The Complexity of Repairing, Adjusting, and Aggregating of Extensions in Abstract Argumentation

Eun Jung Kim[1], Sebastian Ordyniak⋆[2], and Stefan Szeider⋆⋆[2]

[1] LAMSADE, CNRS, Paris, France, eunjungkim78@gmail.com
[2] Faculty of Informatics, Masaryk University, Brno, Czech Republic
sordyniak@gmail.com
[3] Institute of Information Systems, Vienna University of Technology, A-1040 Vienna, Austria
stefan@szeider.net

Abstract. We study the computational complexity of problems that arise in abstract argumentation in the context of dynamic argumentation, minimal change, and aggregation. In particular, we consider the following problems where always an argumentation framework F and a small positive integer k are given.

- The REPAIR problem asks whether a given set of arguments can be modified into an extension by at most k elementary changes (i.e., the extension is of distance k from the given set).
- The ADJUST problem asks whether a given extension can be modified by at most k elementary changes into an extension that contains a specified argument.
- The CENTER problem asks whether, given two extensions of distance k, whether there is a "center" extension that is a distance at most $k-1$ from both given extensions.

We study these problems in the framework of parameterized complexity, and take the distance k as the parameter. Our results covers several different semantics, including admissible, complete, preferred, semi-stable and stable semantics.

1 Introduction

Starting with the seminal work by Dung [10] the area of argumentation has evolved to one of the most active research branches within Artificial Intelligence [4, 26]. Dung's abstract argumentation frameworks, where arguments are seen as abstract entities which are just investigated with respect to how they relate to each other, in terms of "attacks", are nowadays well understood and different semantics (i.e., the selection of sets of arguments which are jointly acceptable) have been proposed. Such sets of arguments are called extensions of the underlying argumentation framework (AF).

Argumentation is an inherently dynamic process, and there has been increasingly interest in the dynamic behavior of abstract argumentation. A first study in this direction was carried out by Cayrol, et al. [6] and was concerned with the impact of additional arguments on extensions. Baumann and Brewka [3] investigated whether it is possible to

⋆ Research supported by Employment of Newly Graduated Doctors of Science for Scientific Excellence (CZ.1.07/2.3.00/30.0009).
⋆⋆ Research supported by the European Research Council, grant reference 239962.

modify a given AF in such a way that a desired set of arguments becomes an extension or a subset of an extension. Baumann [2] further extended this line of research by considering the minimal exchange necessary to enforce a desired set of arguments. In this context, it is interesting to consider notions of *distance* between extensions. Booth et al. [5] suggested a general framework for defining and studying distance measures.

A natural question that arises in the context of abstract argumentation is how computationally difficult it is to decide whether an AF admits an extension at all, or whether a given argument belongs to at least one extension or to all extensions of the AF. Indeed this question has been investigated in a series of papers, and the exact worst-case complexities have been determined for all popular semantics [7, 8, 10, 12–14, 18]. Abstract argumentation has also been studied in the framework of *parameterized complexity* [9] which admits a more fine-grained complexity analysis that can take structural aspects of the AF into account [11, 15, 22, 19, 16].

Surprisingly, very little is known on the computational complexity of problems in abstract argumentation that arise in the context of dynamic behavior of argumentation, such as finding an extension by minimal change. However, as the distance in these problems are assumed to be small, it suggests itself to consider the distance as the parameter for a parameterized analysis.

New Contribution In this paper we provide a detailed complexity map of various problems that arise in abstract argumentation in the context of change and distance.

In particular, we consider the following problems where always an argumentation framework F and a small positive integer k are given, and σ denotes a semantics.

- The σ-REPAIR problem asks whether a given set of arguments can be modified into a σ-extension by at most k elementary changes (i.e., the extension is of distance k from the given set).
  This problem is of relevance, for instance, when a σ-extension E of an argumentation framework is given, and dynamically the argumentation framework changes (i.e., attacks are added or removed, new arguments are added). Now the set E may not any more be a σ-extension of the new framework, and we want to repair it with minimal change to obtain a σ-extension.
- The σ-ADJUST problem asks whether a given σ-extension can be modified by at most k elementary changes into a σ-extension that contains a specified argument.
  This problem is a variant of the previous problem, however, the argumentation framework does not change, but dynamically the necessity occurs to include a certain argument into the extension, by changing the given extension minimally.
- The σ-CENTER problem asks whether, given two σ-extensions of distance k, whether there is a "center" σ-extension that is a distance at most $k-1$ from both given extensions.
  This problem arises in scenarios of judgment aggregations, when, for instance two extensions that reflect the opinion of two different agents are presented, and one tries to find a compromise extension that minimizes the distance to both extensions.

We study these problems in the framework of parameterized complexity, and take the distance k as the parameter. Our results covers several different semantics, including admissible, complete, preferred, semi-stable and stable semantics. The parameterized complexity of the above problems are summarized in Figures 1.

σ	general	bounded degree
adm	W[1]-hard	FPT
com	W[1]-hard	FPT
prf	para-coNP-hard	para-coNP-hard
sem	para-coNP-hard	para-coNP-hard
stb	W[1]-hard	FPT

Fig. 1. Parameterized Complexity of the problems σ-REPAIR, σ-ADJUST, and σ-CENTER for arbitrary AFs and AFs of bounded degree depending on the considered semantics.

2 Preliminaries

An *abstract argumentation system* or *argumentation framework* (*AF*, for short) is a pair (X, A) where X is a (possible infinite) set of elements called *arguments* and $A \subseteq X \times X$ is a binary relation called *attack relation*. In this paper we will restrict ourselves to finite AFs, i.e., to AFs for which X is a finite set. If $(x, y) \in A$ we say that x *attacks* y and that x is an *attacker* of y.

An AF $F = (X, A)$ can be considered as a directed graph, and therefore it is convenient to borrow notions and notation from graph theory. For a set of arguments $Y \subseteq X$ we denote by $F[Y]$ the AF $(Y, \{ (x, y) \in A \mid x, y \in Y \})$ and by $F - Y$ the AF $F[X \setminus Y]$.

We define $\overline{F}$ to be the undirected graph obtained from F that has vertex set X and edge set $\{ \{x, y\} \}(x, y) \in A \}$. We define the degree of an argument $x \in X$ to be the degree (number of neighbors) of the vertex that corresponds to x in $\overline{F}$. We say a class of AFs $\mathcal{C}$ has bounded maximum degree, or shortly bounded degree, if there exists a constant c such that for every $F \in \mathcal{C}$ the maximum degree of the undirected graph $\overline{F}$ is at most c.

If E and E' are 2 sets of arguments of F then we define $E \triangle E'$ to be the symmetric difference between E and E', i.e., $E \triangle E' := \{ x \in X \mid (x \in E \wedge x \notin E') \vee (x \in E' \wedge x \notin E) \}$. We also define $\text{dist}(E, E')$ to be $|E \triangle E'|$.

Let $F = (X, A)$ be an AF, $S \subseteq X$ and $x \in X$. We say that x is *defended* (in F) by S if for each $x' \in X$ such that $(x', x) \in A$ there is an $x'' \in S$ such that $(x'', x') \in A$. We denote by S^+_F the set of arguments $x \in X$ such that either $x \in S$ or there is an $x' \in S$ with $(x', x) \in A$, and we omit the subscript if F is clear from the context. Note that in our setting the set S is contained in S^+_F. We say S is *conflict-free* if there are no arguments $x, x' \in S$ with $(x, x') \in A$.

Next we define commonly used semantics of AFs, see the survey of Baroni and Giacomin [1]. We consider a semantics σ as a mapping that assigns to each AF $F = (X, A)$ a family $\sigma(F) \subseteq 2^X$ of sets of arguments, called *extensions*. We denote by adm, com, prf, sem and stb the *admissible*, *complete*, *preferred*, *semi-stable* and *stable* semantics, respectively. These five semantics are characterized by the following conditions which hold for each AF $F = (X, A)$ and each conflict-free set $S \subseteq X$.

- $S \in \text{adm}(F)$ if and only if each $s \in S$ is defended by S.
- $S \in \text{com}(F)$ if and only if $S \in \text{adm}(F)$ and every argument that is defended by S is contained in S.

- $S \in \text{prf}(F)$ if and only if $S \in \text{adm}(F)$ and there is no $T \in \text{adm}(F)$ with $S \subsetneq T$.
- $S \in \text{sem}(F)$ if and only if $S \in \text{adm}(F)$ and there is no $T \in \text{adm}(F)$ with $S^+ \subsetneq T^+$.
- $S \in \text{stb}(F)$ if and only if $S^+ = X$.

Parameterized Complexity For our investigation we need to take two measurements into account: the input size n of the given AF F and the parameter k given as the input to σ-REPAIR, σ-ADJUST, and σ-CENTER. The theory of *parameterized complexity*, introduced and pioneered by Downey and Fellows [9], provides the adequate concepts and tools for such an investigation. We outline the basic notions of parameterized complexity that are relevant for this paper, for an in-depth treatment we refer to other sources [20, 24].

An instance of a parameterized problem is a pair (I, k) where I is the *main part* and k is the *parameter*; the latter is usually a non-negative integer. A parameterized problem is *fixed-parameter tractable* (FPT) if there exists a computable function f such that instances (I, k) of size n can be solved in time $f(k) \cdot n^{O(1)}$, or equivalently, in fpt-time. Fixed-parameter tractable problems are also called *uniform polynomial-time tractable* because if k is considered constant, then instances with parameter k can be solved in polynomial time where the order of the polynomial is independent of k, in contrast to *non-uniform polynomial-time* running times such as $n^{O(k)}$. Thus we have three complexity categories for parameterized problems: (1) problems that are fixed-parameter tractable (uniform polynomial-time tractable), (2) problems that are non-uniform polynomial-time tractable, and (3) problems that are NP-hard or coNP-hard if the parameter is fixed to some constant (such as k-SAT which is NP-hard for $k = 3$). The major complexity assumption in parameterized complexity is FPT $\subsetneq$ W[1]. Hence, W[1]-hard problems are not fixed-parameter tractable under this assumption. Still, such problems are non-uniform polynomial-time tractable. Problems that fall into (3) above are said to be para-NP-hard or para-coNP-hard. The classes in parameterized complexity are defined by *fpt-reduction*, which performs in fpt-time and bounds the new parameter k' by a computable function in k.

In our proofs of complexity results we will reduce from the following problem, which is W[1]-complete [25].

> MULTICOLORED CLIQUE
> *Instance:* A natural number k, and a k-partite graph $G = (V, E)$ with partition $\{V_1, \ldots, V_k\}$.
> *Parameter:* k.
> *Question:* Does G contain a clique of size k?

W.l.o.g. we may assume that the parameter k of MULTICOLORED CLIQUE is even. To see this, we reduce from MULTICOLORED CLIQUE to itself as follows. Given an instance (G, k) of MULTICOLORED CLIQUE we construct an equivalent instance $(G', 2k)$ of MULTICOLORED CLIQUE where G' is the disjoint union of 2 copies of G.

First-Order Logic We will briefly recall the syntax and semantics of *First-Order formulas* (FO formulas). FO formulas are evaluated over structures defined on a fixed vocabulary. A (relational) *vocabulary* τ is a set of relation symbols. Each relation symbol $R \in \tau$ has an associated arity $\text{arity}(R) \geq 1$. A *structure* $\mathcal{U}$ of vocabulary τ, or τ-structure, consists of a set U called the *universe* and an interpretation $R^{\mathcal{S}} \subseteq A^{\text{arity}(R)}$ for each relational symbol $R \in \tau$.

To define the syntax of FO formulas we fix a countably infinite set of variables which we denote in the following by lower case letters with or without indicies, e.g., x, y, x_1 etc.. Let τ be a vocabulary. The syntax of FO formulas is defined inductively as follows: (1) if x, y, and $x_1, \dots, x_l$ are variables and R is an l-ary relation symbol in τ then $x = y$ and $Rx_1, \dots, x_l$ are FO formulas, (2) if φ and φ' are FO formulas, then $\varphi \wedge \varphi'$, $\varphi \vee \varphi'$, and $\neg\varphi$ are FO formulas, and (3) if x is a variable and φ is a FO formula, then $\exists x\varphi$ and $\forall x\varphi$ are FO formulas. Let x and y be variables and φ and φ' FO formulas. For convenience we also introduce the following abbreviations: (1) $x \neq y$ as a shortcut for $\neg x = y$, (2) $\varphi \rightarrow \varphi'$ as a shortcut for $\neg\varphi \vee \varphi'$, and (3) $\varphi \leftrightarrow \varphi'$ as a shortcut for $(\varphi \rightarrow \varphi') \wedge (\varphi' \rightarrow \varphi)$.

By $\text{free}(\varphi)$ we denote the set of *free variables* of φ, i.e., the set of all variables x that occur in φ but are not in the scope of a quantifier binding in φ. A *sentence* is a formulas without free variables. We write $\varphi(x_1, \dots, x_k)$ to indicate that φ is a first-order formula with $\text{free}(\varphi) \subseteq \{x_1, \dots, x_k\}$. We also use the notation $\varphi(x_1, \dots, x_k)$ to conveniently indicate substitutions. For instance, if $\varphi(x)$ is a formula then $\varphi(y)$ denotes the formula obtained from $\varphi(x)$ by replacing all free occurrences of x by y, renaming bound variables if necessary.

To define the semantics, for each FO formula $\varphi(x_1, \dots, x_k)$ and each structure $\mathcal{U}$ over some vocabulary τ, we define the relation $\varphi(\mathcal{U}) \subseteq U^k$ inductively as follows:

- If $\varphi(x_1, \dots, x_k) = Rx_{i_1} \dots x_{i_l}$ where $R \in \tau$ is l-ary and $i_1, \dots, i_l \in \{1, \dots, k\}$, then

  $$\varphi(\mathcal{U}) := \{\, (u_1, \dots, u_k) \in U^k \mid (u_{i_1}, \dots, u_{i_l}) \in R^{\mathcal{U}} \,\}.$$

  Equalities are treated similarly.
- If $\varphi(x_1, \dots, x_k) = \varphi(x_{i_1}, \dots, x_{i_l}) \wedge \chi(x_{j_1}, \dots, x_{j_r})$ with $i_1, \dots, i_l, j_1, \dots, j_r \in \{1, \dots, k\}$, then

  $$\varphi(\mathcal{U}) := \{\, (u_1, \dots, u_k) \in U^k \mid (a_{i_1}, \dots, a_{i_l}) \in \varphi(\mathcal{U}), \text{ and } (a_{j_1}, \dots, a_{j_r}) \in \chi(\mathcal{U}) \,\}.$$

  The other connectives are treated similarly.
- If $\varphi(x_1, \dots, x_k) = \exists x_{k+1}\varphi(x_{i_1}, \dots, x_{i_l})$ with $i_1, \dots, i_l \in \{1, \dots, k+1\}$, then

  $$\varphi(\mathcal{U}) := \{\, (u_1, \dots, u_k) \in U^k \mid \text{ there exists an } u_{k+1} \in U \text{ such that } (u_{i_1}, \dots, u_{i_l}) \in \varphi(\mathcal{U}) \,\}.$$

  Universal quantifiers are treated similarly.

The above definition also applies for the case that $k = 0$; in this case, $\varphi(\mathcal{A})$ is either the empty set or the set consisting of the empty tuple. We usually write $\mathcal{U} \models \varphi(x_1, \dots, x_k)$ instead of $(a_1, \dots, a_k) \in \varphi(\mathcal{U})$. If φ is a sentence, we simply write $\mathcal{U} \models \varphi$ instead of $\varphi(\mathcal{U}) \neq \emptyset$ and say that $\mathcal{U}$ *satisfies* φ or $\mathcal{U}$ is a *model* of φ. Observe that for a sentence φ the condition $\varphi(\mathcal{U})$ just means that $\varphi(\mathcal{U})$ contains the empty tuple.

In this paper we are mainly interested in structures that model (argument)-labeled AF or equivalently (vertex)-labeled directed graphs. We represent the labels of an AF $F = (X, A)$ by a family of subsets $\{L_1, \dots, L_r\}$ of X, i.e., for every $1 \leq i \leq r$, $L_i \subseteq X$. Let $F = (X, A)$ be a labeled AF with labels $\{L_1, \dots, L_r\}$. We denote by $\tau(F, \{L_1, \dots, L_r\})$ the vocabulary $\{A, L_1, \dots, L_r\}$ and by $\mathcal{F}(\{L_1, \dots, L_r\})$ the $\tau(F, \{L_1, \dots, L_r\})$-structure with universe X, 1 binary relation A with $A^{\mathcal{F}(\{L_1,\dots,L_r\})} := A$, and 1 unary relation L_i with $L_i^{\mathcal{F}(\{L_1,\dots,L_r\})} := L_i$ for every $1 \leq i \leq r$.

Let $\mathcal{C}$ be a class of (possibly labeled) AFs. We consider the following problem.

$\mathcal{C}$-FO MODEL CHECKING
Instance: A labeled AF $F = (X, A)$ with labels $\{L_1, \dots, L_r\}$ such that $F \in \mathcal{C}$ and an FO sentence φ over vocabulary $\tau(F, \{L_1, \dots, L_r\})$.
Parameter: $|\varphi|$.
Question: $\mathcal{F}(\{L_1, \dots, L_r\}) \models \varphi$

3 Problems for Dynamic Argumentation

In this section we present the problems that we consider for dynamic argumentation. Let $\sigma \in \{\text{adm}, \text{com}, \text{prf}, \text{sem}, \text{stb}\}$ and recall that $\triangle$ and dist are defined as the symmetric difference and the cardinality of the symmetric difference between 2 sets of arguments, respectively.

σ-SMALL
Instance: An AF $F = (X, A)$, a nonnegative integer k.
Parameter: k.
Question: Is there a nonempty extension $E \in \sigma(F)$ of size at most k?

σ-REPAIR
Instance: An AF $F = (X, A)$, a set of arguments $S \subseteq X$, a nonnegative integer k.
Parameter: k.
Question: Is there a nonempty extension $E \in \sigma(F)$ s.t. $\text{dist}(E, S) \leq k$?

σ-ADJUST
Instance: An AF $F = (X, A)$, an extension $E_0 \in \sigma(F)$, an argument $t \in X$, a nonnegative integer k.
Parameter: k.
Question: Is there an extension $E \in \sigma(F)$ s.t. $\text{dist}(E, E_0) \leq k$ and $t \in E_0 \triangle E$?

σ-CENTER
Instance: An AF $F = (X, A)$, two extensions $E_1, E_2 \in \sigma(F)$.
Parameter: $\text{dist}(E_1, E_2)$.
Question: Is there an extension $E \in \sigma(F)$ s.t. $\text{dist}(E, E_i) < \text{dist}(E_1, E_2)$ for every $i \in \{1, 2\}$?

4 Hardness Results

This section is devoted to our hardness results. We start by showing that all the problems that we consider in the context of dynamic argumentation are W[1]-hard and hence unlikely to have FPT-algorithms.

Theorem 1. *Let* $\sigma \in \{\text{adm}, \text{com}, \text{prf}, \text{sem}, \text{stb}\}$. *Then the problems* σ-SMALL*,* σ-REPAIR*,* σ-ADJUST*,* σ-CENTER *are* W[1]*-hard.*

We note here that W[1]-hardness also implies the NP-hardness of the considered problems. In the next section we will show that when considering AFs of bounded maximum degree then (fpt)-tractability can be obtained for the admissible, complete, and stable semantics. Unfortunately, this positive result does not hold for the preferred and semi-stable semantics as the following result shows.

Theorem 2. *Let* $\sigma \in \{\text{prf}, \text{sem}\}$. *Then the problems* σ-SMALL*,* σ-REPAIR*,* σ-ADJUST*,* σ-CENTER *are* para-coNP*-hard even for AF of maximum degree* 5.

Due to space-limitations we had to move the proof of Theorem 1 to the appendix.

We will now show Theorem 2.

Lemma 1. *Let* $\sigma \in \{\text{prf}, \text{sem}\}$. *Then the problems* σ-SMALL *and* σ-REPAIR *are* para-coNP*-hard (for parameter equal to* 1*) even if the maximum degree of the AF is bounded by* 5*.*

Proof. We will show the theorem by providing a polynomial reduction from the 3-CNF-2-UNSATISFIABLILY problem which is well-known to be coNP-hard [21]. The 3-CNF-2-UNSATISFIABLILY problem ask whether a given 3-CNF-2 formula Φ, i.e., Φ is a CNF formula where every clause contains at most 3 literals and every literal occurs in at most 2 clauses, is not satisfiable. Let Φ be a such a 3-CNF-2 formula with clauses $C_1, \dots, C_m$ and variables $x_1, \dots, x_n$. We will (in polynomial time) construct an AF $F = (X, A)$ such that (1) $\overline{F}$ has bounded degree and (2) Φ is not satisfiable if and only if there is an $E \in \sigma(F)$ with $|E| = 1$. This implies the theorem.

F contains the following arguments: (1) 2 arguments Φ and $\overline{\Phi}$, (2) 1 argument C_j for every $1 \leq j \leq m$, (3) 2 arguments x_i and $\overline{x_i}$ for every $1 \leq i \leq n$, and (4) 1 argument e. Furthermore, F contains the following attacks: (1) 1 self-attack for the arguments $\overline{\Phi}$ and $C_1, \dots, C_m$, (2) 1 attack from Φ to $\overline{\Phi}$, (3) 1 attack from C_j to Φ for every $1 \leq j \leq m$, (4) 1 attack from x_i to C_j for every $1 \leq i \leq n$ and $1 \leq j \leq m$ such that $x_i \in C_j$, (5) 1 attack from $\overline{x_i}$ to C_j for every $1 \leq i \leq n$ and $1 \leq j \leq m$ such that $\overline{x_i} \in C_j$, (6) 2 attacks from x_i to $\overline{x_i}$ and from $\overline{x_i}$ to x_i for every $1 \leq i \leq n$, and (7) 2 attacks from $\overline{\Phi}$ to x_i and to $\overline{x_i}$ for every $1 \leq i \leq n$.

Note that the constructed AF F does not have bounded degree. Whereas all arguments in $X \setminus \{\Phi, \overline{\Phi}\}$ have degree at most 5, the degree of the arguments Φ and $\overline{\Phi}$ can be unbounded. However, the following simple trick can be used to transform F into an AF with bounded degree.

Let $B(i)$ be an undirected rooted binary tree with root r and i leaves $l_1, \dots, l_i$ and let $B'(i)$ be obtained from B after subdividing every edge of B once, i.e., every edge $\{u, v\}$ is replaced with 2 edges $\{u, n_{uv}\}$ and $\{n_{uv}, v\}$ where n_{uv} is a new vertex for every such edge. We denote by $B(\Phi)$ the rooted directed tree obtained from $B'(m)$ after directing every edge of B' towards the root r and introducing a self-attack for every vertex in $V(B') \setminus V(B)$, i.e., all vertices introduced for subdividing edges of B are self-attacking in $B(\Phi)$. Then to ensure that the argument Φ has bounded degree in F we first delete the attacks from the arguments $C_1, \dots, C_m$ to Φ in F. We then add a copy of $B(\Phi)$ to F and identify Φ with the root r. Finally, we add 1 attack from C_j to l_j for every $1 \leq j \leq m$. Observe that this construction maintains the property of F that if a σ-extension of F contains Φ then it also has to contain at least 1 attacker of every argument $C_1, \dots, C_m$.

Let $B(\overline{\Phi})$ be the rooted directed tree obtained from $B'(2n)$ after directing every edge of B' away from the root r and introducing a self-attack for every vertex in $V(B)$. To ensure that also the argument $\overline{\Phi}$ has bounded degree we first delete the attacks from the argument $\overline{\Phi}$ to $x_1, \overline{x_1}, \dots, x_n, \overline{x_n}$ in F. We then add a copy of $B(\overline{\Phi})$ to F and identify $\overline{\Phi}$ with the root r. Finally, we add 2 attacks from l_i to x_i and from l_{n+i} to $\overline{x_i}$ for every $1 \leq i \leq n$. Observe that this construction maintains the property of F that if a σ-extension of F contains x_i or $\overline{x_i}$ for some $1 \leq i \leq n$ then $\overline{\Phi}$ needs to be attacked by the argument Φ in F and hence such a σ-extension has to contain the argument Φ.

Clearly, after applying the above transformations to F the resulting AF has maximum degree at most 5. However, because it is straightforward to verify but tedious to proof the remaining theorem for the transformed AF with bounded degree we will henceforth prove the theorem for F. We will need the following claim.

Claim 1. *If there is an $E \in \mathrm{adm}(F)$ that contains at least 1 argument in $\{\Phi, x_1, \overline{x_2}, \dots, x_n, \overline{x_n}\}$ then $\Phi \in E$.*

Let $E \in \mathrm{adm}(F)$ with $E \cap \{\Phi, x_1, \dots, x_n\} \neq \emptyset$. If $\Phi \in E$ then the claim holds. So suppose that $\Phi \notin E$. Then there is an $1 \leq i \leq n$ such that either $x_i \in E$ or $\overline{x_i} \in E$. Because both x_i and $\overline{x_i}$ are attacked by the argument $\overline{\Phi}$ and the only argument (apart from $\overline{\Phi}$) that attacks $\overline{\Phi}$ in F is Φ it follows that $\Phi \in E$. This shows the claim.

Claim 2. *There is an $E \in \mathrm{adm}(F)$ that contains at least 1 argument in $\{\Phi, x_1, \overline{x_2}, \dots, x_n, \overline{x_n}\}$ if and only if the formula Φ is satisfiable.*

Suppose there is an $E \in \mathrm{adm}(F)$ with $E \cap \{\Phi, x_1, \dots, x_n\} \neq \emptyset$. Because of the previous claim we have that $\Phi \in E$. Because $\Phi \in E$ and Φ is attacked by every argument $C_1, \dots, C_m$ it follows that every argument $C_1, \dots, C_m$ must be attacked by some argument in E. Let $a(C_j)$ be an argument in E that attacks C_j. Then $a(C_j)$ is an argument that corresponds to a literal of the clause C_j. Furthermore, because E is conflict-free the set $L := \{\, a(C_j) \mid 1 \leq j \leq m \,\}$ does not contain arguments that correspond to complementary literals. Hence, L corresponds to a satisfying assignment of Φ.

For the reverse direction suppose Φ is satisfiable and let L be a set of literals witnessing this, i.e., L is a set of literals that correspond to a satisfying assignment of Φ. It is straightforward to check that $E := \{\Phi\} \cup L$ is in $\mathrm{adm}(F)$. This completes the proof of the claim.

Claim 3. *Let $E \in \sigma(F)$. Then $e \in E$.*

This follows directly from our assumption that $\sigma \in \{\text{prf}, \text{sem}\}$ and the fact that the argument e is isolated in F.

We are now ready to show that Φ is not satisfiable if and only if there is an $E \in \sigma(F)$ with $|E| = 1$. So suppose that Φ is not satisfiable. It follows from the previous claim that $E \cap \{\Phi, x_1, \overline{x_1}, \ldots, x_n, \overline{x_n}\} = \emptyset$ for every $E \in \text{adm}(F)$ and hence also for every $E \in \sigma(F)$. Because of the self-attacks of the arguments in $\{\overline{\Phi}, C_1, \ldots, C_m\}$, we obtain that $E \subseteq \{e\}$. Using the previous claim, we have $E = \{e\}$ as required.

For the reverse direction suppose that there is an $E \in \sigma(F)$ with $|E| = 1$. Because of the previous claim it follows that $E = \{e\}$. Furthermore, because of the maximality condition of the preferred and semi-stable semantics it follows that there is no $E \in \text{adm}(F)$ such that $E \cap \{\Phi, x_1, \overline{x_1}, \ldots, x_n, \overline{x_n}\} \neq \emptyset$ and hence (using Claim 2) the formula Φ is not satisfiable. □

Lemma 2. *Let* $\sigma \in \{\text{prf}, \text{sem}\}$*. Then the problem* σ-ADJUST *is* para-coNP-*hard (for parameter equal to* 2*) even if the maximum degree of the AF is bounded by* 5.

Proof. We use a similar construction as in the proof of Theorem 1. Let F be the AF constructed from the 3-CNF-2 formulas Φ as in the proof of Theorem 1. Furthermore, let F' be the AF obtained from F after removing the argument e and adding 4 novel arguments t_1, t'_1, t_2, and t'_2 and the attacks (t_1, Φ), (Φ, t_1), (t_1, t_2), (t_2, t_1), (t_1, t'_1), (t_2, t'_2), (t'_1, t'_1), and (t'_2, t'_2) to F. Because F has degree bounded by 5 (and the degree of the argument Φ in F is 3) it follows that the maximum degree of F' is 5 as required. We claim that $(F', \{t_1\}, t_1, 2)$ is a YES-instance of σ-ADJUST if and only if Φ is not satisfiable.

It is straightforward to verify that the Claims 1 and 2 also hold for the AF F'. We need the following additional claims.

Claim 4. $\{t_1\} \in \sigma(F')$.

Clearly, $\{t_1\} \in \text{adm}(F')$. We first show that for every $E \in \text{adm}(F')$ with $t_1 \in E$ it holds that $E = \{t_1\}$. Let $E \in \text{adm}(F')$ with $t_1 \in E$. Because of the attacks between t_1 and t_2 and between t_1 and Φ it follows that $\Phi, t_2 \notin E$. Using Claim 1 it follows that also none of the arguments in $\{x_1, \overline{x_1}, \ldots, x_n, \overline{x_n}\}$ are contained in E. Furthermore, because of the self-attacks in F' it also holds that none of the arguments in $\{\overline{\Phi}, C_1, \ldots, C_m, t'_1, t'_2\}$ are contained in E. Hence, $E = \{t_1\}$, as required. This implies that $\{t_1\} \in \text{prf}(F')$. To show that $\{t_1\} \in \text{sem}(F')$ observe that t_1 is the only argument in F (apart from t'_1 itself) that attacks t'_1. Furthermore, because t'_1 attacks itself it cannot be in any semi-stable extension of F'. Hence, $\{t_1\} \in \text{sem}(F')$. This shows the claim.

Claim 5. $\{t_2\} \in \sigma(F')$ *if and only if* Φ *is not satisfiable.*

Suppose that $\{t_2\} \in \sigma(F')$. If $\{t_2\} \in \text{prf}(F')$ then there is no $E \in \text{adm}(F')$ with $\{t_2\} \subsetneq E$. It follows that there is no $E' \in \text{adm}(F')$ with $E' \neq \cap\{\Phi, x_1, \overline{x_1}, \ldots, x_n, \overline{x_n}\} \neq \emptyset$, since such an E' could be added to E. Using Claim 2 it follows that Φ is not satisfiable. If on the other hand $\{t_2\} \in \text{sem}(F')$ then because t_2 is the only argument that attacks t'_2 and because of the self-attack of t'_2 it follows again that there is no $E \in \text{adm}(F')$ with $\{t_2\} \subsetneq E$. Hence, using the same arguments as for the case $\{t_2\} \in \text{prf}(F')$ we again obtain that Φ is not satisfiable.

For the reverse direction suppose that Φ is not satisfiable. Because of Claim 2 we obtain that there every $E \in \text{adm}(F')$ (and hence also every $E \in \sigma(F')$) contains no argument in $\{\Phi, x_1, \overline{x_1}, \ldots, x_n, \overline{x_n}\}$. Because $\{t_2\} \in \text{adm}(F')$ and attacks the only remaining argument t_1 with no self-attack it follows that $\{t_2\} \in \sigma(F')$.

To show the theorem it remains to show that there is an $E' \in \sigma(F')$ with $t \notin E'$ and $\text{dist}(E, E') \leq 2$ if and only if the formula Φ is not satisfiable. First observe that because of Claim 4, $\emptyset \notin \sigma(F')$ and hence E' must contain exactly 1 argument other than t_1. Consequently, it remains to show that there is an argument $x \in X \setminus \{t_1\}$ such that $\{x\} \in \sigma(F')$ if and only if Φ is not satisfiable.

Suppose that there is an $x \in X \setminus \{t_1\}$ with $\{x\} \in \sigma(F')$. If $x \in \{\Phi, x_1, \overline{x_1}, \ldots, x_n, \overline{x_n}\}$ then because of Claim 1 it holds that $x = \Phi$. However, assuming that Φ contains at least 1 clause it follows that $\{x\}$ is not admissible, and hence $x \neq \Phi$. Considering the self-attacks of F we obtain that $x = t_2$. Hence, the forward direction follows from Claim 5.

The reverse direction follows immediately from Claim 5. This concludes the proof of the theorem. ☐

Lemma 3. *Let $\sigma \in \{\text{prf}, \text{sem}\}$. Then the problem σ-CENTER is* para-coNP*-hard (for parameter equal to* 6*) even if the maximum degree of the AF is bounded by* 5.

Proof. We use a similar construction as in the proof of Theorem 1. Let F be the AF constructed from the 3-CNF-2 formulas Φ as in the proof of Theorem 1. Furthermore, let F' be the AF obtained from F after removing the argument e and adding 12 novel arguments t, t', w_1, w_2, w'_1, w'_2, z, z', z_1, z'_1, z_2, z'_2 and the attacks (t, z), (z, z), (t', z'), (z', z'), (w_1, z_1), (z_1, z_1), (w'_1, z'_1), (z'_1, z'_1), (w_2, z_2), (z_2, z_2), (w'_2, z'_2), (z'_2, z'_2), (t, Φ), (Φ, t), (t', Φ), (Φ, t'), (t, t'), (t', t), (w_1, w'_1), (w'_1, w_1), (w_2, w'_2), (w'_2, w_2), (w_1, t), (w_2, t), (w'_1, t'), and (w'_2, t) to F. Because F has degree bounded by 5 (and the degree of the argument Φ of F is 5) it follows that the maximum degree of F' is 6 as required. We claim that $(F', \{t, w'_1, w'_2\}, \{t', w_1, w_2\})$ is a YES-instance of σ-CENTER if and only if Φ is not satisfiable.

It is straightforward to verify that the Claims 1 and 2 also hold for the AF F'. We need the following additional claims.

Claim 6. $\{t, w'_1, w'_2\} \in \sigma(F')$ *and* $\{t', w_1, w_2\} \in \sigma(F')$.

We show that $\{t, w'_1, w'_2\} \in \sigma(F')$. The case for $\{t', w_1, w_2\} \in \sigma(F')$ is analogous due to the symmetry of F'. Clearly, $\{t, w'_1, w'_2\} \in \text{adm}(F')$.

We first show that for every $E \in \text{adm}(F')$ with $t \in E$ it holds that $E = \{t, w'_1, w'_2\}$. Let $E \in \text{adm}(F')$ with $t \in E$. Clearly, E does not contain Φ, t', w_1 or w_2 (since these arguments are neighbors of t in F). Using Claim 1 it follows that also none of the arguments in $\{x_1, \overline{x_1}, \ldots, x_n, \overline{x_n}\}$ are contained in E. Furthermore, because of the self-attacks in F' it also holds that none of the arguments in $\{\overline{\Phi}, C_1, \ldots, C_m, z, z'\}$ are contained in E. Hence, $E \subseteq \{t_1, w'_1, w'_2\}$. However, because t is attacked by w_1 and w_2 in F and w'_1 and w_2 are the only arguments of F that attack w_1 and w_2 it follows that $E = \{t, w'_1, w'_2\}$. This implies that $\{t, w'_1, w'_2\} \in \text{prf}(F')$. To show that $\{t, w'_1, w'_2\} \in \text{sem}(F')$ observe that t is the only argument in F (apart from z itself) that attacks z. Furthermore, because z attacks itself it cannot be in any semi-stable extension of F'. Hence, $\{t, w'_1, w'_2\} \in \text{sem}(F')$. This shows the claim.

The proof of the previous claim actually showed the following slightly stronger statement.

Claim 7. *Let $E \in \sigma(F')$ with $t \in E$. Then $E = \{t, w_1', w_2'\}$. Similarly, if $E \in \sigma(F')$ with $t' \in E$. Then $E = \{t', w_1, w_2\}$.*

We are now ready to show that there is an $E \in \sigma(F')$ with $\text{dist}(E, E_i) < \text{dist}(E_1, E_2) = 6$ for every $i \in \{1, 2\}$ if and only if the formula Φ is not satisfiable.

Suppose that there is an $E \in \sigma(F')$ with $\text{dist}(E, E_i) < \text{dist}(E_1, E_2) = 6$ for every $i \in \{1, 2\}$. Then because of Claim 7 E does not contain t or t'. If there is an $E \in \sigma(F')$ with $\Phi \in E$ then we can assume (because of the maximality properties of the two semantics) that E contains 1 of x_i or $\overline{x_i}$ for every $1 \leq i \leq n$. Hence, if $\Phi \in E$ and the formula Φ contains at least 5 variables (which we can assume w.l.o.g.) then $\text{dist}(E, E_1) > 5$. Consequently, $\Phi \notin E$ and it follows from Claims 1 and 2 that Φ is not satisfiable, as required.

For the reverse direction suppose that Φ is not satisfiable. Let $E := \{w_1, w_1'\}$. Clearly, $\text{dist}(E, E_i) = 3 < 5$, as required. It remains to show that $E \in \sigma(F')$. It is easy to see that $E \in \text{adm}(F')$. Furthermore, Φ is not satisfiable it follows from Claim 2 that no $E' \in \sigma(F')$ can contain an argument in $\{\Phi, x_1, \overline{x_1}, \ldots, x_n, \overline{x_n}\}$ and hence $E \in \text{prf}(F')$. The maximality of E with respect to the semi-stable extension now follows from the fact that w_1 and w_2' are the only arguments that attack the arguments z_1 and z_2' and because of their self-attacks none of z_1 and z_2 can them-self be contained in a semi-stable extension. This completes the proof of the theorem. ☐

Lemmas 1, 2, and 3 together imply Theorem 2.

5 Tractability Results

Unfortunately, the results of the previous section draw a rather negative picture of the complexity of problems important to dynamic argumentation. In particular, Theorem 2 strongly suggests that at least for the preferred and semi-stable semantics these problems remain intractable even under strong structural restrictions. The hardness of these problems under the preferred and semi-stable semantics seems to originate from their maximality conditions. In this section we take a closer look at the complexity of our problems for the 3 remaining semantics, i.e., the admissible, complete, and stable semantics. We show that in contrast to the preferred and semi-stable semantics all our problems become fixed-parameter tractable when restricting the structure of the given AF (e.g. to have bounded degree). In particular, we will show the following result.

Theorem 3. *Let $\sigma \in \{\text{adm}, \text{com}, \text{stb}\}$ and c a natural number. Then the problems σ-*SMALL*, σ-*REPAIR*, σ-*ADJUST*, and σ-*CENTER* are fixed-parameter tractable if the maximum degree of the input AF is bounded by c.*

The rest of this section is devoted to a proof of this theorem.

Our main tool for the proof of the above theorem is the following well-known result for $\mathcal{C}$-FO MODEL CHECKING – this result is originally stated for labeled directed graphs, however, the result directly applies to labeled AFs due to the equivalent nature of labeled AF and labeled directed graphs.

Proposition 1 ([27]). *Let $\mathcal{C}$ be the class of all (possibly labeled) AFs whose maximum degree is bounded by some constant. Then the problem $\mathcal{C}$-FO MODEL CHECKING is fixed-parameter tractable.*

There exists several extensions of the above result to even more general classes, e.g., the class of graphs with locally bounded treewidth. Due to the technicality of the definition of these classes we refrain from stating these results in detail and refer the interested reader to [23]. Results such as the one above are also commonly refereed to as meta-theorems, i.e., they allow us to make statements about a wide variety of algorithmic problems. Similar meta-theorems have been used before in the context of Abstract Argumentation (see, e.g., [11, 22, 17]).

To use this result we need to state our problems in terms of the $\mathcal{C}$-FO MODEL CHECKING problem, i.e., we need to (1) model the input of σ-SMALL, σ-REPAIR, σ-ADJUST, σ-CENTER in terms of labeled AFs, and (2) give a FO sentence that is satisfied if and only if the given instance of σ-SMALL, σ-REPAIR, σ-ADJUST, σ-CENTER (represented by the labeled AF from step (1)) is a YES instance.

To accomplish step (1) we only need to define the labels from the input of the given problem. We do this in the natural way, i.e., for an instance (F, k) of σ-SMALL, F is the corresponding labeled AF, for an instance (F, S, k) of σ-REPAIR, F with labels $\{S\}$ is the corresponding labeled AF, for an instance (F, E_0, t, k) of σ-ADJUST, F with labels $\{E_0, T\}$ where $T := \{t\}$ is the corresponding labeled AF, and for an instance (F, E_1, E_2) of σ-CENTER, F with labels $\{E_1, E_2\}$ is the corresponding labeled AF.

Towards defining the FO formulas for step (2) we start by defining the following auxiliary formulas. Due to the complexity of the FO formulas that we need to define, we will introduce some additional notation that will allow us to reuse formulas by substituting parts of other formulas. We will provide examples how to interpret the notation when these formulas are introduced.

In the following let l be a natural number, and let $\varphi(x)$, $\varphi_1(x)$, and $\varphi_2(x)$ be FO formulas with free variable x.

The formula $\text{SET}[l](x_1, \ldots, x_l, y)$ is satisfied if and only if the argument y is equal to at least 1 of the arguments $x_1, \ldots, x_l$.

$$\text{SET}[l](x_1, \ldots, x_l, y) := y = x_1 \vee \cdots \vee y = x_l$$

We note here that the notation $\text{SET}[l]$ means that the exact definition of the formula $\text{SET}[l]$ depends on the value of l, e.g., if $l = 3$ then $\text{SET}[l]$ is the formula $y = x_1 \vee y = x_2 \vee y = x_3$.

The formula $\text{CF}[\varphi(x)]$ is satisfied if and only if the set of arguments that satisfy the formula $\varphi(x)$ is conflict-free.

$$\text{CF}[\varphi(x)] := \forall x \forall y (\varphi(x) \wedge \varphi(y)) \rightarrow \neg Axy$$

Again we note here that the notation $\text{CF}[\varphi(x)]$ means that the exact definition of the formula $\text{CF}[\varphi(x)]$ depends on the formula $\varphi(x)$, e.g., if $\varphi(x) := \text{SET}[l](x_1, \ldots, x_l, y)$ then $\text{CF}[\varphi(x)]$ is the formula $\forall x \forall y (\text{SET}[l](x_1, \ldots, x_l, x) \wedge \text{SET}[l](x_1, \ldots, x_l, y)) \rightarrow \neg Axy$ which in turn evaluates to $\forall x \forall y (\bigvee_{1 \leq i \leq l} x = x_i \wedge \bigvee_{1 \leq i \leq l} y = x_i) \rightarrow \neg Axy$.

The formula $\text{SYM-DIFF}[\varphi_1(x), \varphi_2(x)](y)$ is satisfied if and only if the argument y is contained in the symmetric difference of the sets of arguments that satisfy the formula $\varphi_1(x)$ and the set of arguments that satisfy the formula $\varphi_2(x)$.

$$\text{SYM-DIFF}[\varphi_1(x), \varphi_2(x)](y) := (\varphi_1(y) \wedge \neg\varphi_2(y)) \vee (\neg\varphi_1(y) \wedge \varphi_2(y))$$

The formula ATMOST$[\varphi(x), k]$ is satisfied if and only if the set of arguments that satisfy the formula $\varphi(x)$ contains at most k arguments.

$$\text{ATMOST}[\varphi(x), k] := \neg(\exists x_1, \ldots, \exists x_{k+1}(\bigwedge_{1\le i<j\le k+1} x_i \neq x_j) \wedge (\bigwedge_{1\le i\le k+1} \varphi(x_i)))$$

The following formulas model the semantics adm, com, stb. These formulas are therefore evaluated over a fixed (possibly labeled) AF $F := (X, A)$.

The formula adm$[\varphi(x)]$ is satisfied by the structure of a possibly labeled AF F if and only if the set of arguments that satisfy the formula $\varphi(x)$ is an admissible extension of F.

$$\text{adm}[\varphi(x)] := \text{CF}[\varphi(x)] \wedge (\forall x \forall z(\varphi(x) \wedge (\neg\varphi(z)) \wedge Azx) \rightarrow (\exists y \varphi(y) \wedge Ayz))$$

The formula com$[\varphi(x)]$ is satisfied by the structure of a possibly labeled AF F if and only if the set of arguments that satisfy the formula $\varphi(x)$ is a complete extension of F.

$$\text{com}[\varphi(x)] := \text{adm}[\varphi(x)] \wedge (\forall z((\forall a Aaz \rightarrow \exists x \varphi(x) \wedge Axa) \wedge (\forall x \varphi(x) \rightarrow \neg(Axz \vee Azx))) \rightarrow \varphi(z)$$

The formula stb$[\varphi(x)]$ is satisfied by the structure of a possibly labeled AF F if and only if the set of arguments that satisfy the formula $\varphi(x)$ is a stable extension of F.

$$\text{stb}[\varphi(x)] := \text{CF}[\varphi(x)] \wedge (\forall z \varphi(z) \vee (\exists a \varphi(a) \wedge Aaz))$$

We are now ready to define the formulas that model the problems σ-SMALL, σ-REPAIR, σ-ADJUST, and σ-CENTER.

Let $\sigma \in \{\text{adm}, \text{com}, \text{stb}\}$. The formula σ-SMALL$[\sigma, k]$ is satisfied by the structure of a possibly labeled AF F if and only if the AF F has a σ-extension that contains at most k arguments, i.e., if and only if (F, k) is a YES instance of σ-SMALL.

$$\sigma\text{-SMALL}[\sigma, k] := \exists x_1, \ldots, \exists x_k \sigma[\text{SET}[k](x_1, \ldots, x_k)]$$

The formula σ-REPAIR$[\sigma, k]$ is satisfied by the structure of a labeled AF $F = (X, A)$ with labels $\{S\}$ if and only if F has a $E \in \sigma(F)$ with dist$(E, S) \le k$, i.e., if and only if (F, S, k) is a YES instance of σ-REPAIR.

$$\sigma\text{-REPAIR}[\sigma, k] := \exists x_1, \ldots, \exists x_k \sigma[\text{SYM-DIFF}[Sx, \text{SET}[k](x_1, \ldots, x_k)]]$$

The formula σ-ADJUST$[\sigma, k]$ is satisfied by the structure of a labeled AF $F = (X, A)$ with labels $\{E_0, T\}$ where $T := \{t\}$ if and only if F has a $E \in \sigma(F)$ such that dist$(E_0, E) \le k$ and $t \in E \triangle E_0$, i.e., if and only if (F, E_0, t, k) is a YES instance of σ-ADJUST..

$$\sigma\text{-ADJUST}[\sigma, k] := \\ \exists t \exists x_1, \ldots, \exists x_{k-1} Tt \wedge \sigma[\text{SYM-DIFF}[E_0 x, \text{SET}[k](t, x_1, \ldots, x_{k-1})]]$$

The formula σ-CENTER$[\sigma, k]$ is satisfied by the structure of a labeled AF $F = (X, A)$ with labels $\{E_1, E_2\}$ if and only if F has a $E \in \sigma(F)$ with dist$(E_i, E) <$ dist(E_1, E_2) for every $i \in \{1, 2\}$, i.e., if and only if (F, E_1, E_2) is a YES instance of σ-CENTER.

$$\sigma\text{-CENTER}[\sigma, k] := \exists x_1, \ldots, \exists x_{k-1} \sigma[\text{SYM-DIFF}[E_1 x, \text{SET}[k-1](x_1, \ldots, x_{k-1})]] \wedge$$
$$\text{ATMOST}[k-1][\text{SYM-DIFF}[\text{SYM-DIFF}[E_1 x, \text{SET}[k-1](x_1, \ldots, x_{k-1})], E_2 x]$$

Because the length of the above FO formulas is easily seen to be bounded in terms of the parameter k of the respective problem, these formulas together with Proposition 1 immediately imply Theorem 3.

Acknowledgement

We would like to thank Stefan Woltran for stimulating discussions.

References

1. Pietro Baroni and Massimiliano Giacomin. Semantics of abstract argument systems. In Iyad Rahwan and Guillermo Simari, editors, *Argumentation in Artificial Intelligence*, pages 25–44. Springer Verlag, 2009.
2. Ringo Baumann. What does it take to enforce an argument? minimal change in abstract argumentation. In Luc De Raedt, Christian Bessière, Didier Dubois, Patrick Doherty, Paolo Frasconi, Fredrik Heintz, and Peter J. F. Lucas, editors, *ECAI 2012 - 20th European Conference on Artificial Intelligence. Including Prestigious Applications of Artificial Intelligence (PAIS-2012) System Demonstrations Track, Montpellier, France, August 27-31 , 2012*, volume 242 of *Frontiers in Artificial Intelligence and Applications*, pages 127–132. IOS Press, 2012.
3. Ringo Baumann and Gerhard Brewka. Expanding argumentation frameworks: Enforcing and monotonicity results. In Pietro Baroni, Federico Cerutti, Massimiliano Giacomin, and Guillermo Ricardo Simari, editors, *Computational Models of Argument: Proceedings of COMMA 2010, Desenzano del Garda, Italy, September 8-10, 2010*, volume 216 of *Frontiers in Artificial Intelligence and Applications*, pages 75–86. IOS Press, 2010.
4. T. J. M. Bench-Capon and Paul E. Dunne. Argumentation in artificial intelligence. *Artificial Intelligence*, 171(10-15):619–641, 2007.
5. Richard Booth, Martin Caminada, Mikolaj Podlaszewski, and Iyad Rahwan. Quantifying disagreement in argument-based reasoning. In Wiebe van der Hoek, Lin Padgham, Vincent Conitzer, and Michael Winikoff, editors, *International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2012, Valencia, Spain, June 4-8, 2012 (3 Volumes)*, pages 493–500. IFAAMAS, 2012.
6. Claudette Cayrol, Florence Dupin de Saint-Cyr, and Marie-Christine Lagasquie-Schiex. Revision of an argumentation system. In Gerhard Brewka and Jérôme Lang, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Eleventh International Conference, KR 2008, Sydney, Australia, September 16-19, 2008*, pages 124–134, 2008.
7. Sylvie Coste-Marquis, Caroline Devred, and Pierre Marquis. Symmetric argumentation frameworks. In Lluis Godo, editor, *Symbolic and Quantitative Approaches to Reasoning with Uncertainty, 8th European Conference, ECSQARU 2005, Barcelona, Spain, July 6-8, 2005, Proceedings*, volume 3571 of *Lecture Notes in Computer Science*, pages 317–328. Springer Verlag, 2005.
8. Yannis Dimopoulos and Alberto Torres. Graph theoretical structures in logic programs and default theories. *Theoretical Computer Science*, 170(1-2):209–244, 1996.
9. R. G. Downey and M. R. Fellows. *Parameterized Complexity*. Monographs in Computer Science. Springer Verlag, New York, 1999.

10. Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, 1995.
11. Paul E. Dunne. Computational properties of argument systems satisfying graph-theoretic constraints. *Artificial Intelligence*, 171(10-15):701–729, 2007.
12. Paul E. Dunne and T. J. M. Bench-Capon. Coherence in finite argument systems. *Artificial Intelligence*, 141(1-2):187–203, 2002.
13. Paul E. Dunne and Martin Caminada. Computational complexity of semi-stable semantics in abstract argumentation frameworks. In Steffen Hölldobler, Carsten Lutz, and Heinrich Wansing, editors, *Proceedings of the 11th European Conference on Logics in Artificial Intelligence JELIA 2008*, volume 5293 of *Lecture Notes in Computer Science*, pages 153–165. Springer Verlag, 2008.
14. Paul E. Dunne and Michael Wooldridge. Complexity of abstract argumentation. In L. Rahwan and G. R. Simari, editors, *Argumentation in Artificial Intelligence*, pages 85–104. Springer Verlag, 2009.
15. Wolfgang Dvorák, Sebastian Ordyniak, and Stefan Szeider. Augmenting tractable fragments of abstract argumentation. *Artificial Intelligence*, 186:157–173, 2012.
16. Wolfgang Dvorák, Reinhard Pichler, and Stefan Woltran. Towards fixed-parameter tractable algorithms for abstract argumentation. *Artificial Intelligence*, 186:1–37, 2012.
17. Wolfgang Dvorák, Stefan Szeider, and Stefan Woltran. Abstract argumentation via monadic second order logic. In Eyke Hüllermeier, Sebastian Link, Thomas Fober, and Bernhard Seeger, editors, *Scalable Uncertainty Management - 6th International Conference, SUM 2012, Marburg, Germany, September 17-19, 2012. Proceedings*, volume 7520 of *Lecture Notes in Computer Science*, pages 85–98. Springer Verlag, 2012.
18. Wolfgang Dvořák and Stefan Woltran. On the intertranslatability of argumentation semantics. In *In Proceedings of the Conference on Thirty Years of Nonmonotonic Reasoning(NonMon@30)*, Lexington, KY, USA, 2010.
19. Wolfgang Dvořák, Stefan Szeider, and Stefan Woltran. Reasoning in argumentation frameworks of bounded clique-width. In Pietro Baroni, Federico Cerutti, Massimiliano Giacomin, and Guillermo R. Simari, editors, *Computational Models of Argumentation, Proceedings of COMMA 2010*, volume 216 of *Frontiers in Artificial Intelligence and Applications*, pages 219–230. IOS, 2010.
20. Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*, volume XIV of *Texts in Theoretical Computer Science. An EATCS Series*. Springer Verlag, Berlin, 2006.
21. Michael R. Garey and David R. Johnson. *Computers and Intractability*. W. H. Freeman and Company, New York, San Francisco, 1979.
22. Eun Jung Kim, Sebastian Ordyniak, and Stefan Szeider. Algorithms and complexity results for persuasive argumentation. *Artificial Intelligence*, 175:1722–1736, 2011.
23. Stephan Kreutzer. Algorithmic meta-theorems. *Electronic Colloquium on Computational Complexity (ECCC)*, 16:147, 2009.
24. Rolf Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford Lecture Series in Mathematics and its Applications. Oxford University Press, Oxford, 2006.
25. Krzysztof Pietrzak. On the parameterized complexity of the fixed alphabet shortest common supersequence and longest common subsequence problems. *J. of Computer and System Sciences*, 67(4):757–771, 2003.
26. Iyad Rahwan and Guillermo R. Simari, editors. *Argumentation in Artificial Intelligence*. Springer Verlag, 2009.
27. Detlef Seese. Linear time computable problems and first-order descriptions. *Mathematical Structures in Computer Science*, 6(6):505–526, 1996.

A Omitted proofs

Proof of Theorem 1

We will have shown Theorem 1 after showing the following 3 Lemmas.

Lemma 4. *Let* $\sigma \in \{\text{adm}, \text{com}, \text{prf}, \text{sem}, \text{stb}\}$. *Then the problems* σ-SMALL *and* σ-REPAIR *are* W[1]*-hard.*

Proof. We start by showing the lemma for the problem σ-SMALL by giving an fpt-reduction from the MULTICOLORED CLIQUE problem to the σ-SMALL problem, when σ is one of the listed semantics. Let (G, k) be an instance of MULTICOLORED CLIQUE with partition $V_1, \ldots, V_k$. We construct in fpt-time an AF F such that there is an $E \in \sigma(F)$ with $|E| = k$ if and only if G has a k-clique. The AF F contains the following arguments: (1) 1 argument y_v for every $v \in V(G)$ and (2) for every $1 \leq i \leq k$, 1 argument z_v^j for every $v \in V_i$ and $1 \leq j \leq k$ with $j \neq i$. For every $1 \leq i < j \leq k$, we denote by $Y[i]$ the set of arguments $\{ a_v \mid v \in V_i \}$ and by $Z[i, j]$ the set of arguments $\{ a_v^j \mid v \in V_i \}$. Furthermore, we set $Y := \bigcup_{1 \leq i \leq k} Y[i]$ and $Z := \bigcup_{1 \leq i < j \leq k} Z[i, j]$. For every $1 \leq i \leq k$, the AF F contains the following attacks:

- 1 attack from y_v to y_u for every $u, v \in Y[i]$ with $u \neq v$;
- 1 self-attack for all arguments in Z;
- For every $v \in V_i$, 1 attack from z_v^j to y_v for every $1 \leq j \leq k$ with $j \neq i$;
- For every $v \in V_i$, 1 attack from y_v to z_u^j for every $u \in V_i \setminus \{v\}$ and $1 \leq j \leq k$ with $j \neq i$.
- For every $\{u, v\} \in E(G)$ with $u \in V_i$ and $v \in V_j$, 1 attack from y_u to z_v^i and 1 attack from y_v to z_u^j.

This completes the construction of F. It remains to show that G has a k-clique if and only if there is an $E \in \sigma(F)$ with $|E| = k$. If $Q \subseteq V(G)$ we denote by Y_Q the set of arguments $\{ y_q \mid q \in Q \}$. We need the following claim.

Claim 8. *A set* $Q \subseteq V$ *is a* k*-clique in* G *if and only if* $Y_Q \in \text{adm}(F)$ *and* $Y_Q \neq \emptyset$.

Suppose that $Q \subseteq V(G)$ is a k-clique in G. Then Y_Q contains exactly 1 argument from $Y[i]$ for every $1 \leq i \leq k$. Because there are no attacks between arguments in $Y[i]$ and $Y[j]$ for every $1 \leq i < j \leq k$ it follows that Y_Q is conflict-free. To see that Y_Q is also admissible let $y_v \in Y_Q \cap V_i$ and suppose that y_v is attacked by an argument x of F. It follows from the construction of F that either $x \in Y[i]$ or $x \in \{ z_v^j \mid 1 \leq j \leq k \text{ and } j \neq i \}$. In the first case x is attacked by y_v. In the second case z_v^j is attacked by the argument in $Y[j] \cap Y_Q$ because Q is a k-clique of G. Hence, $Y_Q \in \text{adm}(F)$ and $Y_Q \neq \emptyset$, as required.

For the opposite direction, suppose that $E \in \text{adm}(F)$ and $E \neq \emptyset$. Because E conflict-free it follows that $E \subseteq Y$ and E contains at most 1 argument from the set $Y[i]$ for every $1 \leq i \leq k$. Because $E \neq \emptyset$ there is an argument $y_v \in Y[i] \cap E$. Because of the construction of F, y_v is attacked by the arguments $\{ z_v^j \mid 1 \leq j \leq k \text{ and } j \neq i \}$. Hence, the arguments $\{ z_v^j \mid 1 \leq j \leq k \text{ and } j \neq i \}$ need to be attacked by arguments in E. However, the only arguments of F that attack an argument z_v^j with $j \neq i$ are the arguments $y_u \in Y[j]$ such that $\{u, v\} \in E(G)$. Hence, for every argument $y_v \in E \cap Y[i]$ and every $1 \leq j \leq k$ with $j \neq i$ there is an argument $y_u \in E \cap Y[j]$ such that $\{u, v\} \in E(G)$. It follows that the set $\{\, v \mid y_v \in E \,\}$ is a k-clique in G. This shows the claim.

The previous claim shows that every non-empty admissible extension of F corresponds to a k-clique of G. It is now straightforward to check that every such extension is not only admissible but also complete, preferred, semi-stable, and stable. This shows the lemma for σ-SMALL. To show the Lemma for the σ-REPAIR problem we note that $(F, \emptyset, k)$ is a YES-instance for σ-REPAIR if and only if (F, k) is a YES-instance for σ-SMALL. ☐

Lemma 5. *Let* $\sigma \in \{\text{adm}, \text{com}, \text{prf}, \text{sem}, \text{stb}\}$. *Then the problem* σ-ADJUST *is* W[1]-*hard.*

Proof. We give an fpt-reduction from the σ-SMALL problem. Let (F, k) be an instance of the σ-SMALL problem where $F = (X, A)$. We construct an equivalent instance (F', E_1, E_2) of the σ-ADJUST problem as follows. $F' = (X', A')$ is obtained from F by adding 1 argument t and 2 attacks (t, x) and (x, t) for every $x \in X$ to F. Because the argument t attacks is attacked by all arguments in X it follows that $\{t\}$ is a σ-extension of F'. In is now straightforward to show that $(F', \{t\}, t, k + 1)$ is a YES-instance of σ-ADJUST if and only if (F, k) is a YES-instance of σ-SMALL. This shows the lemma. ☐

Lemma 6. *Let* $\sigma \in \{\text{adm}, \text{com}, \text{prf}, \text{sem}, \text{stb}\}$. *Then the problem* σ*-Center is* W[1]-*hard.*

Proof. We give an fpt-reduction from the σ-SMALL problem. Let (F, k) be an instance of the σ-SMALL problem where $F = (X, A)$. W.l.o.g. we can assume that k is even. This follows from the remark in Section 2 that MULTICOLORED CLIQUE is W[1]-hard even if k is even and the parameter preserving reduction from MULTICOLORED CLIQUE to σ-SMALL given in Lemma 4. We will construct an equivalent instance (F', E_1, E_2) of the σ-CENTER problem as follows. $F' = (X', A')$ is obtained from F by adding the following arguments and attacks to F.

- 2 arguments t and t';
- the arguments in $W := \{w_1, \ldots, w_k\}$ and $W' := \{w'_1, \ldots, w'_k\}$;
- the arguments in $Z := \{z_1, \ldots, z_k\}$ and $Z' := \{z'_1, \ldots, z'_k\}$;
- attacks from t to all arguments in $X \cup \{t'\} \cup Z \cup Z'$ and attacks from t' to all arguments in $X \cup \{t\} \cup Z \cup Z'$;
- attacks from w_i to $\{t, w'_i\}$ and attacks from w'_i to $\{t', w_i\}$ for every $1 \leq i \leq k$;
- self-attacks for the arguments $z_1, \ldots, z_k$ and $z'_1, \ldots, z'_k$;
- attacks from z_i to $\{w_i, w'_i\}$ and from X to z_i for every $1 \leq i \leq k$;
- attacks from $\{w_i, w'_i\}$ to z'_i and from z'_i to X for every $1 \leq i \leq k$;

We set $E_0 := \{w_1, \ldots, w_{k/2}, w'_{k/2+1}, \ldots, w'_k\}$, $E_1 := \{t\} \cup W'$, $E_2 := \{t'\} \cup W$, and $k' := \text{dist}(E_1, E_2) - 1 = 2(k+1) - 1 = 2k + 1$. Then E_1 and E_2 are σ-extensions and hence (F', E_1, E_2) is a valid instance of the σ-CENTER problem. It remains to show that (F, k) is a YES instance of σ-SMALL if and only if (F', E_1, E_2) is a YES instance of σ-CENTER.

Suppose that (F, k) is a YES instance of σ-SMALL and let E be a non-empty σ-extension of cardinality at most k witnessing this. Then $E' := E \cup E_0$ is a σ-extension of F' and $\text{dist}(E', E_i) = k + k + 1 = 2k + 1 \leq k'$ for $i \in \{1, 2\}$, as required.

For the reverse direction suppose that E' is a σ-extension of F' with $\text{dist}(E', E_i) \leq k'$ for $i \in \{1, 2\}$. We need the following claim.

Claim 9. *E' does not contain t or t'.*

Suppose for a contradiction that E' contains one of t and t'. Because t and t' attack each other E' cannot contain both t and t'. W.l.o.g. we can assume that $t \in E'$. Because E' is a σ-extension E' is also admissible. Since, the arguments $w_1, \dots, w_k$ attack t, there need to be arguments in E' that attack these arguments. It follows that E' contains the arguments $w'_1, \dots, w'_k$. But then $\text{dist}(E', E_2) \geq \text{dist}(E_1, E_2)$ a contradiction.

Claim 10. *$E' \cap X$ is a non-empty σ-extension of F and E' contains exactly one of the arguments w_i and w'_i for every $1 \leq i \leq k$.*

It follows from the previous claim that E' does not contain t or t'. Furthermore, because of the self-loops of the arguments in $Z \cup Z'$, E' contains only arguments from $X \cup W \cup W'$. Since the arguments in X do not attack or are attacked by arguments in $W \cup W'$ it follows that $E' \cap X$ is a σ-extension of F. To see that $E' \cap X$ is also not empty, suppose for a contradiction that this is not the case. Then because E' is non-empty, E' has to contain at least 1 argument from $W \cup W'$. However, any argument in $W \cup W'$ is attacked by an argument in Z and the only arguments that attack arguments in Z are the arguments in $X \cup \{t, t'\}$. Again using the previous claim and the fact that E' is admissible, it follows that E' has to contain at least 1 argument from X, as required. It remains to show that E' contains exactly one of w_i and w'_i for every $1 \leq i \leq k$. Because E' contains at least 1 argument from X and all arguments in X are attacked by all arguments in Z', E' needs to contain arguments that attack all arguments in Z'. However, the only arguments that attack arguments in Z' are the arguments in $\{t, t'\} \cup W \cup W'$. Using the previous claim it follows that the only way for E' to attack all arguments in Z' is to contain at least 1 of w_i and w'_i for every $1 \leq i \leq k$. The claim now follows by observing that because E' is conflict-free, it cannot contain both arguments w_i and w'_i for any $1 \leq i \leq k$. This proves the claim.

Since E' contains exactly 1 of w_i and w'_i for every $1 \leq i \leq k$ we obtain that either $|W \setminus E'| \geq k/2$ or $|W' \setminus E'| \geq k/2$. W.l.o.g. we can assume that $|W \setminus E'| \geq k/2$. But then $\text{dist}(E', E_2) = |E' \cap X| + 1 + 2|W \setminus E'| = |E' \cap X| + k + 1$ and because $\text{dist}(E', E_2) \leq k' = 2k + 1$ it follows that $|E' \cap X| \leq k$. This concludes the proof of the lemma. □