On Concise Encodings of Preferred Extensions

Paul E. Dunne

Department of Computer Science
University of Liverpool
Liverpool L69 7ZF
United Kingdom
ped@csc.liv.ac.uk

Abstract. Argument Systems provide a rich abstraction within which divers concepts of reasoning, acceptability and defeasibility of arguments, etc., may be studied using a unified framework. Much work has focused on the so-called *preferred* extensions of such systems, which define the maximal (with respect to $\subseteq$) collectively defensible subsets of arguments within a given system of arguments and attack relationship. In this article we address problems related to the following issue. Identification and enumeration of preferred extensions of an argument system is (under the usual complexity theoretic assumptions) computationally demanding: there may be exponentially many and deciding if a given subset S of $\mathcal{X}$ *does* define a preferred set is CO-NP–complete. For a domain which is questioned 'frequently' it may be acceptable to invest this computational effort once, but having done so it is desirable to encapsulate these data in a form which is compact and allows, for example, questions concerning the acceptability of specific arguments to be dealt with efficiently. In this article we consider two 'plausible' approaches to reducing the complexity of deciding if S is a preferred extension of a system $\mathcal{H}$ both of which assume some initial potentially 'expensive' precomputation, invested to reduce time needed in subseqent queries to the system. The first approach examines 'reasonable encoding' approaches; the second is to determine the subset of all defensible arguments providing these as additional data when attempting to decide if S is a preferred extension. It is shown that if certain properties are required of the encoding scheme, then the former approach is feasible only if NP $=$ CO-NP. In the latter case, we show that, even if provided with information regarding which arguments are credulously accepted, the question of whether a subset of arguments defines a preferred extension remains CO-NP–complete.

Keywords: Argument Systems, Preferred Extension, Computational Complexity

1 Introduction

Since they were introduced by Dung [5], Argument Systems have provided a fruitful mechanism for studying reasoning in defeasible contexts. They have proved useful both to theorists who can use them as an abstract framework for the study and comparison of non-monotonic logics, e.g. [1–3], and for those who wish to explore more concrete contexts where defeasibility is central. In the study of reasoning in law, for example, they have been used to examine the resolution of conflicting norms, e.g. [10], especially where this is studied through the mechanism of a dispute between two parties, e.g. [9]. The basic definition below is derived from that given in [5].

Definition 1. *An* argument system *is a pair $\mathcal{H} = \langle \mathcal{X}, \mathcal{A} \rangle$, in which $\mathcal{X}$ is a set of* arguments *and $\mathcal{A} \subset \mathcal{X} \times \mathcal{X}$ is the* attack relationship *for $\mathcal{H}$. Unless otherwise stated, $\mathcal{X}$ is assumed to be* finite*, and $\mathcal{A}$ comprises a set of ordered pairs of* distinct *arguments. A pair $\langle x, y \rangle \in \mathcal{A}$ is referred to as 'x* is attacked by *y' or 'y* attacks *(or* is an attacker of*) x'.*

For R, S subsets of arguments in the system $\mathcal{H}(\langle \mathcal{X}, \mathcal{A} \rangle)$, we say that

a) $s \in S$ *is* attacked *by R if there is some $r \in R$ such that $\langle r, s \rangle \in \mathcal{A}$.*
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 argument in S is acceptable with respect to 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 argument $y \notin S$ is attacked by S.*
g) $\mathcal{H}$ is coherent *if every preferred extension in $\mathcal{H}$ is also a stable extension.*

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

The notation $\mathcal{PE}(\mathcal{H})$ is used to describe the set of all subsets of $\mathcal{X}$ which are preferred extensions of $\mathcal{H}$. Similarly, $\mathcal{SE}(\mathcal{H})$ denotes the set of all stable *extensions of $\mathcal{H}$ and $\mathcal{BE}(\mathcal{H})$ refers to an arbitrary one of these sets. We use n to denote $|\mathcal{X}|$.*

The preferred extensions of an Argument System can be taken as being the consistent positions that can be adopted within the Argument System. Any argument that appears in all preferred extensions will be acceptable in every consistent position, and any argument that appears in no preferred extension cannot be held in any consistent position. This means that the notion can be related to varieties of semantics for non-monotonic reasoning: credulously acceptable arguments will be those that appear in at least one preferred extension, and sceptically acceptable arguments will be those which appear in all preferred extensions. In the context of legal reasoning, the notion allows us to distinguish those arguments which must be accepted, those which can be defended, and those which are indefensible.

To avoid repetition we will subsequently refer to the following decision problems:

PREF-EXT (Preferred Extension)
Instance: An argument system $\mathcal{H}(\mathcal{X}, \mathcal{A})$ and $S \subseteq \mathcal{X}$.
Question: Is $S \in \mathcal{PE}(\mathcal{H})$?
STAB-EXT (Stable Extension)
Instance: An argument system $\mathcal{H}(\mathcal{X}, \mathcal{A})$ and $S \subseteq \mathcal{X}$.
Question: Is S a subset of some $T \in \mathcal{SE}(\mathcal{H})$, i.e. can S be expanded to a stable extension?
PREF-EXT-INF (Preferred Extension given Information)
Instance: An argument system $\mathcal{H}(\mathcal{X}, \mathcal{A})$, $S \subseteq \mathcal{X}$, and $\alpha = \langle a_1, a_2, \ldots, a_n\rangle \in \langle \bot, \top\rangle^n$ an n-tuple of Boolean values such that $a_i = \top$ if and only if the argument x_i is credulously accepted in $\mathcal{H}$.
Question: Is $S \in \mathcal{PE}(\mathcal{H})$?
STAB-EXT-INF (Stable Extension given Information)
Instance: An argument system $\mathcal{H}(\mathcal{X}, \mathcal{A})$, $S \subseteq \mathcal{X}$, and $\alpha = \langle a_1, a_2, \ldots, a_n\rangle \in \langle \bot, \top\rangle^n$ an n-tuple of Boolean values such that $a_i = \top$ if and only if the argument x_i is credulously accepted in $\mathcal{H}$.
Question: Is S a subset of some $T \in \mathcal{SE}(\mathcal{H})$, i.e. can S be expanded to a stable extension?

One difficulty with this formalism is the computational intractability of several decision problems that arise within it: results of Dimopoulos and Torres [4] indicate that deciding if p is credulously accepted in $\mathcal{H}$ is NP–complete and that PREF-EXT is CO-NP–complete (even when S is the empty set). Similarly deciding if $\mathcal{H}$ has any stable extension (i.e. the decision problem $\mathcal{SE}(\mathcal{H}) \neq \emptyset$?) is NP-complete (notice that this is the special case of STAB-EXT when $S = \emptyset$). Dunne and Bench-Capon [7] proved that deciding coherence is $\Pi_2^{(p)}$–complete, deducing the same complexity classification for sceptical acceptance as a consequence. Related work, described in [8], has shown that the sound and complete reasoning method for credulous argumentation introduced by Vreeswijk and Prakken [12] in which reasoning proceeds via a dialogue game requires an exponential number of moves to resolve some disputes.

Suppose one is given a specific argument system $\mathcal{H}(\mathcal{X}, \mathcal{A})$ which describes a 'frequently' used application, for example one incorporating complex legal data in which the justification for various different positions may have to be assessed. Rather than deciding acceptability and support for an argument on each new query to $\mathcal{H}$, we may be prepared to invest some computational effort *once* in the hope that the information elicited as a result may help in reducing the time taken for subsequent queries. For example, one could compute all sets in $\mathcal{BE}(\mathcal{H})$ and then have queries on $\mathcal{H}$ performed with respect to a *representation* (encoding) of this set of subsets. There are, of course, several properties which such a representation should, ideally, satisfy. Suppose $\eta(\mathcal{H})$ is some encoding of $\mathcal{BE}(\mathcal{H})$. We define, informally, two such properties that are the main focus of this paper:

R1. η is *terse*, i.e. $|\eta(\mathcal{H})|$ – the number of *bits* needed – is polynomially bounded in n.
R2. η is *extension tractable*, i.e. given any $S \subseteq \mathcal{X}$, the question $S \in \mathcal{BE}$ can be decided from $\eta(\mathcal{H})$ in time polynomial in $|\eta(\mathcal{H})|$.

One obvious representation scheme is simply to use a table $T_{\mathcal{H}}$ of $|\mathcal{BE}(\mathcal{H})|$ rows, each row being n bits in length, so that if $\langle S_1, \ldots, S_k\rangle$ is an ordering of $\mathcal{BE}(\mathcal{H})$, then

$T_{i,j} = 1$ if and only if $x_j \in S_i$. While this representation meets the criterion specified by R2, it will fail to satisfy R1 in those cases where $|\mathcal{BE}(\mathcal{H})|$ is superpolynomial in n. We note that systems may be defined where this number is $\Omega(3^{n/3})$. Alternatively, the system $\mathcal{H}$ itself is a representation of $\mathcal{BE}(\mathcal{H})$: while satisfying R1 it is, however, unlikely to satisfy R2 in the case $\mathcal{BE} = \mathcal{PE}$ (assuming NP $\neq$ CO-NP).

Given such examples, a natural question to raise is whether these extremes are inherent, or are there representation formalisms that are terse *and* extension tractable – a property we subsequently refer to as *concise*. In the next Section we formalise these concepts and, in Section 3, prove some basic results concerning them. In particular it is shown that in the case of $\mathcal{PE}(\mathcal{H})$ 'effective' concise encodings are not possible, in general, unless NP $=$ CO-NP. In contrast concise encoding schemes for $\mathcal{SE}(\mathcal{H})$ are easy to construct. We note that this provides another example of a property which is considerably 'easier' under stable semantics than under preferred semantics for argument systems, cf. [2, 3]. A further indication of the computational difficulties arising in considering preferred extensions is given in the concluding result of Section 3 where the problem PREF-EXT-INF is shown to be CO-NP–complete: thus, even if the defensibility status for every argument in $\mathcal{X}$ is supplied as part of an instance, the problem of deciding whether a given subset S is a preferred extension does not become any easier, i.e. remains CO-NP–complete. Discussion and conclusions occupy Section 4.

2 Definitions

In the remainder of this paper the following notational conventions are used.

$\mathcal{X}_n$ is a set of n arguments $\{x_1, x_2, \ldots, x_n\}$.

For an n element set X_n, $\wp(X_n)$ is the set of all subsets of X_n.

The following definition formalises our abstract concept of *encoding scheme* for the set $\mathcal{BE}$ in an argument system $\mathcal{H}$.

Definition 2. *A $\mathcal{BE}$* encoding scheme *is a pair $\langle \eta, P \rangle$ where*

$$\eta : \bigcup_{n=0}^{\infty} \{\mathcal{H}(\mathcal{X}_n, \mathcal{A}) : \mathcal{H} \textit{ is an argument system } \} \rightarrow \{0,1\}^*$$

is a mapping from argument systems to finite binary words, and P is a deterministic *Turing machine program, that takes as input a pair $\langle \eta(\mathcal{H}), S \rangle$ in which $S \subseteq \mathcal{X}_n$, accepting if and only if $S \in \mathcal{BE}(\mathcal{H})$.*

Before proceeding there are several points that should be noted. First we observe that it is not insisted that schemes be *uniform*, i.e. it is *not* required that there is an algorithm which given $\mathcal{H}$ computes $\eta(\mathcal{H})$. Of course, in 'practical' schemes, one would wish to have some mechanism for automating this translation. In order to capture some sense of 'practical' scheme, we introduce the notion of *verifiability*

Definition 3. *A* verifiable *$\mathcal{BE}$* encoding scheme *is a triple $\langle \eta, P, Q \rangle$ in which $\langle \eta, P \rangle$ is a $\mathcal{BE}$ encoding scheme and Q a (non-deterministic) Turing Machine program, that is given $w \in \{0,1\}^*$ and $\mathcal{H}(\mathcal{X}, \mathcal{A})$ as input, accepting if $w = \eta(\mathcal{H})$.*

Secondly, the definition provides a foundation for introducing more 'sophisticated' schemes other than binary words.

With this abstract idea of encoding scheme we can formalise the notions of *terse* and *extension tractable* outlined earlier.

Definition 4. *Let $\langle \eta, P, Q\rangle$ be a verifiable $\mathcal{BE}$ encoding scheme. We say $\langle \eta, P\rangle$ is* terse *if there is a constant k, such that for all $\mathcal{H}(\mathcal{X}_n, \mathcal{A})$, $|\eta(\mathcal{H})| \leq n^k$; it is* extension tractable *if there is a constant k such that: given $\eta(\mathcal{H})$ and $S \in \wp(X_n)$ as input, P decides if $S \in \mathcal{BE}(\mathcal{H})$ taking at most $(n + |\eta(\mathcal{H})|)^k$ steps. Finally, $\langle \eta, P, Q\rangle$ is* effective *if there is a constant k for which Q has an accepting computation of $w = \eta(\mathcal{H})$ in non-deterministic time bounded by $(n + |w|)^k$. A $\mathcal{BE}$ encoding scheme $\langle \eta, P\rangle$ is* concise *if it is both terse and extension tractable. A* verifiable *$\mathcal{BE}$ encoding scheme $\langle \eta, P, Q\rangle$ is* usefully concise *if it is terse, extension tractable, and effective.*

We now present some examples of encoding approaches.

Tabular Representation Given $\mathcal{H}$, $tab(\mathcal{H})$ is the $n|\mathcal{BE}(\mathcal{H})|$-bit word in which bit $t_{(i-1)n+j} = 1$ if and only if $x_j \in S_i$ where $1 \leq j \leq n$ and $\langle S_1, S_2, \ldots, S_r\rangle$ is an ordering of $\mathcal{BE}(\mathcal{H})$. If the algorithm P in $\langle tab, P\rangle$ is chosen to be an appropriate table look up method, then $\langle tab, P\rangle$ is extension tractable. It is not, however, terse.

Representations via Propositional Logic Functions Given any $\mathcal{H}(\mathcal{X}_n, \mathcal{A})$ there is a unique propositional logic function, $f_{\mathcal{H}}(X_n)$ definable from $\mathcal{BE}(\mathcal{H})$ as follows. For any $S \in \wp(\mathcal{X}_n)$ let the instantiation, α_S of the propositional variables X_n be $x_i = \top$ if $x_i \in S$ and $x_i = \bot$ if $x_i \notin S$. The function $f_{\mathcal{H}}(X_n)$ takes the value $\top$ on exactly those instantiations α_S for which $S \in \mathcal{BE}(\mathcal{H})$.

Given this approach, any representation formalism for arbitrary n-argument propositional logic functions serves as a basis for a $\mathcal{BE}$ encoding scheme, e.g. truth-tables, propositional formulae over a finite complete basis, etc.

3 Properties of Usefully Concise Encoding Schemes

We first observe that construction of usefully concise encoding schemes for stable extensions is trivial: the problem of deciding, given $\mathcal{H}(\mathcal{X}, \mathcal{A})$ and $S \subseteq \mathcal{X}$ whether $S \in \mathcal{SE}(\mathcal{H})$ is polynomial-time solvable; therefore since $\mathcal{H}(\mathcal{X}, \mathcal{A})$ already defines a terse encoding of $\mathcal{SE}(\mathcal{H})$ with an appropriate decision algorithm we have a concise $\mathcal{SE}$ encoding scheme. This can be extended to give a usefully concise scheme, by encoding $\mathcal{H}$ as its n^2 element adjacency matrix, so that $w = \eta(H)$ is decided in $|w|$ steps.

In contrast to the easy construction above, usefully concise encoding schemes for *preferred* extensions are 'unlikely' to exist.

Theorem 1. *If* NP $\neq$ CO-NP *then usefully concise $\mathcal{PE}$ encoding schemes do not exist.*

Proof. Suppose NP $\neq$ CO-NP and that for the sake of contradiction, $\langle \eta, P, Q\rangle$ is a usefully concise $\mathcal{PE}$ encoding scheme. We show that $\langle \eta, P, Q\rangle$ can be used as the basis of an NP decision method for PREF-EXT. Since this problem is CO-NP–complete it follows that the existence of such a decision method would imply NP = CO-NP. Since $\langle \eta, P, Q \rangle$ is terse there is some constant k such that $|\eta(\mathcal{H})| \leq n^k$ for any n-argument system $\mathcal{H}$. Our NP algorithm is as follows: given an instance $\langle \mathcal{H}(\mathcal{X}, \mathcal{A}), S \rangle$ of PREF-EXT non-deterministically choose a sequence $\beta(\mathcal{H})$ of (at most) $|\mathcal{X}|^k$ bits. Then simulate Q on input $\langle \mathcal{H}, \beta(\mathcal{H}) \rangle$. If $\beta(\mathcal{H}) = \eta(\mathcal{H})$ Q will have an accepting computation of polynomial length (since Q is effective). Finally, the program, P, is run with input $\langle \beta(\mathcal{H}), S \rangle$. Since $|\beta(\mathcal{H})| \leq |\mathcal{X}|^k$ and P is a deterministic polynomial time computation, the (non-deterministic) algorithm runs in time polynomial in $\mathcal{X}$. To see that the algorithm accepts instances for which $S \in \mathcal{PE}(\mathcal{H})$ it suffices to observe that *if* $S \in \mathcal{PE}(\mathcal{H})$ then there is *some* choice of $\beta(\mathcal{H})$ that will correspond to $\eta(\mathcal{H})$, be accepted by Q and on which the extension tractable algorithm P will accept $\langle \beta(\mathcal{H}), S \rangle$.□

It should be noted that the argument used in the proof requires the assumption that $\langle \eta, P, Q \rangle$ is effective. The reason being that if P were invoked directly on the word $\beta(\mathcal{H})$ then *every* instance would be accepted: given a subset S there is certainly some argument system, $\mathcal{G}$ for which $\mathcal{PE}(\mathcal{G}) = \{S\}$ – the system with $|S|$ isolated arguments – and if $\beta(\mathcal{H}) = \eta(\mathcal{G})$ then the instance is accepted regardless of whether $S \in \mathcal{PE}(\mathcal{H})$.

Theorem 1 indicates that even if one is prepared to invest considerable computational effort in *constructing* an encoding $\eta(\mathcal{H})$, such effort will not aid in testing $S \in \mathcal{PE}(\mathcal{H})$ if the encoding form is terse *and* effective.

The problem, PREF-EXT-INF in allowing knowledge regarding the set of credulously accepted arguments to be given for free, can be seen as defining an alternative *terse* encoding scheme. We note that this scheme is *not* effective (in our usage) (assuming NP $\neq$ CO-NP) since in the encoding as $\langle \mathcal{H}, \alpha \rangle$, should any bit of α be $\perp$ indicating the associated argument is *not* credulously accepted, we cannot test $w = \eta(\mathcal{H})$ using an NP computation.

Thus, since determining the set of credulously accepted arguments in a system, may at worst involve similar computational expenditure to that of enumerating preferred extensions, the complexity of the problem PREF-EXT-INF is of some interest: the implied encoding scheme is not one which is within the scope of Theorem 1. Our next result shows that PREF-EXT-INF is no easier than PREF-EXT.

Theorem 2. PREF-EXT-INF *is* CO-NP*–complete.*

Proof. Membership in CO-NP is immediate from the fact that PREF-EXT $\in$ CO-NP. To show that PREF-EXT-INF is CO-NP–hard, we give a reduction from the problem of deciding is a propositional formula in 3–CNF is *unsatisfiable*: 3-UNSAT. Let

$$\Phi(X_n) = \bigwedge_{i=1}^{m} C_i \ = \ \bigwedge_{i=1}^{m} (\, y_{i,1} \vee y_{i,2} \vee y_{i,3} \,)$$

be an instance of 3-UNSAT, so that each $y_{i,j}$ is a literal from $\{x_1, \ldots, x_n, \neg x_1, \ldots, \neg x_n\}$. Let $\{x_{n+1}, x_{n+2}\}$ be two new propositional variables (i.e. not belonging to X_n) and consider the CNF formula $\Psi(X_n, x_{n+1}, x_{n+2})$ defined as,

$$\bigwedge_{i=1}^{m} (C_i \vee \neg x_{n+1} \vee x_{n+2}) \wedge (C_i \vee x_{n+1} \vee \neg x_{n+2})$$

The following properties of $\Psi(X_n, x_{n+1}, x_{n+2})$ are easily verified.

a) For each variable $y \in \{X_n, x_{n+1}, x_{n+2}\}$ there is *some* satisfying instantiation for Ψ under which $y = \top$.
b) For each variable $y \in \{X_n, x_{n+1}, x_{n+2}\}$ there is *some* satisfying instantiation for Ψ under which $y = \bot$.
c) There is a satisfying instantiation for Ψ under which $x_{n+1} = \top$ and $x_{n+2} = \bot$ if and only if $\Phi(X_n)$ is satisfiable.

We use $\Psi(X_n, x_{n+1}, x_{n+2})$ to build an argument system $\mathcal{H}_\Psi(\mathcal{V}, \mathcal{A})$. The argument set $\mathcal{V}$ contains $2(m + n + 3)$ arguments labelled

$$V = \{\Psi, \chi\} \cup \{x_i, \bar{x}_i : 1 \le i \le n + 2\} \cup \{C_j^{(1)}, C_j^{(2)} \ : \ 1 \le j \le m\}$$

The attack relationship, $\mathcal{A}$, comprises

1. $\{\langle x_i, \bar{x}_i\rangle, \langle \bar{x}_i, x_i\rangle \ : \ 1 \le i \le n + 2\}$
2. $\{\langle \chi, \bar{x}_i\rangle, \langle \chi, x_i\rangle \ : \ 1 \le i \le n\}$
3. $\{\langle y_{i,j}, C_i^{(k)}\rangle \ : \ 1 \le k \le 2,\ 1 \le i \le m,\ y_{i,j} \in C_i\}$
4. $\{\langle C_i^{(1)}, \Psi\rangle, \langle C_i^{(2)}, \Psi\rangle \ : \ 1 \le i \le m\}$
5. $\{\langle \bar{x}_{n+1}, C_i^{(1)}\rangle, \langle x_{n+2}, C_i^{(1)}\rangle \ : \ 1 \le i \le m\}$
6. $\{\langle x_{n+1}, C_i^{(2)}\rangle, \langle \bar{x}_{n+2}, C_i^{(2)}\rangle \ : \ 1 \le i \le m\}$
7. $\{\langle \Psi, \chi\rangle\}$.

In the system $\mathcal{H}_\Psi(\mathcal{V}, \mathcal{A})$ all *except* the arguments

$$\{\chi\} \cup \{C_j^{(1)}, C_j^{(2)} \ : \ 1 \le j \le m\}$$

are credulously accepted. To see this first observe that a satisfying instantiation α of $\Psi(X_n, x_{n+1}, x_{n+2})$ induces a preferred extension of $\mathcal{H}_\Psi$ containing Ψ together with the $n + 2$ arguments corresponding to literals evaluating to $\top$ under α. From properties (a) and (b) above we deduce that Ψ is credulously accepted and each of the arguments x_i, $\bar{x}_i$ for $1 \le i \le n + 2$. In order for an argument $C_i^{(k)}$ to be credulously accepted, an admissible set containing it, would have to contain $\bar{y}_{i,j}$ for each $y_{i,j} \in C_i$. These arguments, however, are attacked by χ: Ψ (the only attacker of χ) cannot be included in an admissible set containing $C_i^{(k)}$. Similarly χ cannot be credulously accepted, since the only attackers of its attacker - Ψ - are $C_i^{(k)}$.

The instance of PREF-EXT-INF constructed from $\Phi(X_n)$ is $\langle \mathcal{H}_\Psi, \{x_{n+1}, \bar{x}_{n+2}\}, \alpha_{\mathrm{CA}}\rangle$ where α_{CA} is the $2(m + n + 3)$ tuple indicating the credulously accepted arguments in $\mathcal{H}_\Psi$ as described in the preceding paragraph.

We claim that $\{x_{n+1}, \bar{x}_{n+2}\} \in \mathcal{PE}(\mathcal{H}_\Psi)$ if and only if $\Phi(X_n)$ is unsatisfiable. First suppose that $\{x_{n+1}, \bar{x}_{n+2}\} \in \mathcal{PE}(\mathcal{H}_\Psi)$. Assume, for the sake of contradiction, that $\Phi(X_n)$ is satisfiable and let α be a satisfying instantiation of X_n for Φ. Consider the subset, S_α of $\mathcal{V}$ given by,

$$\{x_{n+1}, \bar{x}_{n+2}\} \cup \{\Psi\} \cup \{x_i \ : \ \alpha_i = \top\} \cup \{\bar{x}_i \ : \ \alpha_i = \bot\}$$

We claim that S_α is an admissible set (in fact $S_\alpha \in \mathcal{PE}(\mathcal{H}_\Psi)$). To see this first observe that S_α is conflict-free and consider any argument $p \in \mathcal{V}$ that attacks some argument of S_α. If $p = \chi$, then $\Psi \in S_\alpha$ attacks p; if $p = y \in \{x_i : 1 \leq i \leq n+2\}$, then $\bar{y} \in S_\alpha$ attacks p; similarly if $p = \bar{y} \in \{\bar{x}_i : 1 \leq i \leq n+2\}$ then $y \in S_\alpha$ attacks p. If $p \in \{C_i^{(2)} : 1 \leq i \leq m\}$ then x_{n+1} attacks p. We are left with the case $p \in \{C_i^{(1)}\}$. Since α satisfies Φ some literal $y_{i,j}$ of C_i must take the value $\top$ under α. Now we find an attack on $C_i^{(1)}$ with the corresponding x_i or $\bar{x}_i$ argument in S_α. We deduce that $\Phi(X_n)$ satisfiable would contradict the assumption $\{x_{n+1}, \bar{x}_{n+2}\} \in \mathcal{PE}(\mathcal{H}_\Psi)$.

On the other hand, suppose that $\Phi(X_n)$ is unsatisfiable. We show that $\{x_{n+1}, \bar{x}_{n+2}\} \in \mathcal{PE}(\mathcal{H}_\Psi)$.

Certainly $\{x_{n+1}, \bar{x}_{n+2}\}$ is admissible. Consider any $S \subset \mathcal{V}$ for which $\{x_{n+1}, \bar{x}_{n+2}\} \subseteq S$ and with S admissible. Let $p \in S/\{x_{n+1}, \bar{x}_{n+2}\}$. If $p = y \in \{x_i, \bar{x}_i : 1 \leq i \leq n\}$, then in order to counter-attack the attack by χ on y, the argument Ψ must be in S. If $\Psi \in S$, then for each argument $C_i^{(1)}$, S must contain some argument corresponding to a literal $y_{i,j}$ of C_i. Since S is assumed admissible, it follws that the set of literals identified do not conflict. Choosing an instantiation of X_n which makes each of these literals take the value $\top$ will satisfy Φ. This, however, contradicts the assumption that Φ was unsatisfiable. We deduce that $S/\{x_{n+1}, \bar{x}_{n+2}\} = \emptyset$ and therefore $\{x_{n+1}, \bar{x}_{n+2}\} \in \mathcal{PE}(\mathcal{H}_\Psi)$ as claimed.□

The following Corollary is easily obtained,

Corollary 1. STAB-EXT-INF *is* NP*–complete.*

Proof. Using the construction of the Theorem, $\{x_{n+1}, \bar{x}_{n+2}\}$ can be developed to a *stable* extension, if and only if $\Phi(X_n)$ is satisfiable.□

4 Conclusions

The principal focus of this article has been in deriving negative results concerning various mechanisms for reducing the complexity of deciding $S \in \mathcal{PE}(\mathcal{H})$ though 'expensive' precomputation. Thus for a rather general notion of 'useful encoding scheme' it has been shown that such approaches are unlikely to succeed. There remains, of course, the possibility that specific sub-classes of argument system are amenable to concise encoding approaches.

Another direction for further work arises from the fact that the requirement for $\langle \eta, P, Q \rangle$ to be *effective* is rather strong and, as we have seen in Theorem 2 does not apply to what might be regarded as otherwise 'reasonable' approaches. One possible encoding approach concerning which Theorem 1 will not in general apply is the following.

Recall that $f_{\mathcal{H}}(X_n)$ is the propositional logic *function* for which $f_{\mathcal{H}}(\alpha_S) = \top$ if and only the subset S of $\mathcal{X}$ indicated by α_S is in $\mathcal{PE}(\mathcal{H})$. Given $f_{\mathcal{H}}(X_n)$ one might represent this using a suitable *propositional formula*. Of course, there are (infinitely) many equivalent formulae in this regard. Suppose for a propositional formula $\Phi(X_n)$ over the (binary operation) basis of $\{\wedge, \vee, \neg\}$ we define the *length* of $\Phi(X_n)$ as its total number of occurrences of literals., denoting this $|\Phi(X_n)|$. Now consider the following measures.

$$L(\mathcal{H}) =_{def} \min\{|\Phi(X_n)| \ : \ \Phi(X_n) \text{ represents } f_{\mathcal{H}}(X_n)\}$$

$$L(n) = \max_{\mathcal{H}(\mathcal{X}_n,\mathcal{A})} \{L(\mathcal{H})\}$$

Informally, $L(n)$ is given by: for each different n-argument $\mathcal{H}$ identify the shortest formula representing $f_{\mathcal{H}}$; $L(n)$ is then the maximum of these values. The function $L(n)$ is well-defined (and computable, albeit by highly infeasible mechanisms). It is certainly the case that encoding $\mathcal{PE}(\mathcal{H})$ by a propositional formula $\Phi(X_n)$ is an *extension tractable* approach: to test $S \in \mathcal{PE}(\mathcal{H})$ simply evaluate $\Phi(\alpha_S)$. This is unlikely to be *effective*: given $w \in \{0,1\}^*$, even though (assuming some standard encoding of propositional formulae, e.g. [6, p.273]) one could determine whether w encodes *some* Φ, it is unlikely that one can test within NP if this represents $f_{\mathcal{H}}$. In summary we have *open* the posssibility that propositional formulae offer *concise* encodings of $\mathcal{PE}(\mathcal{H})$ since such are not ruled out by Theorem 1.

Problem 1. Do propositional formulae admit concise $\mathcal{PE}$ encodings, i.e. is there any $k \in \mathbb{N}$ such that $L(n) = O(n^k)$?

It should be noted that the classic information-theoretic argument of Riordan and Shannon[11], (cf [6, pp.273–274]) does not help in proving superpolynomial lower bounds for the class of propositional functions considered: even if it is assumed that each distinct n argument system encodes a different preferred extension set, the lower bound on $L(n)$ implied by this is only $n^2/\log n$.

As a final open question we mention the following decision problem. As our starting point for building an encoding of $\mathcal{PE}(\mathcal{H})$ we have assumed that the argument system $\mathcal{H}$ is provided as the instance. One might 'relax' this and assume that an arbitrary *subset* of $\wp(X_n)$ is given and we wish to encode *only* those subsets that correspond to $\mathcal{PE}(\mathcal{H})$ for *some* $\mathcal{H}(\mathcal{X},\mathcal{A})$. Thus, we have the following decision problem:
REALISABLE
Instance: $\mathcal{S} = \{S_1, S_2, \ldots, S_k\} \subseteq \wp(\mathcal{X}_n)$.
Question: Does there exist an argument system $\mathcal{H}(\mathcal{X}_n,\mathcal{A})$ for which $\mathcal{PE}(\mathcal{H}) = \mathcal{S}$?

Problem 2. Determine the complexity classifaction of REALISABLE.

We conjecture that REALISABLE $\in$ P, which would follow by proving that the following condition (which is easily shown to be necessary) is also sufficient for $\mathcal{S}$ to be realisable.

$$\forall T \subseteq \mathcal{X}_n \forall \{x,y\} \subseteq T\ (\exists S_i \in \mathcal{S}\ S_i \subset T) \Rightarrow (\forall S_i \in \mathcal{S}\ \{x,y\} \not\subseteq S_i)$$

This condition can be tested in polynomial-time simply by restricting T to range over those supersets of $S_i \in \mathcal{S}$ formed by adding a single new argument.

References

1. A. Bondarenko, P. M. Dung, R. A. Kowalski, and F. Toni. An abstract, argumentation-theoretic approach to default reasoning. *Artificial Intelligence*, 93(1–2):63–101, 1997.
2. 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.
3. 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.
4. Y. Dimopoulos and A. Torres. Graph theoretical structures in logic programs and default theories. *Theoretical Computer Science*, 170:209–244, 1996.
5. P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reason, logic programming, and N-person games. *Artificial Intelligence*, 77:321–357, 1995.
6. P. E. Dunne. *The complexity of Boolean networks*. Academic Press, 1988.
7. P.E. Dunne and T.J.M. Bench-Capon. Coherence in finite argument systems. Technical Report ULCS-01-006, Dept. of Comp. Sci., Univ. of Liverpool, (submitted), December 2001.
8. P.E. Dunne and T.J.M. Bench-Capon. Two party immediate response disputes: Properties and efficiency. Technical Report ULCS-01-005, Dept. of Comp. Sci., Univ. of Liverpool, (submitted), October 2001.
9. H. Jakobovits and D. Vermeir. Dialectic semantics for argumentation frameworks. In *Proceedings of the Seventh International Conference on Artificial Intelligence and Law (ICAIL-99)*, ACM SIGART, pages 53–62, N.Y., June 1999. ACM Press.
10. H. Prakken. *Logical Tools for Modelling Legal Argument*. Kluwer Academic Publishers, 1997.
11. J. Riordan and C.E. Shannon. The number of two–terminal series–parallel networks. *Jnl. Math. and Phys.*, 21:83–93, 1942.
12. 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.