \setminus \{\langle y, \langle u, v\rangle\rangle\}$;
19:   **end if**
20: **until** No change in $\mathcal{D}$
21: Report whether x is acceptable w.r.t. S in the AF $\langle \mathcal{X}, \mathcal{A}\rangle$

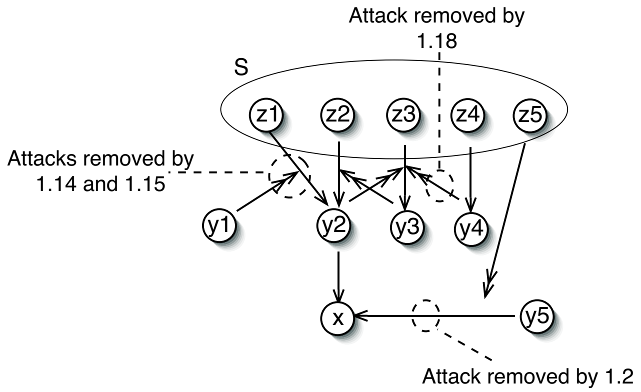

Figure 1. x is EAF-acceptable w.r.t. $\{z1, \ldots, z5\}$ via Alg. 1. Note that x would *not* be EAF-acceptable w.r.t. $\{z1\}$ via Alg. 1.

Theorem 1 *Given $\langle \mathcal{X}, \mathcal{A}, \mathcal{D}\rangle$, $S \subseteq \mathcal{X}$ and $x \in \mathcal{X}$*

a. Alg. 1 reports true if and only if x is EAF *acceptable w.r.t. S.*
b. Alg. 1 reports in time polynomial in $|\mathcal{X}|$.

Proof: (Outline) For part (a), first suppose that x is EAF acceptable w.r.t. S and consider any y_1 for which $y_1 \rightarrow^S x$. Then there is some $z_1 \in S$ with $z_1 \rightarrow^S y_1$ and $R_S = \{z_1 \rightarrow^S y_1, \ldots, z_k \rightarrow^S y_k\}$ defining a reinstatement set for $z_1 \rightarrow^S y_1$. It is easy to see that each of these is coloured **RED** in l. 6 of Alg. 1. Thus, letting $\langle \mathcal{X}, \mathcal{B}\rangle$ be the AF against which acceptability of x w.r.t. S is tested in l. 21, it suffices to argue that $\langle z_1, y_1\rangle \in \mathcal{B}$. Suppose this were not so. Then $\langle z_1, y_1\rangle$ must have been removed at l. 14: this, however contradicts the form of R_S as a reinstatement set and, hence, we deduce that if x is EAF acceptable w.r.t. S then Alg. 1 reports true.

On the other hand, suppose that the algorithm returns true at l. 21. Consider the AF, $\langle \mathcal{X}, \mathcal{B}\rangle$ tested in l. 21 and any $y_1 \rightarrow^S x \in \mathcal{A}$. From the assumption that Alg. 1 reported true we find $z_1 \in S$ and

[2] Fact 1(d) is said to hold in [16]. Thanks to Dung (personal communication) for pointing out this is not the case.
[3] The term is by analogy with the property of coherence as defined in [11], whereby an AF is said to be *coherent* whenever every preferred extension is also stable.

$\langle z_1, y_1 \rangle \in \mathcal{B}$ so that $z_1 \rightarrow^S y_1$. It follows that $\langle z_1, y_1 \rangle$ cannot have been removed from $\mathcal{A}$ at l. 14 so that either no attack $\langle y_2, \langle z_1, y_1 \rangle \rangle$ is in $\mathcal{D}$ (i.e. $\{z_1 \rightarrow^S y_1\}$ is a reinstatement set for itself) or there is some $z_2 \in S$ with $z_2 \rightarrow^S y_2 \in \mathcal{A}$ that also survives as an attack in $\mathcal{B}$. Continuing thus we identify a reinstatement set for some counter-attack $z \rightarrow^S y$ on y with $z \in S$ for every $y \rightarrow^S x$. It follows that x is EAF acceptable w.r.t. S.

Part (b) follows easily by observing that for each $\langle p, q \rangle \in \mathcal{A}$ and $\langle r, \langle p, q \rangle \rangle \in \mathcal{D}$ only a constant number of operations are ever performed. Since $|\mathcal{D}| + |\mathcal{A}|$ is polynomially bounded in $|\mathcal{X}|$ the upper bound follows. ▫

Corollary 1 *Given an* EAF $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$

a. The EAF *grounded extension of $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$ can be constructed in polynomial time.*
b. Determining if $S \subseteq \mathcal{X}$ is EAF *admissible can be decided in polynomial time.*
c. Deciding if $x \in \mathcal{X}$ is credulously accepted w.r.t. EAF *admissibility is* NP*–complete.*
d. Deciding if $S \subseteq \mathcal{X}$ is an EAF *preferred extension is co*NP*–complete.*
e. Deciding if $x \in \mathcal{X}$ is sceptically accepted w.r.t. EAF *preferred extensions is Π_2^p–complete.*

Proof: Parts (a) and (b) follow immediately from Thm. 1 with the definitions of EAF admissibility and EAF grounded extension. The lower bounds for (c)–(e) follow from analogous lower bounds in AFs: simply choose $\mathcal{D} = \emptyset$. As regards upper bounds: for (c) it suffices to guess a subset S of $\mathcal{X}$ and verify both $x \in S$ and S is EAF admissible; for (d), the complementary problem is in NP simply by guessing $T \subseteq \mathcal{X}$ with $S \subset T$ and verifying that either T is EAF admissible or S is not so. Finally (e) follows by testing that for all $S \subseteq \mathcal{X}$ that if $x \notin S$ then S is not an EAF preferred extension. ▫

Corollary 1 establishes that there is no computational overhead in moving from standard AF semantics to EAFs as regards admissibility or grounded semantics.

4 Decision Problems Specific to EAFs

We have seen that the extended concepts of acceptability and admissibility in EAFs raise no additional computational overheads in comparison with AF concepts. However, it turns out that the EAF specific properties described in the decision problems GS and MF raise rather more challenging computational issues.

Theorem 2 GS *is co*NP*–hard.*

Proof: (Outline) We use a reduction from CNF unsatisfiability (UNSAT). Let $\Psi(Z_n) = C_1 \wedge \ldots C_m$ be an instance of UNSAT where, without loss of generality it may be assumed that each clause, C_j is defined using at most three literals over $Z_n = \{z_1, \ldots, z_n\}$. We construct an EAF, $\langle \mathcal{X}_\Psi, \mathcal{A}_\Psi, \mathcal{D}_\Psi \rangle$ as follows given $\Psi(Z_n)$. The set $\mathcal{X}_\Psi$ contains $5n + m + 2$ arguments

$$\{z_i, \neg z_i, +_i, \neg_i, s_i \,:\, 1 \leq i \leq n\} \cup \{C_j \,:\, 1 \leq j \leq m\} \cup \{\Psi, \Phi\}$$

The set of attacks, $\mathcal{A}_\Psi$ contains

$$\begin{array}{l}\{\langle z_i, \neg z_i \rangle,\ \langle \neg z_i, z_i \rangle \,:\, 1 \leq i \leq n\} \cup \\ \{\langle +_i, z_i \rangle,\ \langle \neg_i, \neg z_i \rangle \,:\, 1 \leq i \leq n\} \cup \\ \{\langle \Phi, z_i \rangle,\ \langle \Phi, \neg z_i \rangle\ 1 \leq i \leq n\} \cup \\ \{\langle s_i, +_i \rangle, \langle s_i, \neg_i \rangle \,:\, 1 \leq i \leq n\} \cup \\ \{\langle s_i, \Psi \rangle \,:\, 1 \leq i \leq n\} \cup \\ \{\langle y_k, C_j \rangle \,:\, \text{when } y_k \in \{z_k, \neg z_k\} \text{ occurs in } C_j\} \cup \\ \{\langle C_j, \Psi \rangle \,:\, 1 \leq j \leq m\} \cup \{\langle \Psi, \Phi \rangle\}\end{array}$$

Finally $\mathcal{D}_\Psi$ contains

$$\begin{array}{l}\{\langle +_i, \langle s_i, +_i \rangle \rangle,\ \langle \neg_i, \langle s_i, \neg_i \rangle \rangle\ 1 \leq i \leq n\} \cup \\ \{\langle z_i, \langle +_i, z_i \rangle \rangle,\ \langle \neg z_i, \langle \neg_i, \neg z_i \rangle \rangle \,:\, 1 \leq i \leq n\} \cup \\ \{\langle \Psi, \langle s_i, \Psi \rangle \rangle \,:\, 1 \leq i \leq n\}\end{array}$$

The construction is illustrated in Fig. 2 (note that, in order to improve clarity, some attacks are not shown).

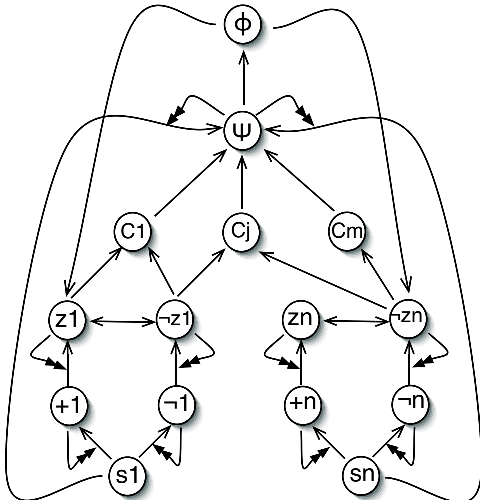

Figure 2. The EAF used in reduction from UNSAT

We claim that $\langle \mathcal{X}_\Psi, \mathcal{A}_\Psi, \mathcal{D}_\Psi \rangle$ is accepted as an instance of GS if and only if $\Psi(Z_n)$ is unsatisfiable.

First observe that $GE(\langle \mathcal{X}_\Psi, \mathcal{A}_\Psi, \mathcal{D}_\Psi \rangle)$ contains exactly $\{s_1, \ldots, s_n\} \cup \{\Phi\}$: the s_i arguments are unattacked, hence form $\mathcal{F}_{\langle \mathcal{X}_\Psi, \mathcal{A}_\Psi, \mathcal{D}_\Psi \rangle}(\emptyset)$; the only argument that is EAF acceptable w.r.t. $\{s_1, \ldots, s_n\}$ is Φ the attack $\langle \Psi, \Phi \rangle$ being countered by $\langle s_1, \Psi \rangle$ (noting that $s_1 \rightarrow^{\{s_1, \ldots, s_n\}} \Psi$ is its own reinstatement set). Now noting that no EAF admissible set can contain *both* Ψ and Φ (and every preferred extension exactly one of these), in order to complete the proof it suffices to show that there is an EAF admissible set containing Ψ if and only if $\Psi(Z_n)$ is satisfiable. Thus, suppose $\langle \alpha_1, \ldots, \alpha_n \rangle \in \langle \top, \bot \rangle^n$ is a satisfying assignment. Then the set

$$S_\alpha = \{\, z_i \,:\, \alpha_i = \top\} \cup \{\neg z_i \,:\, \alpha_i = \bot\} \cup \{\Psi\} \cup \{s_1, \ldots, s_n\}$$

is both EAF conflict free (since $\langle \Psi, \langle s_i, \Psi \rangle \rangle \in \mathcal{D}_\Psi$ for each $1 \leq i \leq n$) and admissible: each C_j must be succesfully attacked by some $z \in S_\alpha$ (α satisfies $\Psi(Z_n)$) and these are the only (remaining) attackers of Ψ. Furthermore the only attacks on $z \in S_\alpha$ are from $\neg z$ (defended by z itself) and Φ (countered by $\Psi \rightarrow^{S_\alpha} \Phi$). Notice that if $z_i \in S_\alpha$ then it is *not* the case that $+_i \rightarrow^{S_\alpha} z_i$. We deduce that should $\Psi(Z_n)$ be satisfiable then there is an EAF admissible set containing Ψ (thus an EAF preferred extension that does not contain Φ) and $\langle \mathcal{X}_\Psi, \mathcal{A}_\Psi, \mathcal{D}_\Psi \rangle$ fails to be a positive instance of GS.

For the converse direction suppose that $\langle \mathcal{X}_\Psi, \mathcal{A}_\Psi, \mathcal{D}_\Psi \rangle$ is *not* a positive instance of GS. We show that in this case there is an EAF admissible set containing Ψ whence we deduce that $\Psi(Z_n)$ is satisfiable. It is immediate that $\{s_1, \ldots, s_n\} \subseteq S$ for every EAF preferred extension of $\langle \mathcal{X}_\Psi, \mathcal{A}_\Psi, \mathcal{D}_\Psi \rangle$ since these arguments are unattacked. From our assumption that $\langle \mathcal{X}_\Psi, \mathcal{A}_\Psi, \mathcal{D}_\Psi \rangle$ is not a positive instance of GS there must be some EAF preferred extension, S say, for which $\Phi \notin S$. Were it the case that $\Psi \notin S$ such would contradict S being *maximal* since Φ is EAF acceptable w.r.t. any EAF conflict free subset that does not contain Ψ and does contain $\{s_1, \ldots, s_n\}$. Hence from the premise it follows that $\Psi \in S$ and thus each attack $C_j \rightarrow^S \Psi$ must be countered. The only arguments available for this purpose, however, must come from $\{z_1, \ldots, z_n, \neg z_1, \ldots, \neg z_n\}$. From this subset a satisfying assignment for $\Psi(Z_n)$ is easily formed. We deduce that $\langle \mathcal{X}_\Psi, \mathcal{A}_\Psi, \mathcal{D}_\Psi \rangle$ is accepted as an instance of GS if and only if $\Psi(Z_n)$ is unsatisfiable. □

Corollary 2

a. MF *is* NP*–complete.*
b. SC *is co*NP*-complete.*

Proof: (Outline) For (a) membership in NP follows from the fact that $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$ is accepted as an instance of MF if and only if

$$\exists\, S,\ T \subseteq \mathcal{X},\ x \in \mathcal{X} \text{ s.t. } \quad S \subset T,\ x \text{ is EAF acceptable w.r.t. } S \text{ and } x \text{ is } \textit{not} \text{ EAF acceptable w.r.t. } T$$

From Thm. 1 validating $\langle S, T, x \rangle$ satisfy such conditions is polynomial time computable. To establish MF is NP–hard, it suffices to use a similar reduction (from SAT) to that described in the proof of Thm 2: fixing $S = \{s_1, \ldots, s_n\}$ there is a (strict) superset T of S with Φ not EAF acceptable w.r.t. T – the set T containing, in addition to S, $\{\Psi\}$ and arguments from $\{z_1, \ldots, z_n, \neg z_1, \ldots, z_n\}$ defining a satisfying assignment of $\Psi(Z_n)$. The upper bound for part (b) follows since the complementary problem is in NP: guess subsets, S and T of $\mathcal{X}$ and verify that $S \subset T$, S is stable and T is admissible.[4] The lower bound uses a simplification of the construction in Fig. 2 (without the $\{s_i, +_i, \neg_i\}$ arguments and with $\langle y_i, \langle \Phi, y_i \rangle\rangle$ for each literal $y_i \in \{z_i, \neg z_i\}$, $\langle \Phi, \langle \Psi, \Phi \rangle\rangle$ added to $\mathcal{D}_\Psi$). We omit the straightforward argument that $\Psi(Z_n)$ is unsatisfiable if and only if the constructed EAF is semi-coherent. □

Finally, we address the CFE problem. Firstly, recall that the analogous problem in AFs is trivial since the property of S being non-conflict free is monotonic, i.e. if $S \subseteq \mathcal{X}$ fails to be conflict free in $\langle \mathcal{X}, \mathcal{A} \rangle$ then every superset of S also fails to be conflict free. This property does not hold for VAFs, weighted argument systems [12], and EAFs. The corresponding problems in VAFs – given $S \subseteq \mathcal{X}$ in the VAF $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle$ is there an audience, R, and subset T of $\mathcal{X}$ containing S, such that T is conflict free w.r.t. the audience R – are polynomial time decidable [7], as is the related formulation for weighted argument systems.

The capability of constructing a conflict free extension of S (or deciding that none such is possible) is a prerequisite for determining the scope for S to be part of an EAF admissible set. Thus consideration of the computational complexity of CFE is well motivated. Perhaps surprisingly – in view of the polynomial time procedures in VAFs and weighted argument systems – CFE turns out to be computationally intractable: including justifications for preferences *within* the framework results in complexity akin to that of finding a *particular* audience in a VAF for which a given argument is acceptable.

[4] The authors thank the anonymous reviewer who observed this.

Theorem 3 CFE *is* NP*–complete.*

Proof: Membership in NP is immediate simply by guessing $T \subseteq \mathcal{X}$ and verifying that both T is EAF conflict free and $S \subseteq T$ hold. For NP–hardness we use a reduction from CNF SAT. Given an instance $\Phi(Z_n)$ of CNF SAT with clauses $\{C_1, \ldots, C_m\}$, form the EAF $\langle \mathcal{X}_\Phi, \mathcal{A}_\Phi, \mathcal{D}_\Phi \rangle$ in which

$$\begin{aligned}
\mathcal{X}_\Phi &= \{z_i,\ \neg z_i\ :\ 1 \le i \le n\} \cup \{\ C_j\ :\ 1 \le j \le m\} \cup \{\Phi\} \\
\mathcal{A}_\Phi &= \{\langle z_i, \neg z_i \rangle,\ \langle \neg z_i, z_i \rangle\ :\ 1 \le i \le n\} \cup \\
&\quad\ \{\langle C_j, \Phi \rangle\ :\ 1 \le j \le m\} \\
\mathcal{D}_\Phi &= \{\langle y_k, \langle C_j, \Phi \rangle\rangle\ :\ y_k \in \{z_k, \neg z_k\} \text{ is a literal in } C_j\}
\end{aligned}$$

The instance of CFE is completed by fixing $S = \{\Phi, C_1, \ldots, C_m\}$.[5]

We claim that there is an EAF conflict free extension of $\{\Phi, C_1, \ldots, C_m\}$ in $\langle \mathcal{X}_\Phi, \mathcal{A}_\Phi, \mathcal{D}_\Phi \rangle$ if and only if there is an assignment $\alpha \in \langle \bot, \top \rangle^n$ to Z_n that satisfies Φ.

Suppose first that T with $\{\Phi, C_1, \ldots, C_m\} \subseteq T$ is an EAF conflict free set. Consider the set $T \cap \{z_i,\ \neg z_i\ :\ 1 \le i \le n\}$ and the assignment $\langle \alpha_1, \ldots, \alpha_n \rangle$ for which $\alpha_i = \top$ if $z_i \in T$ and $\alpha_i = \bot$ if $\neg z_i \in T$. Noting that we cannot have both z_i and $\neg z_i$ in T, the assignment $\langle \alpha_1, \ldots, \alpha_n \rangle$ is well defined. We claim that this assignment satisfies $\Phi(Z_n)$. To see this it suffices to observe that since T is EAF conflict free and $\{\Phi, C_1, \ldots, C_m\}$ is not (by reason of the attacks $\{\langle C_j, \Phi \rangle\ :\ 1 \le j \le m\}$) it must be the case that each attack $\langle C_j, \Phi \rangle$ fails with respect to some $y_i \in T \cap \{z_i,\ \neg z_i\ :\ 1 \le i \le n\}$. This can only happen if the corresponding literal occurs in the clause C_j so that the chosen assignment to z_i will render C_j true. From the fact that T is EAF conflict free, every attack $\langle C_j, \Phi \rangle$ must fail; thus the assignment $\langle \alpha_1, \ldots, \alpha_n \rangle$ will satisfy every clause of $\varphi(Z_n)$, i.e. $\varphi(Z_n)$ is satisfiable as required.

Conversely, suppose $\langle \alpha_1, \ldots, \alpha_n \rangle$ satisfies $\varphi(Z_n)$. Consider the subset, T_α of $\{z_i,\ \neg z_i\ :\ 1 \le i \le n\}$ given by

$$T_\alpha \;=\; \{\ z_i\ :\ \alpha_i = \top\} \cup \{\ \neg z_i\ :\ \alpha_i = \bot\}$$

With this, $\{\Phi, C_1, \ldots, C_m\} \cup T_\alpha$ is EAF conflict free: each attack $\langle C_j, \Phi \rangle$ fails given $\langle y_i, \langle C_j, \Phi \rangle\rangle$ where $y_i \in \{z_i, \neg z_i\}$ is the literal in C_j that takes the value $\top$ under $\langle \alpha_1, \ldots, \alpha_n \rangle$.

□

5 Relationship to EAF+

Before addressing developments of our results, we briefly discuss a recently proposed alternative approach: the treatment of EAFs (referred to as EAF+) deriving from the argumentation framework with recursive attacks (AFRA) formalism from Baroni *et al.* [2]. This, as with EAFs, develops the concept of $\mathcal{D}$ (a subset of $\mathcal{X} \times \mathcal{A}$) to one in which arguments in $\mathcal{X}$ may also attack elements of $\mathcal{D}$. Formally,

Definition 5 *([2]) An* EAF+ *is described by* $\langle \mathcal{X}, \mathcal{R}, \mathcal{D}+ \rangle$ *where* $\mathcal{X}$ *is a set of arguments,* $\mathcal{R} \subseteq \mathcal{X} \times \mathcal{X}$*, and* $\mathcal{D}+$ *is a set of pairs* $\langle x, \delta \rangle$

[5] It may be noted that the structure of $\langle \mathcal{X}, \mathcal{A}, \mathcal{D} \rangle$ is similar to that in the standard translation of CNF to an AF originally presented in [10] to show deciding credulous acceptance (w.r.t. admissibility) is NP–complete: instead of literal arguments attacking the clause arguments in which the literal appears, such arguments attack *the attack* by the clause on Φ.

in which $x \in \mathcal{X}$ *and* ($\delta \in \mathcal{R}$ or $\delta \in \mathcal{D}+$). *Given* $\alpha = \langle x, \mathcal{H} \rangle \in \mathcal{D}+$, *the* source *(src) and* target *(trg) of* α *are described by,* $src(\alpha) = x$ *and* $trg(\alpha) = \mathcal{H}$. *Letting* $\mathcal{C} = \mathcal{R} \cup \mathcal{D}+$, *for* $x,\ y \in \mathcal{X} \cup \mathcal{C}$, x *is said to* defeat y ($x \rightarrow_{\mathcal{C}} y$) *whenever any of the following hold*

1. $x,\ y \in \mathcal{X}$ *and* $\langle x, y \rangle \in \mathcal{R}$.
2. $x \in \mathcal{X}$, $y \in \mathcal{C}$ *and* $\langle x, y \rangle \in \mathcal{C}$
3. $x,\ y \in \mathcal{C}$ *and* $trg(x) = y$.
4. $x \in \mathcal{C}$, $y \in \mathcal{X}$ *and* $trg(x) = y$

The basic structures of interest within an EAF+ are derived from the so-called *self-contained pairs*: $\langle \mathcal{S}, \mathcal{T} \rangle$ for $\mathcal{S} \subseteq \mathcal{X}$ and $\mathcal{T} \subseteq \mathcal{C}$ being a self-contained pair whenever $src(\alpha) \in \mathcal{S}$ for all $\alpha \in \mathcal{T}$. Analogues of conflict free, acceptability and admissibility are presented w.r.t. self-contained pairs.

Definition 6 *([2]) Given an* EAF+ $\langle \mathcal{X}, \mathcal{A}, \mathcal{D}+ \rangle$, *a self-contained pair* $\langle \mathcal{S}, \mathcal{T} \rangle$ *is* conflict free *if*

a. $\forall\ x, y \in \mathcal{S}$ *should* $x \rightarrow_{\mathcal{C}} y$ *(i.e.* $\alpha = \langle x, y \rangle \in \mathcal{A}$*) then* $\exists \beta \in \mathcal{T}$ *s.t.* $\beta \rightarrow_{\mathcal{C}} \alpha$.
b. $\forall \gamma \in \mathcal{T},\ \forall \delta \in \mathcal{S} \cup \mathcal{T}\ \ \neg(trg(\gamma) = \delta)$.

An argument $x \in \mathcal{X}$ *is acceptable w.r.t. a self-contained pair* $\langle \mathcal{S}, \mathcal{T} \rangle$ *if* $\forall\ y \in \mathcal{X}$ *s.t.* $\langle y, x \rangle \in \mathcal{A}$ *there is some* $\alpha \in \mathcal{T}$ *s.t.* $\alpha \rightarrow_{\mathcal{C}} y$ or $\alpha \rightarrow_{\mathcal{C}} \langle y, x \rangle$.
In a related way, $\alpha \in \mathcal{C}$ *is acceptable w.r.t. a self-contained pair* $\langle \mathcal{S}, \mathcal{T} \rangle$ *if* $src(\alpha) \in \mathcal{X}$ *is acceptable w.r.t.* $\langle \mathcal{S}, \mathcal{T} \rangle$ and $\forall \beta \in \mathcal{D}+$ *s.t.* $\beta \rightarrow_{\mathcal{C}} \alpha$, $\exists \gamma \in \mathcal{T}$ *for which* $\gamma \rightarrow_{\mathcal{C}} src(\beta)$ *or* $\gamma \rightarrow_{\mathcal{C}} \beta$.
A self-contained pair, $\langle \mathcal{S}, \mathcal{T} \rangle$*, is* admissible *if it is conflict free, and every* $x \in \mathcal{S}$, $\alpha \in \mathcal{T}$ *is acceptable w.r.t.* $\langle \mathcal{S}, \mathcal{T} \rangle$. *Using the (partial) ordering of self-contained pairs, under which* $\langle \mathcal{S}, \mathcal{T} \rangle$ is included in $\langle \mathcal{S}', \mathcal{T}' \rangle$ *if* $\mathcal{S} \subseteq \mathcal{S}'$ and $\mathcal{T} \subseteq \mathcal{T}'$ *gives rise to the concept of* preferred extensions *of* $\langle \mathcal{X}, \mathcal{A}, \mathcal{D}+ \rangle$ *as maximal admissible self-contained pairs.*

Acceptability is defined without recourse to the notion of reinstatement set. However, [2] do not present proposals for defining characteristic functions of self-contained pairs mirroring Defn. 4. Thus a direct comparison of issues such as, e.g. membership in the "grounded extension" and sceptical acceptance are, at present, not possible. We note that by considering the EAF constructed in the proof of Thm. 2, formulations of grounded extension of an EAF+ (even with the constraint $\mathcal{D}+ = \mathcal{D}$ imposed) will either result in different $\mathcal{S} \subseteq \mathcal{X}$ forming the EAF (resp. EAF+) "grounded" extensions or exhibit behaviours whereby membership of this does not guarantee sceptical acceptance w.r.t. preferred extensions. We note that given $\langle \langle \mathcal{X}, \mathcal{A}, \mathcal{D}+ \rangle, \mathcal{S} \rangle$, by a near identical construction to Thm. 3, deciding if $\langle \mathcal{S}, \emptyset \rangle$ can be extended to a *conflict free* self-contained pair $\langle \mathcal{S}', \mathcal{T} \rangle$ (with $\mathcal{S} \subseteq \mathcal{S}'$) is NP–hard (and in NP for finite $\mathcal{D}+$)

6 Conclusions

We have considered a number of computational issues arising in the EAF model of Modgil [16] from both an algorithmic and computational complexity viewpoint. In particular it has been demonstrated that, in keeping with Dung's abstract AF model and other developments of Dung [11], such as those from [4, 6, 12], the central question of deciding acceptability of x w.r.t. a subset S is polynomial time decidable and so, consequently, verifying if S is EAF admissible and construction of the EAF grounded extension may also be carried out by polynomial time procedures. However, in contrast to these positive properties, a number of questions specifically arising from properties of EAF semantics have been shown to be unlikely to admit efficient decision algorithms: in particular determining if the EAF grounded extension is a subset of every EAF preferred extension is coNP–hard; as well as problems relating to acceptability w.r.t. subsets S, T for which $S \subset T$, and the relationship between stable and preferred extensions. In addition, one problem whose analogues in AFs, VAFs, and weighted systems is polynomial time decidable, turns out to be hard within EAFs: that of deciding if a given set S can be extended to an EAF conflict free set.

There are a number of natural directions in which our results can be pursued. In addition to *exact* complexity bounds on GS, it would be of some interest to characterise forms of EAFs that are free from the extreme behaviours giving rise to GS, MF, and SC, or, find classes of EAF in which their presence can be decided efficiently. Finally, there is considerable scope for exploring detailed semantic and algorithmic issues arising from the interaction of the EAF+ and EAF approaches: given suitable formulations for concepts of "grounded extension" in EAF+ may well provide valuable insights not only with respect to purely computational concerns but also into the metalevel modelling of argumentation that is the focus of [17].

REFERENCES

[1] L. Amgoud and C. Cayrol, 'A reasoning model based on the production of acceptable arguments', *Annals of Math. and AI*, **34**, 197–215, (2002).
[2] P. Baroni, F. Cerutti, M. Giacomin, and G. Guida, 'Encompassing attacks to attacks in abstract argumentation frameworks', in *Proc. ECSQARU*, volume 5590 of *LNAI*, pp. 83–94. Springer Verlag, (2009).
[3] P. Baroni, P. E. Dunne, and M. Giacomin, 'Computational properties of resolution-based grounded semantics', in *Proc. 21st IJCAI*, pp. 683–689, (2009).
[4] P. Baroni and M. Giacomin, 'Resolution-based argumentation semantics', in *Proc. 2nd COMMA*, volume 172 of *FAIA*, pp. 25–36. IOS Press, (2008).
[5] P. Baroni and M. Giacomin, 'Semantics of abstract argument systems', in *Argumentation in AI*, eds., I. Rahwan and G. Simari, chapter 2, 25–44, Springer-Verlag, (2009).
[6] T. J. M. Bench-Capon, 'Persuasion in Practical Argument Using Value-based Argumentation Frameworks', *Journal of Logic and Computation*, **13**(3), 429–448, (2003).
[7] T. J. M. Bench-Capon, S. Doutre, and P. E. Dunne, 'Audiences in argumentation frameworks', *Artificial Intelligence*, **171**, 42–71, (2007).
[8] T. J. M. Bench-Capon and P. E. Dunne, 'Argumentation in artificial intelligence', *Artificial Intelligence*, **171**, 619–641, (2007).
[9] T. J. M. Bench-Capon and S. Modgil, 'Case law in extended argumentation frameworks', in *ICAIL*, pp. 118–127, (2009).
[10] Y. Dimopoulos and A. Torres, 'Graph theoretical structures in logic programs and default theories', *Th. Comp. Sci.*, **170**, 209–244, (1996).
[11] 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).
[12] P. E. Dunne, A. Hunter, P. McBurney, S. Parsons, and M. Wooldridge, 'Inconsistency tolerance in weighted argument systems', in *Proc. 8th AAMAS*, pp. 851–858, (2009).
[13] P. E. Dunne and M. Wooldridge, 'Complexity of abstract argumentation', in *Argumentation in AI*, eds., I. Rahwan and G. Simari, chapter 5, 85–104, Springer-Verlag, (2009).
[14] S. Modgil, 'An argumentation based semantics for agent reasoning', in *LADS 07*, pp. 37–53, Durham, UK, (2007).
[15] S. Modgil, 'Labellings and games for extended argumentation frameworks', in *Proc. 21st IJCAI*, pp. 873–878, (2009).
[16] S. Modgil, 'Reasoning about preferences in argumentation frameworks', *Artificial Intelligence*, **173**(9–10), 901–934, (2009).
[17] S. Modgil and T. Bench-Capon, 'Metalevel argumentation', Technical Report ULCS-09-018, Dept. of Comp. Sci., Univ. of Liverpool, (2009).
[18] S. Modgil and M. Luck, 'Argumentation based resolution of conflicts between desires and normative goals', in *Proc. 5th ARGMAS*, pp. 252–263, Estoril, Portugal, (2008).