Argumentation Frameworks with Strong and Weak Constraints: Semantics and Complexity

Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, Irina Trubitsyna
Department of Informatics, Modeling, Electronics and System Engineering,
University of Calabria, Italy
{g.alfano, greco, fparisi, i.trubitsyna}@dimes.unical.it

Abstract

Dung's abstract Argumentation Framework (AF) has emerged as a central formalism in formal argumentation. Key aspects of the success and popularity of Dung's framework include its simplicity and expressiveness. Integrity constraints help to express domain knowledge in a compact and natural way, thus keeping easy the modeling task even for problems that otherwise would be hard to encode within an AF. In this paper, after providing an intuitive semantics based on Lukasiewicz's logic for AFs with (strong) constraints, called Constrained AFs (CAFs), we propose Weak constrained AFs (WAFs) that enhance CAFs with weak constraints. Intuitively, these constraints can be used to find "optimal" solutions to problems defined through CAFs. We provide a detailed complexity analysis of CAFs and WAFs, showing that strong constraints do not increase the expressive power of AFs in most cases, while weak constraints systematically increase the expressive power of CAFs under several well-known argumentation semantics.

Introduction

Formal argumentation is an important field of research within Artificial Intelligence (Bench-Capon and Dunne 2007; Simari and Rahwan 2009; Atkinson et al. 2017). In particular, Dung's abstract Argumentation Framework (AF) has emerged as a central formalism for modelling disputes between two or more agents (Dung 1995). An AF consists of a set of *arguments* and a binary *attack* relation over the set of arguments that specifies conflicts between arguments (if argument a attacks argument b, then b is acceptable only if a is not). Hence, arguments are abstract entities whose role is determined by attacks. We can think of an AF as a directed graph whose nodes represent arguments and edges represent attacks. As for graph theory, an important aspect of the success of Dung's framework is that it is a simple yet powerful formalism. The meaning of an AF is given in terms of argumentation *semantics*, which intuitively tell us the sets of arguments (called *extensions*) that can collectively be accepted to support a point of view in a dispute.

Despite the expressive power and generality of AFs, in some cases it is difficult to accurately model domain knowledge by an AF in a natural and easy-to-understand way.

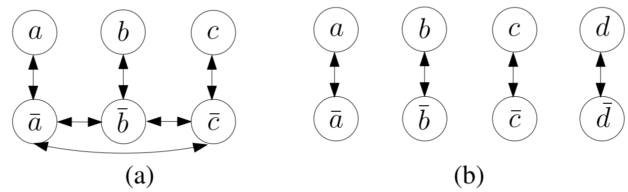

Figure 1: (a) AF Λ of Example 1; (b) AF Λ' of Example 3.

For this reason, Dung's framework has been extended by the introduction of further constructs, such as preferences (Amgoud and Cayrol 1998; Modgil and Prakken 2013) and *integrity constraints* (Coste-Marquis, Devred, and Marquis 2006; Arieli 2015), to achieve more comprehensive, natural, and compact ways of representing useful relationships among arguments. In particular, enhancing AFs with constraints allow us to naturally and compactly express domain conditions that need to be taken into account to filter out unfeasible solutions, as illustrated in the following example.

Example 1. Albert, Betty and Charlie wish to attend a basketball game on Saturday evening, but only two tickets are available. In an attempt to model this situation by an AF Λ, the following six arguments can be used: a (resp., b, c) states that Albert (resp., Betty, Charlie) attends the game, whereas $\bar{a}$ (resp., $\bar{b}$, $\bar{c}$) states that Albert (resp., Betty, Charlie) does not attend the game. The direct graph encoding Λ is shown in Figure 1(a), where double arrows are used to represent mutually attacks between arguments. Specifically, every argument a (resp., b, c) attacks and is attacked by argument $\bar{a}$ (resp., $\bar{b}$, $\bar{c}$), i.e., only one of them can be accepted. Moreover, every argument $\bar{a}$ (resp., $\bar{b}$, $\bar{c}$) is attacked by the other two arguments $\bar{b}$ and $\bar{c}$ (resp., $\bar{a}$ and $\bar{c}$; $\bar{a}$ and $\bar{b}$) since the argument that Albert (resp., Betty, Charlie) attends the game can be accepted only if one of the arguments stating that Betty or Charlie (resp., Albert or Charlie; Albert or Betty) do not attend the game is accepted. Thus, the set of attacks between every pair in $\{\bar{a}, \bar{b}, \bar{c}\}$ models the fact that at most one argument among $\bar{a}$, $\bar{b}$ and $\bar{c}$ can be accepted and then, as a consequence, at least two arguments among a, b and c can be accepted, i.e., all available tickets are sold.

The extensions of AF Λ under the well-known preferred and stable semantics are $E_1 = \{a, b, \bar{c}\}$, $E_2 = \{a, \bar{b}, c\}$, $E_3 = \{\bar{a}, b, c\}$, and $E_4 = \{a, b, c\}$, where the presence of an

		Framework									
		AF		CAF		WAF			LWAF	NCAF	NWAF
	$\mathcal{S}$	$CA_{\mathcal{S}}$	$SA_{\mathcal{S}}$	$CA_{\mathcal{S}}$	$SA_{\mathcal{S}}$	$CA_{\mathtt{ms}\mathcal{S}}$	$SA_{\mathtt{ms}\mathcal{S}}$	$CA_{\mathtt{mc}\mathcal{S}}/SA_{\mathtt{mc}\mathcal{S}}$	$CA_{\mathcal{S}}/SA_{\mathcal{S}}$	$CA_{\mathcal{S}}$	$CA_{\mathtt{ms}\mathcal{S}}$
Semantics	co	NP-c	*P*	NP-c	coNP-c	Σ_2^P-c	Π_2^P-c	$\Delta_2^P[\log n]$-c	Δ_2^P-c	NP-c	Σ_2^P-c
	st	NP-c	coNP-c	NP-c	coNP-c	Σ_2^P-c	Π_2^P-c	$\Delta_2^P[\log n]$-c	Δ_2^P-c	NP-c	Σ_2^P-c
	pr	NP-c	Π_2^P-c	NP-h, Σ_2^P	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_3^P-c	$\Delta_3^P[\log n]$-c	Δ_3^P-c	NP-c	Σ_2^P-c
	sst	Σ_2^P-c	Π_2^P-c	Σ_2^P-c	Π_2^P-c	Σ_3^P-c	Π_3^P-c	$\Delta_3^P[\log n]$-c	Δ_3^P-c	Σ_2^P-c	Σ_3^P-c

Table 1: Complexity of $CA_{\mathcal{S}}$ and $SA_{\mathcal{S}}$ under complete (co), stable (st), preferred (pr), and semi-stable (sst) semantics. For any complexity class C, C-c (resp., C-h) means C-complete (resp., C-hard). An interval C-h, C' means C-hard and in C'.

argument in one of the 4 solutions means that it is accepted. However, there is a solution (i.e., E_4) which is not feasible, because only two tickets are available, meaning that only two people could attend the game.

To overcome such a situation, and thus providing a natural and compact way for expressing such kind of conditions, the use of constraints has been proposed. Considering our example, the constraint

$$\kappa = a \wedge b \wedge c \Rightarrow \mathtt{f}$$

can be used. It states that the propositional formula $a \wedge b \wedge c$ must be false. That is, feasible solutions must satisfy the condition that the 3 arguments a, b, and c are not jointly accepted, i.e., Albert, Betty and Charlie cannot attend the game together. The effect of using constraint κ is that E_4 is discarded from the set of solutions of our problem. □

The use of constraints in AFs has been firstly proposed in (Coste-Marquis, Devred, and Marquis 2006) and then further investigated in (Arieli 2013, 2015, 2016). The constrained argumentation frameworks in (Arieli 2013) and (Arieli 2016) are particular cases of those in (Arieli 2015) as the set of constraints is restricted to atomic formulae only.

We call an AF with constraints a *Constrained AF* (CAF).

Although constraints in a CAF allow restricting the set of feasible solutions, they do not help in finding "best" or preferable solutions. Considering our running example, Albert, Betty and Charlie may agree on the fact that *"if there are only two tickets available then Albert and Betty should preferably attend the game"*. To express this kind of conditions, in this paper we introduce *weak* constraints, that is, constraints that are required to be satisfied *if possible*. Syntactically, they have the same form of (strong) constraint except that the implication symbol $\rightarrow$ is used (instead of $\Rightarrow$). Intuitively, these constraints can be used to find "optimal" solutions to a problem defined by means of an AF or a CAF.

A CAF with the addition of weak constraints is said to be a *Weak constrained Argumentation Framework* (WAF).

Example 2. Consider a WAF obtained by adding to the CAF of Example 1 the weak constraint $\mathtt{t} \rightarrow a \wedge b$, stating that is desirable that Albert and Betty attend the game together. Herein, t denotes the truth value true. Then, extension $\{a, b, \bar{c}\}$ is selected as the "best" one. □

Weak constraints (also called relaxed constraints in some contexts) have been considered in several research areas, including Mathematical Programming with Equilibrium Constraints (Ramos 2019), Answer Set Programming (Buccafurri, Leone, and Rullo 2000; Greco 1998), and for modelling and solving optimization problems (Faber et al. 2016; Li and Lv 2014). In particular, concerning the field of Answer Set Programming, weak constraints have been implemented in DLV (Alviano et al. 2017), a disjunctive logic programming system with (total) stable models semantics.

The use of strong and weak constraints substantially reduces the effort needed to figure out how to define an AF that models a given problem. In fact, as said before, constraints facilitate to express knowledge in a more compact and easy to understand way. For instance, the problem presented in Example 1, has been represented through an AF which expresses the condition that "at most one argument among $\bar{a}$, $\bar{b}$ and $\bar{c}$ can be accepted" and then, as a consequence, at least two arguments among a, b and c can be accepted. However, this condition is not easy to be generalized if we have more than three people. Suppose there is a fourth guy, David, who wish to attend the game, and there are again only two available tickets. After adding the arguments d (David attends the game) and $\bar{d}$ (David does not attend the game) to AF Λ of Figure 1(a), we cannot use the same reasoning as in Example 1 to model the fact that two of the four people attend the game. In fact, having the attacks between every pair in $\{\bar{a}, \bar{b}, \bar{c}, \bar{d}\}$ does not model this situation (it models that at least three of the four people attend the game). Remarkably, using strong and weak constraints allow for using a common reasoning pattern to generalize to this more complex situation, even starting from an AF having a simpler structure.

Example 3. Consider a WAF consisting of AF Λ' of Figure 1(b) and the following sets of strong and weak constraints:
$\mathcal{C} = \{\kappa, a \wedge b \wedge d \Rightarrow \mathtt{f}, a \wedge c \wedge d \Rightarrow \mathtt{f}, b \wedge c \wedge d \Rightarrow \mathtt{f}\}$;
$\mathcal{W} = \{\mathtt{t} \rightarrow a, \mathtt{t} \rightarrow b, \mathtt{t} \rightarrow c, \mathtt{t} \rightarrow d\}$.
The strong constraints in $\mathcal{C}$ (that includes κ of Example 1) filter out from the (16 preferred extensions) of Λ' the solutions where more than two people attend the game, whereas the weak constraints maximize the set (or number) of people attending the game. □

Although in this paper we consider ground constraints, the proposed framework can be easily extended for general formulae with variables, whose ground version is a propositional set. For instance, the strong and weak constraints in Example 3 could be written using only one strong constraint $X \wedge Y \wedge Z \wedge (X \neq Y) \wedge (X \neq Z) \wedge (Y \neq Z) \Rightarrow \mathtt{f}$ and only one weak constraint $\mathtt{t} \rightarrow X$, where X, Y and Z are variables whose domain is the set of arguments.

Contributions. In this paper, after introducing CAFs and WAFs, we investigate the complexity of both credulous and skeptical reasoning in these frameworks, providing the results which are summarized in Table 1, where $CA_{\mathcal{S}}$ (resp., $SA_{\mathcal{S}}$) denotes the credulous (resp., skeptical) acceptance problem under argumentation semantics $\mathcal{S}$.

Specifically, we make the following main contributions:

- We propose CAFs relying on a simple yet expressive form of constraints that are interpreted using Lukasiewicz's logic, leading to an intuitive constraints' semantics.
- We investigate the complexity of $CA_{\mathcal{S}}$ and $SA_{\mathcal{S}}$ for CAFs under the most well-known semantics, showing that it remains the same as for AFs in most cases.
- We introduce WAFs and propose two criteria for interpreting weak constraints, under any argumentation semantics $\mathcal{S}$: *maximal-set* ($\mathtt{ms}\mathcal{S}$) and *maximum-cardinality* ($\mathtt{mc}\mathcal{S}$) according to which the best/optimal $\mathcal{S}$-extensions are those satisfying a maximal set, or a maximum number, of weak constraints, respectively.
- We investigate the complexity of $CA_{\mathcal{S}}$ and $SA_{\mathcal{S}}$ for WAFs, showing that differently from strong constraints the introduction of weak constraints typically increases the complexity of one level in the polynomial hierarchy.
- Moreover, we investigate restricted forms of WAFs, that is, Linearly WAFs (LWAFs) where constraints are linearly ordered, for which the complexity of $CA_{\mathcal{S}}$ and $SA_{\mathcal{S}}$ generally decreases, though it is higher than that of CAFs.
- Finally, we investigate the case of NCAFs and NWAFs, that is, CAFs and WAFs, respectively, where constraints are expressed by negative constraints (i.e., denials constraints whose body is a conjunction of literals, used in several contexts such as databases and logic programming) and show that the complexity of $CA_{\mathcal{S}}$ for the preferred semantics decreases.

Preliminaries

In this section, we briefly review Dung's framework and some basic notions about computational complexity.

Argumentation Frameworks

An abstract *Argumentation Framework* (AF) is a pair $\langle \mathcal{A}, \mathcal{R} \rangle$, where $\mathcal{A}$ is a set of *arguments* and $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ is a set of *attacks*. If $(a, b) \in \mathcal{R}$ then we say that a attacks b.

Given an AF $\Lambda = \langle \mathcal{A}, \mathcal{R} \rangle$ and a set $S \subseteq \mathcal{A}$ of arguments, an argument $a \in \mathcal{A}$ is said to be *i*) *defeated* w.r.t. S iff $\exists b \in S$ such that $(b, a) \in \mathcal{R}$, and *ii*) *acceptable* w.r.t. S iff for every argument $b \in \mathcal{A}$ with $(b, a) \in \mathcal{R}$, there is $c \in S$ such that $(c, b) \in \mathcal{R}$. The sets of defeated and acceptable arguments w.r.t. S are defined as follows (where Λ is understood):

- $Def(S) = \{a \in \mathcal{A} \mid \exists b \in S \,.\, (b, a) \in \mathcal{R}\}$;
- $Acc(S) = \{a \in \mathcal{A} \mid \forall b \in \mathcal{A} \,.\, (b, a) \notin \mathcal{R} \vee b \in Def(S)\}$.

Given an AF $\langle \mathcal{A}, \mathcal{R} \rangle$, a set $S \subseteq \mathcal{A}$ of arguments is said to be:

- *conflict-free* iff $S \cap Def(S) = \emptyset$;
- *admissible* iff it is conflict-free and $S \subseteq Acc(S)$.

Different argumentation semantics have been proposed to characterize collectively acceptable sets of arguments, called *extensions* (Dung 1995; Caminada 2006b). Every extension is an admissible set satisfying additional conditions. Specifically, the complete, preferred, stable, semi-stable, and grounded extensions of an AF are defined as follows.

Given an AF $\langle \mathcal{A}, \mathcal{R} \rangle$, a set $S \subseteq \mathcal{A}$ is an *extension* called:

- *complete* (co) iff it is an admissible set and $S = Acc(S)$;
- *preferred* (pr) iff it is a maximal (w.r.t. $\subseteq$) complete extension;
- *stable* (st) iff it is a total preferred extension, i.e. a preferred extension such that $S \cup Def(S) = \mathcal{A}$;
- *semi-stable* (sst) iff it is a preferred extension such that $S \cup Def(S)$ is maximal;
- *grounded* (gr) iff it is the smallest (w.r.t. $\subseteq$) complete extension.

It is well-known that the set of complete extensions forms a complete semilattice with respect to set inclusion. Arguments occurring in an extension are said to be accepted, whereas arguments attacked by accepted arguments are said to be rejected; remaining arguments are said to be undecided (w.r.t. the considered extension). Argumentation semantics can be also defined in terms of *labelling* (Caminada 2006a; Baroni, Caminada, and Giacomin 2011). A labelling for an AF $\langle \mathcal{A}, \mathcal{R} \rangle$ is a total function $L : \mathcal{A} \rightarrow \{\text{IN}, \text{OUT}, \text{UNDEC}\}$ assigning to each argument a label. $L(a) = \text{IN}$ (resp., OUT, UNDEC) means that argument a is accepted (resp., rejected, undecided). Thus, (complete) extensions can be also denoted as triples of sets of arguments containing accepted, rejected, and undecided arguments, respectively, as shown below.

Example 4. Let $\Lambda = \langle \mathcal{A}, \mathcal{R} \rangle$ be an AF where $\mathcal{A} = \{a, b, c\}$ and $\mathcal{R} = \{(a, b), (b, a), (b, c), (c, c)\}$. AF Λ has three complete extensions: $E_1 = \emptyset, E_2 = \{a\}, E_3 = \{b\}$, whose corresponding complete labellings are $L_1 = \langle \emptyset, \emptyset, \{a, b, c\} \rangle$, $L_2 = \langle \{a\}, \{b\}, \{c\} \rangle$, and $L_3 = \langle \{b\}, \{a, c\}, \emptyset \rangle$. Moreover, the set of preferred extensions is $\{E_2, E_3\}$, whereas the set of stable (and semi-stable) extensions is $\{E_3\}$, and the grounded extension is E_1. □

Given an AF-based framework Λ (e.g., AF, CAF, WAF), an argument a, and an argumentation semantics $\mathcal{S} \in \{$co, pr, st, sst, gr$\}$,

- the *credulous acceptance* problem, denoted as $CA_{\mathcal{S}}$, is the problem of deciding whether argument a is credulously accepted, that is, deciding whether a belongs to at least an $\mathcal{S}$-extension of Λ.
- the *skeptical acceptance* problem, denoted as $SA_{\mathcal{S}}$, is the problem of deciding whether argument a is skeptically accepted, that is, deciding whether a belongs to *every* $\mathcal{S}$-extension of Λ.

Clearly, for the grounded semantics, which admits exactly one extension, these problems become identical.

Complexity Classes. We recall here the main complexity classes used in the paper and, in particular, the definition of the classes Σ_k^P, Π_k^P and Δ_k^P, with $k \geq 0$ (see e.g. (Papadimitriou 1994)):

- $\Sigma_0^P = \Pi_0^P = \Delta_0^P = P$;
- $\Sigma_1^P = NP$ and $\Pi_1^P = coNP$;
- $\Delta_k^P = P^{\Sigma_{k-1}^P}$, $\Sigma_k^P = NP^{\Sigma_{k-1}^P}$, and $\Pi_k^P = co\Sigma_k^P$, $\forall k > 0$.

Thus, P^C (resp., NP^C) denotes the class of problems that can be solved in polynomial time using an oracle in the class C by a deterministic (resp., non-deterministic) Turing machine. The class $\Delta_k^P[log\ n]$ denotes the subclass of Δ_k^P containing the problems that can be solved in polynomial time by a deterministic Turing machine by performing a number of calls bounded by $O(log\ n)$ to an oracle in the class Σ_{k-1}^P. It is known that:

- $\Sigma_k^P \subset \Delta_{k+1}^P[log\ n] \subset \Delta_{k+1}^P \subset \Sigma_{k+1}^P \subseteq PSPACE$ and
- $\Pi_k^P \subset \Delta_{k+1}^P[log\ n] \subset \Delta_{k+1}^P \subset \Pi_{k+1}^P \subseteq PSPACE$.

For AFs, the complexity of the credulous and skeptical acceptance problems has been investigated in (Dung 1995) for the grounded semantics, in (Dimopoulos and Torres 1996) for the stable semantics, in (Dimopoulos and Torres 1996; Dunne and Bench-Capon 2002) for the preferred semantics, and in (Dunne and Caminada 2008; Dvorák and Woltran 2010) for the semi-stable semantics. The results for AFs are summarized in the second and third column of Table 1.

Constrained Argumentation Frameworks

We briefly review the Constrained Argumentation Framework (CAF) introduced in (Coste-Marquis, Devred, and Marquis 2006) and further investigated in (Arieli 2015).

We assume that, given a set of propositional symbols S, $\mathcal{L}_S$ denotes the propositional language defined in the usual inductive way from S using the built-in constants `f`, `u`, and `t` denoting the truth values `false`, `undef` (*undefined*), and `true`, and the connectives $\wedge$, $\vee$, $\neg$, $\Rightarrow$ and $\Leftrightarrow$.

Definition 1 (CAF). *A* Constrained Argumentation Framework (CAF) *is a triple $\Omega = \langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ where $\langle \mathcal{A}, \mathcal{R} \rangle$ is an AF and $\mathcal{C}$ is a set of propositional formulae built from $\mathcal{L}_A$.*

CAF Semantics

Given an AF $\langle \mathcal{A}, \mathcal{R} \rangle$ and a set $S \subseteq \mathcal{A}$, the truth value of an argument $a \in \mathcal{A}$ w.r.t. S is denoted by $\vartheta_S(a)$, or simply $\vartheta(a)$ whenever S is understood, and is defined as follows:

$$\vartheta(a) = \begin{cases} \texttt{true} & if\ a \in S \\ \texttt{false} & if\ \exists\, b \in S\ s.t.\ (b, a) \in \mathcal{R} \\ \texttt{undef} & otherwise \end{cases}$$

Observe that, for a given extension E whose corresponding labelling is L, $\vartheta_E(a)$ is `true` (resp., `false`, `undef`) iff $L(a) = \text{IN}$ (resp., OUT, UNDEC).

Assuming that $\neg\texttt{undef} = \texttt{undef}$, and the ordering on truth values $\texttt{false} < \texttt{undef} < \texttt{true}$, using the classical Kleene's 3-valued logic we have that $\vartheta(\varphi \wedge \psi) = min(\vartheta(\varphi), \vartheta(\psi))$, $\vartheta(\varphi \vee \psi) = max(\vartheta(\varphi), \vartheta(\psi))$, $\vartheta(\neg a) = \neg\vartheta(a)$ and $\vartheta(\varphi \Leftrightarrow \psi) = \vartheta(\phi \Rightarrow \psi) \wedge \vartheta(\psi \Rightarrow \varphi)$. It is important to note that, regarding the operator $\Rightarrow$, there is no convergence on its semantics in CAFs. In the following we discuss the proposals of (Coste-Marquis, Devred, and Marquis 2006) and (Arieli 2015); the former relying on classical 2-valued semantics, the latter relying on 3-valued semantics. With a little abuse of notation, in the rest of this section, we use the same symbol ϑ to denote the truth value of an argument even under different CAF semantics.

Coste-Marquis et al. Semantics The semantics proposed in (Coste-Marquis, Devred, and Marquis 2006) is based on the concept of *completion* of extensions. For any set of arguments $S \subseteq \mathcal{A}$, the completion of S is $\hat{S} = S \cup \{\neg a \mid a \in \mathcal{A} \setminus S\}$. Then, we say that S satisfies $\mathcal{C}$ if and only if $\hat{S}$ is a (2-valued) model of $\mathcal{C}$, denoted as $\hat{S} \models \mathcal{C}$, that is, $\vartheta_{\hat{S}}\left(\bigwedge_{\varphi \in \mathcal{C}} \varphi\right) = \texttt{true}$.

Definition 2 (C-admissible set). *Given a CAF $\Omega = \langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ and a set $S \subseteq \mathcal{A}$, S is a C-admissible set for Ω if and only if S is admissible for $\langle \mathcal{A}, \mathcal{R} \rangle$ and satisfies $\mathcal{C}$.*

A constrained argumentation framework $\Omega = \langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ is *consistent* when it has a C-admissible set for Ω.

Definition 3 (Preferred/Stable C-extension). *Let $\Omega = \langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ be a CAF. A C-admissible set $E \subseteq \mathcal{A}$ for Ω is*

- *a* preferred *C-extension of Ω if and only if $\nexists\, E' \subseteq \mathcal{A}$ such that $E \subset E'$ and E' is C-admissible for Ω;*
- *a* stable *C extension if and only if it is a total preferred C-extension.*

A drawback of the semantics proposed in (Coste-Marquis, Devred, and Marquis 2006) is that in checking whether an extension E satisfies a set of constraints it does not distinguish between false and undefined arguments. Thus, a constraint of the form $a \wedge \neg a \Rightarrow \texttt{f}$ is always satisfied, even when $\vartheta_E(a) = \texttt{undef}$. Indeed, the completion of E converts undefined truth values to negated truth values (e.g. assuming an extension $E = \emptyset$ stating that a is undefined, $\hat{E} = \{\neg a\}$) and then the 2-valued semantics is applied w.r.t. $\hat{E}$.

Arieli Semantics. The semantics proposed in (Arieli 2015) assumes the standard 3-valued Kleene's semantics for the connectives $\wedge$, $\vee$ and $\neg$, whereas for $\Rightarrow$ it assumes the Slupecki's interpretation which is defined as follows:

$$\vartheta(\varphi \Rightarrow \psi) = \begin{cases} \texttt{true} & if\ \vartheta(\varphi) \in \{\texttt{false}, \texttt{undef}\} \\ \vartheta(\psi) & otherwise \end{cases}$$

A natural requirement for constraints applied to argumentation frameworks is that they should have admissible interpretations: the constraints themselves should not be contradictory and every argument that is satisfied by their models should not be exposed to undefended attacks.

Definition 4 (Admissible constraint). *Let $\Lambda = \langle \mathcal{A}, \mathcal{R} \rangle$ be an AF. A set of formulae $\mathcal{C}$ in $\mathcal{L}_\mathcal{A}$ is called admissible (for Λ) if it has a 3-valued model which is an admissible set of Λ.*

Assuming that constraints are in $\mathcal{L}_\mathcal{A}$ and they are admissible, extensions for a CAF are defined as follows.

		$\vartheta(\psi)$			$\vartheta(\psi)$			$\vartheta(\psi)$		
		f	u	t	f	u	t	f	u	t
$\vartheta(\varphi)$	f	t	t	t	t	t	t	t	t	t
	u	u	t	t	t	t	t	u	u	t
	t	f	u	t	f	u	t	f	u	t
		Lukasiewicz			Slupecki			Kleene		

Figure 2: Semantics of the implication operator $\varphi \Rightarrow \psi$

Definition 5 ($\mathcal{S}$-extension of a CAF)**.** *Let $\Omega = \langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ be a CAF and let $\mathcal{S}$ be a semantics for $\langle \mathcal{A}, \mathcal{R} \rangle$. Then $E \subseteq \mathcal{A}$ is an $\mathcal{S}$-extension of Ω if it is an $\mathcal{S}$-extension of $\langle \mathcal{A}, \mathcal{R} \rangle$ and is a 3-valued model of $\mathcal{C}$.*

The main difference between the two CAF semantics reviewed in this section is as follows. In (Coste-Marquis, Devred, and Marquis 2006) the interpretations are determined by completion semantics and satisfiability of constraints is evaluated with respect to two-valued semantics. It follows, e.g., that a constraint of the form $\texttt{t} \Rightarrow a \vee \neg a$ is useless according to (Coste-Marquis, Devred, and Marquis 2006) (since it is always satisfied). In contrast, in the Arieli's 3-valued semantics this constraint indicates that argument a cannot have a neutral, undefined, status. The use of 3-valued semantics allows us to distinguish between different conditions on arguments. For instance, the constraint $\texttt{t} \Rightarrow \neg a$ means that a should be rejected, while the constraint $a \Rightarrow \texttt{f}$ is a somewhat weaker demand: a should not be accepted, and so its status may be undecided.

A drawback of Arieli's semantics, due to the assumption of the Slupecki's logic for interpreting the implication operator, is that it does not distinguish two constraints of the form $\varphi \Rightarrow \texttt{f}$ and $\varphi \Rightarrow \texttt{u}$, though it distinguishes two constraints of the form $\texttt{t} \Rightarrow \varphi$ and $\texttt{u} \Rightarrow \varphi$.

Revisiting the CAF Semantics

In order to avoid the problems discussed in the above section, we consider a simpler yet sufficiently general form of constraints and the classical interpretation of Lukasiewicz's logic for the implication operator. The tables in Figure 2 report three different semantics for the implication operator.

Let $\mathcal{L}'_{\mathcal{A}}$ be the propositional language defined from $\mathcal{A}$ and the connectives $\wedge$, $\vee$, $\neg$, where $\mathcal{A}$ is a set of arguments.

Definition 6 ((Strong) constraint)**.** *A* (strong) constraint *is a formula of one of the following forms: (i) $\varphi \Rightarrow v$, or (ii) $v \Rightarrow \varphi$, where φ is a propositional formula in $\mathcal{L}'_{\mathcal{A}}$ and $v \in \{\texttt{f}, \texttt{u}, \texttt{t}\}$. A constraint is said* boolean *when $v \in \{\texttt{f}, \texttt{t}\}$. A boolean constraint of the form $\varphi \Rightarrow \texttt{f}$ where φ is a conjunction containing arguments or negated arguments is called* denial *(or* negative*) constraint.*

We consider a less general form of constraints than that of Coste-Marquis et al. and Arieli (e.g. we do not deal with constraints with multiple implications); despite that, the expressive power of WAFs turns out to increase at least one level of the polynomial hierarchy w.r.t. that of AFs.

In our context, checking whether a constraint is satisfied (under the Lukasiewicz'logic) is equivalent to check whether the truth value of the head is greater than or equal to the truth value of the body.

In the following, we assume that $\mathcal{C}$ is a set of satisfiable constraints built from $\mathcal{L}'_{\mathcal{A}}$ as defined in Definition 6.

Observe that Kleene's logic interprets $\varphi \Rightarrow \psi$ as $\neg\varphi \vee \psi$, whereas Lucasiewicz'logic interprets the implication as $\neg\varphi \vee \psi \vee (\varphi \equiv \psi)$, i.e. the implication is satisfied whenever the truth value of the head is greater or equal than the truth value of the body. For boolean constraints Kleene and Lucasiewicz semantics coincide. A nice property of both Lucasiewicz and Kleene interpretations is that literals can be moved from the head to the body (after negating them), and vice versa, analogously to the case of 2-valued semantics. For formulae defining constraints we believe that Lucasiewicz interpretation is more appropriate as, for instance, it allows to distinguish $\varphi \Rightarrow \texttt{f}$ from $\varphi \Rightarrow \texttt{u}$, and avoids problems existing in other interpretations (Avron 1991).

Example 5. The constraint $a \wedge b \wedge c \Rightarrow \texttt{f}$ states that at least one of the arguments a, b and c must be false, whereas $a \wedge b \wedge c \Rightarrow \texttt{u}$ states that a, b and c cannot be all true. □

Clearly, constraints of the forms $\texttt{f} \Rightarrow \varphi$ and $\varphi \Rightarrow \texttt{t}$ are useless because always satisfied. Regarding the stable semantics, which is 2-valued, only the symbols f and t can be used and all interpretation of the implication operator coincide with the classical 2-valued interpretation.

Definition 7 ((Revised) CAF semantics)**.** *Given a CAF $\Omega = \langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ where $\mathcal{C}$ contains constraints built from $\mathcal{L}'_{\mathcal{A}}$, a set of arguments $S \subseteq \mathcal{A}$ is a complete (resp., grounded, preferred, stable, semi-stable) extension for Ω if S is a complete (resp., grounded, preferred, stable, semi-stable) extension for $\langle \mathcal{A}, \mathcal{R} \rangle$ and $S \models \mathcal{C}$.*

Note that, given a CAF $\Omega = \langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$, if we consider the corresponding AF $\Lambda = \langle \mathcal{A}, \mathcal{R} \rangle$, then the set of complete extensions of Λ that satisfy $\mathcal{C}$ does not always form a complete meet-semilattice. This is an important difference between CAFs and AFs, and it also holds for the CAF semantics reviewed in the previous section. Roughly speaking, the constraints break the lattice by marking as unfeasible some extensions. As a consequence, even the grounded extension is not guaranteed to exists, as shown below.

Example 6. Consider the CAF $\Omega = \langle \{a, b, c\}, \{(a, b), (b, a), (b, c), (c, c)\}, \{\texttt{t} \Rightarrow a \wedge b\} \rangle$ derived from the AF Λ of Example 4 by adding a strong constraint. As shown in Example 4, AF Λ has three complete extensions/labellings, $E_1 = \emptyset$, $E_2 = \{a\}$ and $E_3 = \{b\}$, but all extensions (resp., labellings) do not satisfy the constraint stating that both a and b must belong to them (resp., must be IN). Thus Ω has no complete extensions, and thus no grounded extension. □

Complexity of Credulous and Skeptical Acceptance

Although the presence of constraints in CAFs break the meet-semilattice of complete extensions, reasoning under the grounded semantics remains tractable.

Proposition 1. *The complexity of checking whether a CAF admits a grounded extension is in P.*

Therefore, since if a grounded extension for a CAF exists then it is unique, computing the credulous (or, equivalently, the skeptical) acceptance of an argument under the grounded semantics is still polynomial.

However, the credulous and skeptical acceptance of an argument w.r.t a CAF $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ may differ from that of the associated AF $\langle \mathcal{A}, \mathcal{R} \rangle$, independently of the semantics adopted, as shown in the following example.

Example 7. Continuing from Example 6, there are no arguments in the CAF that are credulously accepted under the complete semantics. In contrast, for the AF of Example 4, a (and also b) is credulously accepted under the complete, preferred and stable semantics, whereas b is skeptically accepted under the semi-stable semantics. □

The fact that the grounded extension may not exist for CAFs impacts on the complexity of the skeptical acceptance problem under complete semantics, which cannot be longer decided by simply looking at the grounded extension as for the case of AFs (where an argument is skeptically accepted under complete semantics if and only if it is in the grounded extension). Similarly, credulous acceptance under preferred semantics for CAFs can no longer be decided by checking credulous acceptance under complete semantics. In all the other cases we can show that the complexity of credulous and skeptical reasoning for CAFs and AFs coincides, as stated in the following theorem.

Theorem 1. *For any CAF $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$, the problem*

- *$CA_{\mathcal{S}}$ is: (i) NP-complete for any semantics $\mathcal{S} \in \{\texttt{co}, \texttt{st}\}$, (ii) NP-hard and in Σ_2^P for $\mathcal{S} = \texttt{pr}$, and (iii) Σ_2^P-complete for $\mathcal{S} = \texttt{sst}$.*
- *$SA_{\mathcal{S}}$ is: (i) coNP-complete for $\mathcal{S} \in \{\texttt{co}, \texttt{st}\}$, and (ii) Π_2^P-complete for $\mathcal{S} \in \{\texttt{pr}, \texttt{sst}\}$.*

Weak Constrained AFs

In this section, we introduce a generalization of CAFs where *weak* constraints are also considered. Differently from the strong constraints discussed in the previous section, weak constraints are propositional formulae that should be satisfied *if possible*. Specifically, weak constraints are logical formulae of the form $\varphi \rightarrow v$ (or, equivalently, $v \rightarrow \varphi$), where φ is a propositional formula built using the symbols of a given set $\mathcal{A}$ and the connectives $\wedge$, $\vee$ and $\neg$. Herein, $\rightarrow$ denotes the logical implication connective. Observe that we use the symbol $\rightarrow$ (instead of $\Rightarrow$) to have different syntaxes for weak and strong constraints.

Definition 8. *(Weak constrained AF) A* Weak constrained Argumentation Framework (WAF) *is a tuple $\Upsilon = \langle \mathcal{A}, \mathcal{R}, \mathcal{C}, \mathcal{W} \rangle$, where $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ is a CAF and $\mathcal{W}$ is a set of weak constraints built from $\mathcal{L}'_{\mathcal{A}}$.*

The semantics of a WAF is defined by considering two possible criteria for selecting the preferable extensions w.r.t. weak constraints—only weak constraints are considered when selecting the preferable extensions since strong constraints must be all satisfied. The two criteria considered for assessing to which extent an extension satisfies a set of weak constraints are: (*i*) *maximal set* criterion, considering as preferable (or "best") extensions the ones that satisfy a maximal set of weak constraints, and (*ii*) *maximum-cardinality* criterion, considering as preferable (or "optimal") extensions the ones that satisfy a maximal number of weak constraints. Clearly, the selection of preferable extensions make sense only for semantics admitting multiple extensions, that is, complete, preferred, stable, and semi-stable semantics. Thus, in the following, whenever we consider a generic semantics $\mathcal{S}$, we refer to $\mathcal{S} \in \{\texttt{co}, \texttt{pr}, \texttt{st}, \texttt{sst}\}$.

In the next subsections, after formally defining the meaning of a WAF under the maximal-set and maximum-cardinality semantics, we investigate the complexity of credulous and skeptical reasoning in the new framework.

Maximal-Set Semantics for WAFs

A WAF using the maximal-set criterion is defined as follows.

Definition 9 (Maximal-Set Semantics)**.** *Given a WAF $\Upsilon = \langle \mathcal{A}, \mathcal{R}, \mathcal{C}, \mathcal{W} \rangle$, an $\mathcal{S}$-extension E for $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ is a maximal-set $\mathcal{S}$-extension ($\texttt{ms}\mathcal{S}$-extension) for Υ if, let $\mathcal{W}_E \subseteq \mathcal{W}$ be the set of weak constraints that are satisfied by E (that is, $E \models \mathcal{W}_E$) , there is no $\mathcal{S}$-extension F for $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ and $\mathcal{W}_F \subseteq \mathcal{W}$ such that $F \models \mathcal{W}_F$ and $\mathcal{W}_E \subset \mathcal{W}_F$.*

Given a semantics $\mathcal{S}$, $\texttt{ms}\mathcal{S}$ denotes the maximal-set version of $\mathcal{S}$ (e.g. $\texttt{msco}$ denotes the $\texttt{ms}$ complete semantics).

Example 8. Consider the WAF $\Upsilon = \langle \mathcal{A}, \mathcal{R}, \mathcal{C}, \mathcal{W} \rangle$ with $\mathcal{A} = \{a, b, c, d\}$, $\mathcal{R} = \{(a, b), (b, a), (c, d), (d, c)\}$, $\mathcal{C} = \emptyset$ and $\mathcal{W} = \{w_1 = c \rightarrow \texttt{f}, w_2 = a \vee \neg a \rightarrow \texttt{u}\}$ stating that c should preferably be false (w_1) and a should preferably be undefined (w_2).

$\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ has 9 complete extensions:

$E_0 = \{\}$, $E_1 = \{a\}$, $E_2 = \{b\}$, $E_3 = \{c\}$, $E_4 = \{d\}$, $E_5 = \{a, c\}$, $E_6 = \{a, d\}$, $E_7 = \{b, c\}$ and $E_8 = \{b, d\}$.

In particular, E_0 is the grounded extension, whereas E_5, E_6, E_7, E_8 are preferred, stable, and semi-stable extensions of $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$. These are also extensions of AF $\langle \mathcal{A}, \mathcal{R} \rangle$, since $\mathcal{C} = \emptyset$.

Regarding the satisfaction of weak constraints, we have that $E_0 \models \{w_2\}$, $E_4 \models \{w_1, w_2\}$, $E_6 \models \{w_1\}$, and $E_8 \models \{w_1\}$, whereas the other complete extensions do not satisfy any constraint. Therefore, the maximal-set preferred (stable, semi-stable) extensions are E_6 and E_8, whereas there is only one maximal-set complete extension, which is E_4. □

Intuitively, given an $\mathcal{S}$-extension, checking satisfaction of a maximal-set of weak constraints means ensuring that no any other $\mathcal{S}$-extension is better according to that criterion. This is an additional source of complexity that makes, in most cases, credulous and skeptical reasoning in WAFs one level higher in the polynomial-time hierarchy than CAFs.

Theorem 2. *For any WAF $\langle \mathcal{A}, \mathcal{R}, \mathcal{C}, \mathcal{W} \rangle$, the problem*

- *$CA_{\texttt{ms}\mathcal{S}}$ is: (i) Σ_2^P-complete for any semantics $\mathcal{S} \in \{\texttt{co}, \texttt{st}\}$, (ii) Σ_2^P-hard and in Σ_3^P for $\mathcal{S} = \texttt{pr}$, and (iii) Σ_3^P-complete for $\mathcal{S} = \texttt{sst}$.*
- *$SA_{\texttt{ms}\mathcal{S}}$ is: (i) Π_2^P-complete for $\mathcal{S} \in \{\texttt{co}, \texttt{st}\}$, and (ii) Π_3^P-complete for $\mathcal{S} \in \{\texttt{pr}, \texttt{sst}\}$.*

Maximum-Cardinality Semantics

Maximum-cardinality semantics for WAFs prescribes as preferable extensions those satisfying the highest number of weak constraints. This is similar to the semantics of weak constraints in DLV (Alviano et al. 2017) where, in addition, each constraint has assigned a weight.

Definition 10 (Maximum-Cardinality Semantics)**.** *Given a WAF $\Upsilon = \langle \mathcal{A}, \mathcal{R}, \mathcal{C}, \mathcal{W} \rangle$, an $\mathcal{S}$-extension E for $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ is a maximum-cardinality $\mathcal{S}$-extension ($\mathtt{mc}\mathcal{S}$-extension) for Υ if, let $\mathcal{W}_E \subseteq \mathcal{W}$ be the set of weak constraints in $\mathcal{W}$ that are satisfied by E, there is no $\mathcal{S}$-extension F for $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ and $\mathcal{W}_F \subseteq \mathcal{W}$ such that $F \models \mathcal{W}_F$ and $|\mathcal{W}_E| < |\mathcal{W}_F|$.*

Example 9. Consider the WAF $\Upsilon = \langle \mathcal{A}, \mathcal{R}, \mathcal{C}, \mathcal{W} \rangle$ with $\mathcal{A} = \{a, b, c\}$, $\mathcal{R} = \{(a, b), (b, a), (b, c), (c, c)\}$, $\mathcal{C} = \emptyset$ and $\mathcal{W} = \{w_1 = \mathtt{t} \to a, w_2 = \mathtt{t} \to b, w_3 = c \to \mathtt{f}\}$ stating that it is desirable that a is true, b is true, and c is false.

$\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ has three complete extensions: $E_1 = \{\}$, $E_2 = \{a\}$, and $E_3 = \{b\}$, whose corresponding labellings are $L_1 = \langle \emptyset, \emptyset, \{a, b, c\} \rangle$, $L_2 = \langle \{a\}, \{b\}, \{c\} \rangle$, and $L_3 = \langle \{b\}, \{a, c\}, \emptyset \rangle$. Herein, E_2 and E_3 are the preferred extensions of $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$, whereas the unique stable (and semi-stable) extension is E_3. Regarding the satisfactions of weak constraints we have that $E_1 \models \mathcal{W}_0 = \emptyset$, $E_2 \models \mathcal{W}_1 = \{w_1\}$, and $E_3 \models \mathcal{W}_3 = \{w_2, w_3\}$. Therefore, the only maximum-cardinality preferred extension of Υ is E_3 (as $|\mathcal{W}_3| = 2 > |\mathcal{W}_1| = 1 > |\mathcal{W}_0| = 0$). Note that, according to the maximal-set semantics, both E_2 and E_3 are maximal-set preferred extensions. Regarding the stable (and semi-stable) semantics, as there is only one extension, E_3 is both a maximal-set and a maximum-cardinality extension. □

It turns out that, under standard complexity assumptions, computing credulous and skeptical acceptance in WAFs under maximum-cardinality semantics is easier than using maximal-set semantics. Roughly speaking, this follows from the fact that a binary search strategy can be used for deciding whether the cardinality of a set of constraints satisfied by an $\mathcal{S}$-extension containing a given argument is maximum.

Theorem 3. *For any WAF $\langle \mathcal{A}, \mathcal{R}, \mathcal{C}, \mathcal{W} \rangle$ with $|\mathcal{W}| = n$, the problems $CA_{\mathtt{mc}\mathcal{S}}$ and $SA_{\mathtt{mc}\mathcal{S}}$ are:*

- $\Delta_2^P[\log n]$*-complete for any semantics* $\mathcal{S} \in \{\mathtt{co}, \mathtt{st}\}$,
- $\Delta_3^P[\log n]$*-complete for* $\mathcal{S} \in \{\mathtt{pr}, \mathtt{sst}\}$.

Stratified Weak Constrained AFs

In this section we consider WAFs where weak constraints define a partial order. The partial order on the set $\mathcal{W}$ of weak constraints is defined by partitioning $\mathcal{W}$ into strata $\mathcal{W}_1, \ldots, \mathcal{W}_n$ (with $n \geq 1$) so that weak constraints in a stratum i have higher priority than those in a stratum $j > i$.

Definition 11. *(Stratified WAF) A* Stratified Weak constrained Argumentation Framework (SWAF) *is a tuple $\langle A, \mathcal{R}, \mathcal{C}, \mathcal{W} \rangle$ where $\langle A, \mathcal{R}, \mathcal{C} \rangle$ is a CAF and $\mathcal{W}$ is a list of sets of weak constraints $(\mathcal{W}_1, \ldots, \mathcal{W}_n)$ built from $\mathcal{L}'_{\mathcal{A}}$.*

Note that whenever $n = 1$ we have a unique stratum and, then, SWAFs coincide with standard WAFs, which in turns implies that the complexity of SWAFs coincides with that of WAFs in the worst case.

Regarding the semantics of a SWAF $\langle A, \mathcal{R}, \mathcal{C}, \mathcal{W} \rangle$, the underlying idea is that weak constraints are applied one stratum at a time. Therefore, given a set S of $\mathcal{S}$-extensions of $\langle A, \mathcal{R}, \mathcal{C} \rangle$, the best/optimal $\mathcal{S}$-extensions are obtained by first computing the set $S_1 \subseteq S$ which are best/optimal solutions w.r.t. $\mathcal{W}_1$, then the set $S_2 \subseteq S_1$ of $\mathcal{S}$-extensions which are best/optimal solutions w.r.t. $\mathcal{W}_2$ is selected, and so on.

Definition 12 (SWAF Semantics)**.** *Let $\Upsilon = \langle \mathcal{A}, \mathcal{R}, \mathcal{C}, (\mathcal{W}_1, \ldots, \mathcal{W}_n) \rangle$ be a SWAF and $\mathcal{S}$ a semantics. An $\mathcal{S}$-extension E for $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ is an $\mathtt{ms}\mathcal{S}$-extension ($\mathtt{mc}\mathcal{S}$-extension) for Υ if:*

- *$n = 1$ and E is an $\mathtt{ms}\mathcal{S}$-extension (resp., $\mathtt{mc}\mathcal{S}$-extension) for $\langle \mathcal{A}, \mathcal{R}, \mathcal{C}, \mathcal{W}_1 \rangle$;*
- *$n > 1$, E is an $\mathtt{ms}\mathcal{S}$-extension (resp., $\mathtt{mc}\mathcal{S}$-extension) for $\langle \mathcal{A}, \mathcal{R}, \mathcal{C}, (\mathcal{W}_1, ..., \mathcal{W}_{n-1}) \rangle$ and E satisfies a maximal set of constraints in $\mathcal{W}_n$ (resp., the maximum number of constraints in $\mathcal{W}_n$).*

Example 10. Consider the SWAF derived from the AF Λ of Example 1 by adding the following list of sets of weak constraints $(\{a \wedge b \wedge c \to \mathtt{f}\}, \{\mathtt{t} \to a\}, \{\mathtt{t} \to b\}, \{\mathtt{t} \to c\})$. Since each stratum contains only one weak constraint, maximal-set and maximum-cardinality semantics give the same result. The weak constraints are applied one (set) at a time to discard extensions. After applying the first constraint the extension containing a, b and c is discarded. At the second step only extensions containing a are selected from the ones selected at the first step. At the third step only the extension containing a and b is selected. Thus, the best/optimal extension is the one containing exactly a and b.

Note that, assuming that weak constraints are not stratified, we would have the three extensions $\{a, b\}$, $\{a, c\}$ and $\{b, c\}$ under both maximal-set and maximum-cardinality preferred and stable semantics □

A particular form of SWAFs are the ones used in the above example, where every stratum is a singleton, meaning that weak constraints define a total order.

Definition 13 (LWAF)**.** *A SWAF $\langle \mathcal{A}, \mathcal{R}, \mathcal{C}, (\mathcal{W}_1, ..., \mathcal{W}_n) \rangle$ is said to be a linearly ordered if every $\mathcal{W}_i$ ($1 \leq i \leq n$) contains only one weak constraint.*

Observe that for linearly ordered SWAFs, $CA_{\mathtt{ms}\mathcal{S}} = CA_{\mathtt{mc}\mathcal{S}}$ and $SA_{\mathtt{ms}\mathcal{S}} = SA_{\mathtt{mc}\mathcal{S}}$. Thus, for this class of constrained AFs, we simply use the notations $CA_{\mathcal{S}}$ and $SA_{\mathcal{S}}$.

Theorem 4. *For any linearly ordered SWAF $\langle \mathcal{A}, \mathcal{R}, \mathcal{C}, (\mathcal{W}_1, \ldots, \mathcal{W}_n) \rangle$, the problems $CA_{\mathcal{S}}$ and $SA_{\mathcal{S}}$ are:*

- Δ_2^P*-complete for any semantics* $\mathcal{S} \in \{\mathtt{co}, \mathtt{st}\}$,
- Δ_3^P*-complete for* $\mathcal{S} \in \{\mathtt{pr}, \mathtt{sst}\}$.

CAFs and WAFs with Denial Constraints

In several contexts (e.g. database, logic programming, inconsistent knowledge management) constraints are expressed by denial constraints. In this section we investigate credulous and skeptical acceptance when constraints are expressed by negative constraints only. In the following, with the acronym NCAF (resp., NWAF) we denote any CAF (resp., WAF) where weak and strong constraints are defined by denials.

Example 11. The WAF of Example 3 is an NWAF, since the constraints in $\mathcal{C}$ are denials and those in $\mathcal{W}$ can be equivalently written as the denials $\neg a \Rightarrow \mathtt{f}$, $\neg b \Rightarrow \mathtt{f}$, $\neg c \Rightarrow \mathtt{f}$, and $\neg d \Rightarrow \mathtt{f}$. Moreover, if $\mathcal{W} = \emptyset$ then we obtain an NCAF. □

The following theorem shows that, if only denial constraints are used, the complexity of credulous acceptance under preferred semantics does not increase for NCAFs w.r.t. AFs; moreover, the complexity of skeptical acceptance for NCAFs does not change w.r.t CAFs.

Theorem 5. *For any NCAF* $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$*, the problem*

- $CA_{\mathcal{S}}$ *is:* (i) *NP-complete for any semantics* $\mathcal{S} \in \{\texttt{co}, \texttt{st}, \texttt{pr}\}$*, and* (ii) Σ_2^P*-complete for* $\mathcal{S} = \texttt{sst}$.
- $SA_{\mathcal{S}}$ *is:* (i) *coNP-complete for* $\mathcal{S} \in \{\texttt{co}, \texttt{st}\}$*, and* (ii) Π_2^P*-complete for* $\mathcal{S} \in \{\texttt{pr}, \texttt{sst}\}$.

As stated next, the introduction of weak denial constraints increases the complexity of one level in the polynomial hierarchy w.r.t. that of NCAFs, for each semantics.

Theorem 6. *For any NWAF* $\langle \mathcal{A}, \mathcal{R}, \mathcal{C}, \mathcal{W} \rangle$*, the problem*

- $CA_{\texttt{ms}\mathcal{S}}$ *is:* (i) Σ_2^P*-complete for any semantics* $\mathcal{S} \in \{\texttt{co}, \texttt{st}, \texttt{pr}\}$*, and* (ii) Σ_3^P*-complete for* $\mathcal{S} = \texttt{sst}$.
- $SA_{\texttt{ms}\mathcal{S}}$ *is:* (i) Π_2^P*-complete for* $\mathcal{S} \in \{\texttt{co}, \texttt{st}\}$*, and* (ii) Π_3^P*-complete for* $\mathcal{S} \in \{\texttt{pr}, \texttt{sst}\}$.

Related Work

Besides being related to the proposals for CAFs in (Coste-Marquis, Devred, and Marquis 2006; Arieli 2015), that we discussed in the paper, our work is also related to the approach in (Booth et al. 2013) that provides a method for generating non-empty conflict-free extensions for CAFs. Constraints have been also used in the context of dynamic AFs to refer to the enforcement of some change (Doutre and Mailly 2018). In this context, extension enforcement aims at modifying an AF to ensure that a given set of arguments becomes (part of) an extension for the chosen semantics (Baumann and Brewka 2010; Baumann 2012; Coste-Marquis et al. 2015; Wallner, Niskanen, and Järvisalo 2017; Niskanen, Wallner, and Järvisalo 2018). This is different from our approach where integrity constrains are used to discard unfeasible solutions (extensions), without enforcing that a new set of arguments becomes an extension.

Weak constraints allow for selecting "best" or "optimal" extensions satisfying some conditions on arguments, if possible. This can be viewed as expressing a kind of preference over the set of extensions. Dung's framework has been extended in several ways for allowing preferences over arguments (Amgoud and Cayrol 2002; Kaci and van der Torre 2008; Modgil 2009; Amgoud and Vesic 2011). In particular, preferences relying to so-called critical attacks, i.e., attacks from a less preferred argument to a more preferred one, can be handled by removing or invalidating such attacks or by "inverting" them (Amgoud and Vesic 2014). Such kind of preferences can be encoded into AFs, possible through reductions relying on additional (meta)-arguments and attacks (Kaci, van der Torre, and Villata 2018). Thus these preferences do not increase expressiveness of AFs from a computational standpoint.

Preferences can be also expressed in value-based AFs (Bench-Capon 2003; Dunne and Bench-Capon 2004), where each argument is associated with a numeric value, and a set of possible orders (preferences) among the values is defined. In (Dunne et al. 2011) weights are associated with attacks, and new semantics extending the classical ones on the basis of a given numerical threshold are proposed. (Coste-Marquis et al. 2012) extends (Dunne et al. 2011) by considering other aggregation functions over weights apart from sum. Except for weighted solutions under grounded semantics (that prescribes more than one weighted solution), the complexity of credulous and skeptical reasoning in the above-considered AF-based frameworks is lower than that of WAFs, which suggests that WAFs are more expressive and can be used to model those frameworks (we plan to formally investigate these connections in future work).

Conclusions and Future Work

We have introduced a general argumentation framework where both strong and weak constraints can be easily expressed. Our complexity analysis shows how the several forms of constrains (including restricted forms, e.g. denials) impact on the complexity of credulous and skeptical reasoning. It turns out that constraints, especially weak ones, generally increase the expressivity of AFs. In fact, WAFs allow us to model optimization problems such as for instance Min Coloring and Maximum Satisfiability, where some kind of preferences (e.g. use the minimum number of colors) are expressed on solutions. This is not possible for AFs/CAFs whose expressivity is lower than that of WAFs (AFs/CAFs can capture simpler problems such as k-coloring and SAT).

We envisage implementations of the proposed WAF semantics by exploiting ASP-based systems and analogies with logic programs with weak constraints (Buccafurri, Leone, and Rullo 2000; Greco 1998) (the relationship between the semantics of some frameworks extending AF and that of logic programs has been recently investigated in (Alfano et al. 2020c,d)). For WAFs, DLV system (Alviano et al. 2017) could be used for computing maximum-cardinality stable semantics.

Future work will be also devoted to considering more general forms of constraints, not only using variables ranging on the sets of arguments, but also constraints allowing to express conditions on aggregates (e.g., at least n arguments from a given set S should be accepted/rejected). The basic idea of adding weak constraints could be also applicable for structured argumentation formalisms (Bondarenko et al. 1997; Garcia, Prakken, and Simari 2020).

Finally, given the inherent nature of argumentation and the typical high computational complexity of most of the reasoning tasks (Alfano et al. 2020a), there have been several recent efforts toward the investigation of incremental techniques that use AF solutions (e.g. extensions, skeptical acceptance) at time t to recompute updated solutions at time $t + 1$ after that an update (e.g. adding/ removing an attack) is performed (Greco and Parisi 2016a,b; Alfano, Greco, and Parisi 2017, 2019; Alfano and Greco 2021; Doutre and Mailly 2018). These approaches have been extended to argumentation frameworks more general than AFs (Alfano, Greco, and Parisi 2020; Alfano et al. 2020b, 2018a,b). Following this line of research, we plan to investigate incremental techniques for recomputing CAF and WAF semantics after performing updates consisting of changes to the AF component or to the sets of strong and weak constraints.

References

Alfano, G.; Calautti, M.; Greco, S.; Parisi, F.; and Trubitsyna, I. 2020a. Explainable Acceptance in Probabilistic Abstract Argumentation: Complexity and Approximation. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR*, 33–43.

Alfano, G.; Cohen, A.; Gottifredi, S.; Greco, S.; Parisi, F.; and Simari, G. R. 2020b. Dynamics in Abstract Argumentation Frameworks with Recursive Attack and Support Relations. In *Proc. of the 24th European Conference on Artificial Intelligence (ECAI)*, 577–584.

Alfano, G.; and Greco, S. 2021. Incremental Skeptical Preferred Acceptance in Dynamic Argumentation Frameworks. *IEEE Intelligent Systems* doi:10.1109/MIS.2021.3050521.

Alfano, G.; Greco, S.; and Parisi, F. 2017. Efficient Computation of Extensions for Dynamic Abstract Argumentation Frameworks: An Incremental Approach. In *Proc. of the Twenty-Sixth International Joint Conference on Artificial Intelligence, (IJCAI)*, 49–55.

Alfano, G.; Greco, S.; and Parisi, F. 2019. An Efficient Algorithm for Skeptical Preferred Acceptance in Dynamic Argumentation Frameworks. In *Proc. of the Twenty-Eighth International Joint Conference on Artificial Intelligence, (IJCAI)*, 18–24.

Alfano, G.; Greco, S.; and Parisi, F. 2020. Computing Skeptical Preferred Acceptance in Dynamic Argumentation Frameworks with Recursive Attack and Support Relations. In *In Proc. of International Conference on Computational Models of Argument (COMMA)*, 67–78.

Alfano, G.; Greco, S.; Parisi, F.; Simari, G. I.; and Simari, G. R. 2018a. An Incremental Approach to Structured Argumentation over Dynamic Knowledge Bases. In *Principles of Knowledge Representation and Reasoning: Proc. of the Sixteenth International Conference (KR)*, 78–87.

Alfano, G.; Greco, S.; Parisi, F.; Simari, G. I.; and Simari, G. R. 2018b. Incremental computation of warranted arguments in dynamic defeasible argumentation: the rule addition case. In *Proc. of the 33rd Annual ACM Symposium on Applied Computing, SAC*, 911–917.

Alfano, G.; Greco, S.; Parisi, F.; and Trubitsyna, I. 2020c. On the Semantics of Abstract Argumentation Frameworks: A Logic Programming Approach. *Theory Pract. Log. Program.* 20(5): 703–718.

Alfano, G.; Greco, S.; Parisi, F.; and Trubitsyna, I. 2020d. On the Semantics of Recursive Bipolar AFs and Partial Stable Models. In *Proceedings of the Workshop on Advances In Argumentation In Artificial Intelligence 2020 co-located with the 19th International Conference of the Italian Association for Artificial Intelligence (AIxIA 2020), Online, November 25-26, 2020*, volume 2777 of *CEUR Workshop Proceedings*, 16–30. CEUR-WS.org.

Alviano, M.; Calimeri, F.; Dodaro, C.; Fuscà, D.; Leone, N.; Perri, S.; Ricca, F.; Veltri, P.; and Zangari, J. 2017. The ASP System DLV2. In *Proc. of 14th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR)*, 215–221.

Amgoud, L.; and Cayrol, C. 1998. On the Acceptability of Arguments in Preference-based Argumentation. In *Proc. of the Fourteenth Conference on Uncertainty in Artificial Intelligence (UAI)*, 1–7.

Amgoud, L.; and Cayrol, C. 2002. A Reasoning Model Based on the Production of Acceptable Arguments. *Ann. Math. Artif. Intell.* 34(1-3): 197–215.

Amgoud, L.; and Vesic, S. 2011. A new approach for preference-based argumentation frameworks. *Ann. Math. Artif. Intell.* 63(2): 149–183.

Amgoud, L.; and Vesic, S. 2014. Rich preference-based argumentation frameworks. *Int. J. Approx. Reason.* 55(2): 585–606.

Arieli, O. 2013. Towards Constraints Handling by Conflict Tolerance in Abstract Argumentation Frameworks. In *Proc. of the Twenty-Sixth International Florida Artificial Intelligence Research Society Conference (FLAIRS)*.

Arieli, O. 2015. Conflict-free and conflict-tolerant semantics for constrained argumentation frameworks. *J. Appl. Log.* 13(4): 582–604.

Arieli, O. 2016. On the acceptance of loops in argumentation frameworks. *J. Log. Comput.* 26(4): 1203–1234.

Atkinson et al., K. 2017. Towards Artificial Argumentation. *AI Magazine* 38(3): 25–36.

Avron, A. 1991. Natural 3-valued logics–characterization and proof theory. *The Journal of Symbolic Logic* 56(1): 276–294.

Baroni, P.; Caminada, M.; and Giacomin, M. 2011. An introduction to argumentation semantics. *The Knowledge Engineering Review* 26(4): 365–410.

Baumann, R. 2012. What Does it Take to Enforce an Argument? Minimal Change in abstract Argumentation. In *Proc. of the 20th European Conference on Artificial Intelligence (ECAI)*, 127–132.

Baumann, R.; and Brewka, G. 2010. Expanding Argumentation Frameworks: Enforcing and Monotonicity Results. In *Proc. of International Conference on Computational Models of Argument (COMMA)*, 75–86.

Bench-Capon, T.; and Dunne, P. E. 2007. Argumentation in Artificial Intelligence. *Artif. Intell.* 171: 619 – 641.

Bench-Capon, T. J. M. 2003. Persuasion in Practical Argument Using Value-based Argumentation Frameworks. *J. Log. Comput.* 13(3): 429–448.

Bondarenko, A.; Dung, P. M.; Kowalski, R. A.; and Toni, F. 1997. An Abstract, Argumentation-Theoretic Approach to Default Reasoning. *AI* 93: 63–101.

Booth, R.; Kaci, S.; Rienstra, T.; and van der Torre, L. W. N. 2013. A Logical Theory about Dynamics in Abstract Argumentation. In *Proc. of International Conference Scalable Uncertainty Management (SUM)*, 148–161.

Buccafurri, F.; Leone, N.; and Rullo, P. 2000. Enhancing Disjunctive Datalog by Constraints. *IEEE Trans. Knowl. Data Eng.* 12(5): 845–860.

Caminada, M. 2006a. On the Issue of Reinstatement in Argumentation. In *Proc. of 10th European Conference on Logics in Artificial Intelligence (JELIA)*, 111–123.

Caminada, M. 2006b. Semi-Stable Semantics. In *Proc. of Computational Models of Argument (COMMA)*, 121–130.

Coste-Marquis, S.; Devred, C.; and Marquis, P. 2006. Constrained Argumentation Frameworks. In *Proc. of the Tenth International Conference on Principles of Knowledge Representation and Reasoning (KR)*, 112–122.

Coste-Marquis, S.; Konieczny, S.; Mailly, J.; and Marquis, P. 2015. Extension Enforcement in Abstract Argumentation as an Optimization Problem. In *Proc. of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI)*, 2876–2882.

Coste-Marquis, S.; Konieczny, S.; Marquis, P.; and Ouali, M. A. 2012. Weighted Attacks in Argumentation Frameworks. In *Proc. of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning (KR)*.

Dimopoulos, Y.; and Torres, A. 1996. Graph Theoretical Structures in Logic Programs and Default Theories. *Theor. Comput. Sci.* 170(1-2): 209–244.

Doutre, S.; and Mailly, J. 2018. Constraints and changes: A survey of abstract argumentation dynamics. *Argument & Computation* 9(3): 223–248.

Dung, P. M. 1995. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. *AI* 77: 321–358.

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

Dunne, P. E.; and Bench-Capon, T. J. M. 2004. Complexity in Value-Based Argument Systems. In *Proc. of 9th European Conference on Logics in Artificial Intelligence (JELIA)*, 360–371.

Dunne, P. E.; and Caminada, M. 2008. Computational Complexity of Semi-stable Semantics in Abstract Argumentation Frameworks. In *Proc. of 11th European Conference on Logics in Artificial Intelligence (JELIA)*, volume 5293, 153–165.

Dunne, P. E.; Hunter, A.; McBurney, P.; Parsons, S.; and Wooldridge, M. J. 2011. Weighted argument systems: Basic definitions, algorithms, and complexity results. *Artif. Intell.* 175(2): 457–486.

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

Faber, W.; Vallati, M.; Cerutti, F.; and Giacomin, M. 2016. Solving Set Optimization Problems by Cardinality Optimization with an Application to Argumentation. In *Proc. of 22nd European Conference on Artificial Intelligence (ECAI)*, 966–973.

Garcia, A. J.; Prakken, H.; and Simari, G. R. 2020. A Comparative Study of Some Central Notions of ASPIC+ and DeLP. *Theory Pract. Log. Program.* 20(3): 358–390.

Greco, S. 1998. Non-Determinism and Weak Constraints in Datalog. *New Gener. Comput.* 16(4): 373–396.

Greco, S.; and Parisi, F. 2016a. Efficient Computation of Deterministic Extensions for Dynamic Abstract Argumentation Frameworks. In *Proc. of 22nd European Conference on Artificial Intelligence (ECAI)*, 1668–1669.

Greco, S.; and Parisi, F. 2016b. Incremental Computation of Deterministic Extensions for Dynamic Argumentation Frameworks. In *Proc. of 15th European Conference on Logics in Artificial Intelligence (JELIA)*, 288–304.

Kaci, S.; and van der Torre, L. W. N. 2008. Preference-based argumentation: Arguments supporting multiple values. *Int. J. Approx. Reason.* 48(3): 730–751.

Kaci, S.; van der Torre, L. W. N.; and Villata, S. 2018. Preference in Abstract Argumentation. In *In Proc. of International Conference on Computational Models of Argument (COMMA)*, 405–412.

Li, M.; and Lv, R. 2014. The Binary Algorithm for finding the best weak-constraints in contradictory weak-constrained optimization problems. In *Proc. of 10th International Conference on Natural Computation (ICNC)*, 475–479.

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

Modgil, S.; and Prakken, H. 2013. A general account of argumentation with preferences. *Artif. Intell.* 195: 361–397.

Niskanen, A.; Wallner, J. P.; and Järvisalo, M. 2018. Extension Enforcement under Grounded Semantics in Abstract Argumentation. In *Proc. of the Sixteenth International Conference on Principles of Knowledge Representation and Reasoning (KR)*, 178–183.

Papadimitriou, C. H. 1994. *Computational complexity.* Addison-Wesley.

Ramos, A. 2019. Two New Weak Constraint Qualifications for Mathematical Programs with Equilibrium Constraints and Applications. *J. Optim. Theory Appl.* 183(2): 566–591.

Simari, G. R.; and Rahwan, I., eds. 2009. *Argumentation in Artificial Intelligence*. Springer.

Wallner, J. P.; Niskanen, A.; and Järvisalo, M. 2017. Complexity Results and Algorithms for Extension Enforcement in Abstract Argumentation. *J. Artif. Intell. Res.* 60: 1–40.