FernUniversität in Hagen
Faculty of Mathematics and Computer Science

AIG
Artificial Intelligence Group

Serializability of Argumentation Semantics for Bipolar Argumentation Frameworks

Master's Thesis

in partial fulfillment of the requirements for
the degree of Master of Science (M.Sc.)
in Practical Computer Science

submitted by
Nikolaos Mitrakis

Examiner: Prof. Dr. Matthias Thimm
Artificial Intelligence Group

Statement

I hereby declare that I have written this master's thesis independently and without any improper assistance from third parties. I have used only the sources and aids listed and have clearly indicated all passages taken from these sources either verbatim or in substance. This declaration of independent work also applies to any drawings, sketches, or graphical representations included. The thesis has not previously been submitted, in the same or a similar form, to this or any other examination authority, nor has it been published. By submitting the electronic version of the final thesis, I acknowledge that it may be checked for plagiarism using a plagiarism detection service and will be stored exclusively for examination purposes.

	Yes	No
I agree to have this thesis published in the library.	☒	☐
I agree to have this thesis published on the webpage of the artificial intelligence group.	☒	☐
The thesis text is available under a Creative Commons License (CC BY-SA 4.0).	☒	☐
The source code is available under a GNU General Public License (GPLv3).	☒	☐
The collected data is available under a Creative Commons License (CC BY-SA 4.0).	☒	☐

Wuppertal, 18.02.2026
(Place, Date)

(Signature)

Abstract

Abstract Argumentation Frameworks (AFs) [10] model conflicts between arguments by a directed attack relation. Bipolar argumentation frameworks (BAFs) [6] extend this setting by adding support relations alongside attacks, enabling richer forms of argumentative reasoning. Serialisation [14] refers to the constructive, stepwise computation of extensions via initial sets and reducts, providing a decomposable view on argument acceptability.

This thesis formalises serialisation for BAFs under the deductive interpretation of support. We define initial sets and reducts via derived attacks, namely supported and secondary attacks, and we show that coherence-based semantics, including coherently admissible, complete, grounded, preferred, and stable extensions, admit uniform serialisation characterisations. Building on these results, we extract selection and termination functions that provide an operational view on extension construction in BAFs.

We analyse the complexity of core decision problems. In particular, we show that verification of coherent admissibility is decidable in polynomial time, and that the existence and uniqueness of initial sets are **NP**- and **DP**-hard, respectively, with matching completeness results on embedded AF instances. Thus, for the problems studied, adding deductive support does not increase asymptotic computational complexity beyond the classical AF case. Overall, the thesis provides a constructive characterisation of coherence-based semantics for bipolar argumentation and clarifies its relationship to the AF setting.

Contents

1 Introduction

Abstract argumentation frameworks (AFs) [10] model conflicts between arguments by an attack relation. A semantics assigns to an AF sets of arguments that can be jointly accepted. Prominent examples include admissible, complete, grounded, preferred, and stable semantics.

Bipolar argumentation frameworks (BAFs) extend AFs by adding a support relation alongside attacks [7, 8, 6]. Support represents positive dependencies between arguments and interacts with attack in ways that affect acceptability. Coherence-based approaches address this interaction by strengthening the classical notions of conflict-freeness and defence and by taking support-induced forms of indirect conflict into account [7, 6].

A constructive perspective on admissibility-based semantics for AFs is provided by serialisation. Thimm [14] shows that many such semantics can be understood as stepwise constructions that repeatedly select minimal admissible building blocks, called initial sets, and continue the construction in reducts of the framework. This viewpoint decomposes extensions into smaller components and provides an operational reading of how extensions are formed.

The aim of this thesis is to develop an analogous serialisation characterisation for BAFs under the deductive interpretation of support. In the bipolar setting, this transfer is not immediate. Indirect attack patterns depend on support chains and these chains can be destroyed by reduction steps. As a consequence, the construction primitives and the reduct operation have to be chosen so that they remain faithful to coherence-based acceptability.

1.1 Motivation

Serialisation is appealing because it links semantics to a concrete construction process. For AFs, selecting initial sets and applying reducts provides a uniform way of understanding admissibility-based semantics and supports algorithmic developments [14]. For BAFs, the additional support relation suggests that a comparable stepwise view could offer similar benefits, provided that it correctly captures the interaction of attack and support described by coherence-based approaches.

The key question is therefore whether one can define bipolar counterparts of initial sets and reducts such that coherent semantics can again be characterised as results of serialisation sequences. A further question is whether the resulting decision problems retain the same asymptotic complexity behaviour known from the AF case.

1.2 Research Gap and Contributions

Serialisation and initial sets have been studied in detail for AFs [14], while coherence-based semantics provide standard acceptability notions for BAFs based on the interaction of attack and support (e.g. [7, 8, 6]). This thesis connects these strands for deductive support by developing a serialisation-based characterisation of the coherence-based semantics considered here.

Concretely, we introduce BAF-initial sets and BAF reducts that take the interaction of attack and deductive support into account and that coincide with the AF notions when the support relation is empty. On this basis, we establish serialisation characterisations for coherently admissible, complete, grounded, preferred, and stable semantics in the bipolar setting, and we derive corresponding selection and termination conditions that make these characterisations operational. We also analyse the complexity of core decision problems, including coherent admissibility verification and existence and uniqueness of initial sets, and we relate the resulting bounds to the embedded AF case.

Worked examples with small BAFs show how serialisation sequences proceed in the presence of support. They also show why the induced indirect attack relation (via supported and secondary attack patterns) has to be recomputed after each reduct.

1.3 Thesis structure

Section 2 recalls AFs and BAFs, introduces deductive support and the coherence-based semantics used throughout the thesis, and summarises the AF serialisation framework.

Section 3 develops the serialisation theory for BAFs by introducing bipolar initial sets and reducts and by proving the serialisation characterisations for the coherent semantics.

Section 3.6 studies the computational complexity of coherent admissibility verification and of initial-set existence and uniqueness, and relates the results to the AF case.

Section 4 discusses the main findings and outlines directions for future work.

2 Background

2.1 Abstract Argumentation Frameworks (AFs)

Following Dung's presentation, let $\mathfrak{A}$ denote a universal set, where the elements $a \in \mathfrak{A}$ are called *arguments* [10].

Definition 2.1 (Abstract Argumentation Framework (AF))**.** *An* abstract argumentation framework *(AF) is a pair*

$$AF = (A, R)$$

where $A \subseteq \mathfrak{A}$ is a finite set of arguments *and $R \subseteq A \times A$ is an* attack relation*. The set of all abstract argumentation frameworks (over $\mathfrak{A}$) is denoted by $\mathfrak{F}_{\mathfrak{A}}$.*

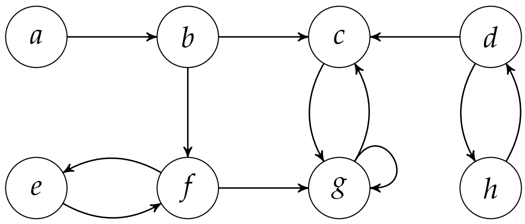

Figure 1: Example AF_0

Notation 2.1 (Attack Relations, Restrictions, Range and σ-Extensions)**.** *For $AF = (A, R)$ and two arguments $a, b \in A$, whenever $(a, b) \in R$ and the reference to the relation R is unambiguous, we say that a* attacks *b in AF by writing:*

$$a \overset{R}{\to} b \quad \Leftrightarrow \quad (a, b) \in R \quad \Leftrightarrow \quad aRb$$

If the relation is clear from context, we use $a \to b$.

A set $S \subseteq A$ attacks an argument $a \in A$ (symbolically: $S \to a$ or SRa) if any member $s \in S$ attacks a. An argument $a \in A$ attacks a set $S \subseteq A$ if a attacks some $s \in S$. We write $a \to S$ or aRS. For sets $S, T \subseteq A$, $S \to T$ or SRT denotes that there are $s \in S, t \in T$ such that $s \to t$ or sRt.

For $AF = (A, R)$ and $AF' = (A', R')$ we write $AF' \sqsubseteq AF$ iff $A' \subseteq A$ and $R' = R \cap (A' \times A')$. For $X \subseteq A$, the restriction *or* projection *of AF to X is defined as:*

$$AF|_X = (X, R \cap (X \times X))$$

which retains only arguments in X and attacks between them.

Given a set $S \subseteq A$, we denote $S^+ = \{a \in A \mid S \to a\}$ the set of all arguments attacked by S, while $S^- = \{a \in A \mid a \to S\}$ is the set of all attackers of S. Formally:

$$S^+ = \{a \in A \mid \exists b \in S : bRa\}, \qquad S^- = \{a \in A \mid \exists b \in S : aRb\}$$

For a set $S \subseteq A$, we denote

$$S^{\oplus} = S \cup S^{+}$$

as the range *of S, representing all arguments in S together with all arguments attacked by S.*

For a semantics σ, $\sigma(AF)$ denotes the set of σ-extensions of AF.

If S is a singleton set, we omit braces for readability, i.e. we write a^- and a^+ instead of $\{a\}^-$ and $\{a\}^+$, respectively. For two sets S and S' we write SRS' iff $S^+ \cap S' \neq \emptyset$.

Definition 2.2 (Conflict-Free). *Let $(A, R) \in \mathfrak{F}_{\mathfrak{A}}$ be an AF. A set $S \subseteq A$ is* conflict-free *iff there is no $a, b \in S$ such that $(a, b) \in R$ as in $a \to b$ or aRb, meaning no argument in S attacks another argument in S. Formally:*

$$\forall a, b \in S : (a, b) \notin R$$

Let $cf(AF)$ denote the set of all conflict-free sets of AF.

Definition 2.3 (Defence). *Let $AF = (A, R) \in \mathfrak{F}_{\mathfrak{A}}$. A set $D \subseteq A$* defends *a set $S \subseteq A$ iff:*

$$S^- \subseteq D^+$$

Equivalently:

$$\forall a \in S, \forall b \in A \quad (bRa \Rightarrow \exists d \in D : dRb)$$

Explicitly, we can say a set D *defends* an argument $a \in S$ iff for all $b \in A$ with bRa there is $d \in D$ with dRb or alternatively, D defends S iff for every attacker $b \in A$ with bRs $(b \to s)$ for some $s \in S$, we have DRb $(D \to b)$.

The following definition expands on the concept of defence.

Definition 2.4. *(Characteristic Function) Let $AF = (A, R) \in \mathfrak{F}_{\mathfrak{A}}$. The* characteristic function

$$\tau_{AF} : 2^A \to 2^A$$

defines, a set $S \subseteq A$, the arguments that are defended by S. Specifically:

$$\tau_{AF}(S) = \{a \in A \mid \{a\}^- \subseteq S^+\}$$

In particular, every set S also defends all arguments that have no attackers, since for such arguments a we have $a^- = \emptyset \subseteq S^+$. Consequently, $\tau_{AF}(S)$ always includes all arguments with no attackers, regardless of S.

2.2 Semantics in AFs

Building on this foundation, we now turn to the formalised semantics in AFs, where different mappings of frameworks to sets of arguments are introduced. These semantics help to identify acceptable sets of arguments, which can then be categorised into specific types of extensions [10].

We want an acceptable set of arguments to not only be conflict-free, but to also exclude external attacks on the set or at least have those attacks mitigated by a defence. Thus, the concepts of Definitions 2.2 and 2.3 are fundamental to the following Definition 2.5.

Definition 2.5 (Admissibility). *A set $S \subseteq A$ is admissible (ad), if it is conflict-free and defends all $a \in S$ against any attacks from within A. Let ad(AF) denote the set of admissible sets of AF.*

Note that the empty set $\emptyset$ is always admissible. In order to distinguish between different notions of acceptability, we now introduce *semantics* as mappings from AFs to sets of acceptable sets of arguments, called *extensions*.

Definition 2.6 (Semantics and Extensions). *A* semantics *is a mapping $\sigma : \mathfrak{F}_{\mathfrak{A}} \to 2^{2^{\mathfrak{A}}}$ such that for every $AF = (A, R) \in \mathfrak{F}_{\mathfrak{A}}$ we have $\sigma(AF) \subseteq 2^A$. For a given framework $AF \in \mathfrak{F}_{\mathfrak{A}}$, we call $\sigma(AF)$ the set of σ-*extensions *of AF.*

Different semantics can be phrased by imposing constraints on admissible sets and thus are given special names. Below, in Definition 2.7, we list some Dung-style standard semantics for AFs [10, 2], which we will refer to as *classical semantics*.

Definition 2.7 (AF Classical Semantics). *Let $AF = (A, R) \in \mathfrak{F}_{\mathfrak{A}}$. A set $E \subseteq A$ is called*

- complete *(co), iff $\forall a \in A$, if E defends a then $a \in E$,*
  i.e. it is admissible and includes all arguments that it defends. Thus $\tau_{AF}(E) = E$;
- grounded *(gr), iff E is the $\subseteq$-minimal complete extension of AF,*
  i.e. it is the smallest complete extension. Equivalently, $E = lfp(\tau_{AF})$, the least fixed point *of the characteristic function;*
- preferred *(pr), iff E is $\subseteq$-maximal among admissible sets,*
  i.e. it is the largest admissible set wrt. set inclusion, there is no larger admissible set E' such that $E \subset E'$;
- stable *(st), iff E is conflict-free and $E^{\oplus} = A$,*
  i.e. it attacks every argument in $A \setminus E$, leaving no argument status neutral or undecided.

For a semantics $\sigma \in \{ad, co, gr, pr, st\}$, we denote by $\sigma(AF)$ the set of all σ-extensions of AF.[1]

[1] Additional semantics, e.g. *semi-stable, ideal*, and *strongly admissible* [2] are beyond this thesis' scope.

These designations likewise extend to the corresponding semantics. For example, the semantics that assigns to an AF the set of its preferred extensions is referred to as the *preferred semantics*. All statements on minimality or maximality are meant to be wrt. set inclusion.

We say an argument of an AF is *credulously accepted* with regard to a semantics σ if it appears in at least one extension of $\sigma(AF)$, and that it is *sceptically accepted* if it appears in all extensions of $\sigma(AF)$.

Example 2.1 (AF Semantics). *For AF_0 in Figure 1, the corresponding extensions under the classical semantics of Definition 2.7 are*

$$\begin{aligned}
cf(AF_0) = \{&\{a\},\{b\},\{c\},\{d\},\{e\},\{f\},\{h\},\\
&\{a,c\},\{a,d\},\{a,e\},\{a,f\},\{a,h\},\{b,d\},\{b,e\},\{b,h\},\\
&\{c,e\},\{c,f\},\{c,h\},\{d,e\},\{d,f\},\{e,h\},\{f,h\},\\
&\{a,c,e\},\{a,c,f\},\{a,c,h\},\{a,d,e\},\{a,d,f\},\{a,e,h\},\{a,f,h\},\\
&\{b,d,e\},\{b,e,h\},\{c,e,h\},\{c,f,h\},\\
&\{a,c,e,h\},\{a,c,f,h\}\}
\end{aligned}$$

$$\begin{aligned}
ad(AF_0) = \{&\{a\},\{d\},\{e\},\{h\},\{a,d\},\{a,e\},\{a,f\},\{a,h\},\{d,e\},\{e,h\},\\
&\{a,c,h\},\{a,d,e\},\{a,d,f\},\{a,e,h\},\{a,f,h\},\{a,c,e,h\},\{a,c,f,h\}\}
\end{aligned}$$

$$\begin{aligned}
co(AF_0) = \{&\{a\},\{a,d\},\{a,e\},\{a,f\},\{a,h\},\{a,c,h\},\{a,d,e\},\{a,d,f\},\\
&\{a,e,h\},\{a,c,e,h\},\{a,c,f,h\}\}
\end{aligned}$$

$$gr(AF_0) = \{\{a\}\}$$

$$pr(AF_0) = \{\{a,d,e\},\{a,d,f\},\{a,c,e,h\},\{a,c,f,h\}\}$$

$$st(AF_0) = \{\{a,d,f\},\ \{a,c,e,h\},\ \{a,c,f,h\}\}$$

Example 2.2 (Informal Reading of AF_0). *The framework AF_0 in Figure 1 can be read, for instance, as a discussion about whether to adopt a certain policy a and how to address various objections and side effects:*

- *a: "We should adopt the policy."*
- *b: "The policy is too expensive."*
- *c: "The budget report shows the costs are unsustainable."*
- *d: "The budget report is outdated and misleading."*
- *e: "External funding will cover the main costs."*
- *f: "External funding has not been approved yet."*
- *g: "Past projects with similar funding have failed financially."*
- *h: "Independent auditors confirm the budget report is valid."*

The attack relation of AF_0 captures how statements can undercut each other. For example, a attacks b, and d attacks c, while d and h form a direct dispute. The extensions in Example 2.1 can then be read as coherent positions that defend their elements against attacks, e.g. stable extensions correspond to positions that defeat every incompatible argument.

2.3 Initial Sets in AFs

Non-empty minimal admissible sets have been coined *initial sets* by Xu and Cayrol [19, 17, 18].

Definition 2.8 (Initial Sets [19]). *For $AF = (A, R) \in \mathfrak{F}_{\mathfrak{A}}$, a nonempty subset $S \subseteq A$ is called an* initial set *of AF iff*

(i) S is admissible in AF, and

(ii) there is no nonempty admissible $S' \subsetneq S$.

Formally, iff

$$S \in ad(AF) \quad and \quad \nexists S' \subsetneq S \; \big(S' \neq \emptyset \wedge S' \in ad(AF)\big)$$

We write $\mathrm{IS}(AF)$ for the set of all initial sets of AF. While every initial set is admissible, not every credulously accepted argument wrt. admissibility necessarily appears in an initial set (as illustrated by argument b in Example 1 of [14]). Following Xu, Cayrol [19, 17, 18] and Thimm [14], we distinguish three types of initial sets:

Definition 2.9 (Classification of Initial Sets). *For $AF = (A, R)$ and $S \in IS(AF)$, we say that*

1. *S is* unattacked *iff $S^- = \emptyset$, i.e. S is not attacked by any argument in A;*
2. *S is* unchallenged *iff $S^- \neq \emptyset$ and $\nexists S' \in IS(AF)$ with $S' \neq S$ and $S'RS$, i.e. S is attacked but not by another initial set;*
3. *S is* challenged *iff $\exists S' \in IS(AF)$ with $S' \neq S$ and $S'RS$, i.e. S is attacked by another initial set.*

Note that $S'RS$ iff $S' \cap S^- \neq \emptyset$, since $S^- = \{\, a \in A \mid aRS \,\}$.

According to Thimm [14], only unattacked initial sets have been considered explicitly in prior work [18]. Moreover, every unattacked initial set S is necessarily a singleton $S = \{a\}$. The notions of unattacked, unchallenged, and challenged initial sets are mutually exclusive and exhaustive. In the following, we denote with $\mathrm{IS}^{\not\leftarrow}(AF)$, $\mathrm{IS}^{\not\leftrightarrow}(AF)$, and $\mathrm{IS}^{\leftrightarrow}(AF)$ the set of unattacked, unchallenged, and challenged initial sets, respectively. So we have $\mathrm{IS}(AF) = \mathrm{IS}^{\not\leftarrow}(AF) \cup \mathrm{IS}^{\not\leftrightarrow}(AF) \cup \mathrm{IS}^{\leftrightarrow}(AF)$. Moreover, for $S \in \mathrm{IS}(AF)$ let

$$\mathrm{conflicts}(S, AF) = \{S' \in \mathrm{IS}(AF) \mid S'RS\}$$

denote the set of conflicting initial sets of S, which is always empty in the case of unattacked and unchallenged initial sets. Also, SRS' implies $S'RS$ for any $S, S' \in$ $\text{IS}(AF)$ as S' is admissible and therefore defends itself.

Example 2.3 (Initial Sets). *Consider AF_0 in Figure 1. According to Definition 2.8, initial sets are the minimal non-empty admissible sets. For AF_0, we find four initial sets:*

$$IS(AF_0) = \{\{a\}, \{d\}, \{e\}, \{h\}\}$$

Each initial set resolves a distinct, atomic conflict in the framework, but no initial set itself constitutes a solution to the AF as a whole. Hence all three initial-set types from Definition 2.9 occur in AF_0:

- $\{a\}$ *is an* unattacked initial set*, as it is not attacked by any argument.*
- $\{e\}$ *is an* unchallenged initial set*, as it is attacked by f, which is not an initial set.*
- $\{d\}$ *and* $\{h\}$ *are* challenged initial sets*, as they attack each other and both are initial sets.*

Explicitly:

$$IS^{\not\leftarrow}(AF_0) = \{\{a\}\}, \quad IS^{\not\leftrightarrow}(AF_0) = \{\{e\}\}, \quad IS^{\leftrightarrow}(AF_0) = \{\{d\}, \{h\}\}$$

The conflicts among challenged initial sets are:

$$\mathit{conflicts}(\{d\}, AF_0) = \{\{h\}\}, \quad \mathit{conflicts}(\{h\}, AF_0) = \{\{d\}\}$$

Thimm also extends the analysis from [19, 17, 18] by exposing that initial are always completely contained in a single strongly connected component:

Definition 2.10 (Strongly Connected Component (SCC) [14]). *Let $AF = (A, R) \in \mathfrak{F}_{\mathfrak{A}}$. A subframework $AF' = (A', R')$ is a* strongly connected component *(SCC) of AF if $AF' \sqsubseteq AF$ such that*

1. *For all $a, b \in A'$, there exists a directed path from a to b in AF', and*
2. *There is no larger subframework AF'' with $AF' \sqsubseteq AF'' \sqsubseteq AF$ satisfying property (1).*

Let $SCC(AF)$ denote the set of all strongly connected components of AF.

This means, a SCC is a maximal set of arguments where every argument can reach every other via some directed sequence of attacks.

Example 2.4 (SCC Decomposition). *For AF_0 in Figure 1, the SCCs are*

$$SCC(AF_0) = \{AF_0|_{\{a\}}, AF_0|_{\{b\}}, AF_0|_{\{c,g\}}, AF_0|_{\{d,h\}}, AF_0|_{\{e,f\}}\}$$

Each initial set is fully contained in a single SCC:

- $\{a\}$ *in* $AF_0|_{\{a\}}$,
- $\{e\}$ *in* $AF_0|_{\{e,f\}}$,
- $\{d\}$ *and* $\{h\}$ *in* $AF_0|_{\{d,h\}}$

This is consistent with Proposition 2.1.

Proposition 2.1 (Initial Sets Within SCCs [14]). *If S is an initial set of AF, then there exists $AF' = (A', R') \in SCC(AF)$ such that $S \subseteq A'$.*

If S is an initial set, denote by $\mathrm{SCC}(S)$ its SCC as defined in Proposition 2.1. Initial sets can be further characterised as follows.

Proposition 2.2 (Characterisation via SCC [14]). *S is an initial set of AF iff S is an initial set of $SCC(S) = (A', R')$ and $S^- \subseteq A'$.*

This means, S is an initial set iff it is an initial set of a single SCC and it is not attacked by arguments outside that SCC. Finally, Thimm observes useful properties regarding the types of initial sets:

Proposition 2.3 (Types of Initial Sets [14]). *Let $S \in IS(AF)$ and $SCC(S) = (A', R')$.*

1. *If S is unattacked then $|A'| = 1$.*
2. *If S is challenged or unchallenged then $|A'| > 1$.*
3. *If S is challenged and $S' \in conflicts(S, AF)$, then $SCC(S) = SCC(S')$.*

In particular, the third property shows that conflicts between initial sets always occur within a single SCC.

2.4 Serialisability in AFs

We revisit the concept of *serialisability* [14, 5], a property of argumentation semantics enabling a constructive characterisation of extensions. This approach relies on the reduct of an argument set in an argumentation framework [3]. So, to build extensions, we recall the definition of the reduct.

Definition 2.11 (Reduct [14]). *Let $AF = (A, R)$ and $S \subseteq A$. The S-*reduct *AF^S is defined as*

$$AF^S := AF|_{A \setminus S^{\oplus}}$$

Building from the reduct, we define the *serialisation sequence* $\mathcal{S}$ as a decomposition of an extension into a series of initial sets.

Definition 2.12 (Serialisation Sequence). *A* serialisation sequence *for $AF = (A, R) \in \mathfrak{F}_{\mathfrak{A}}$ is a sequence $\mathcal{S} = (S_1, \ldots, S_n)$ where*

- $S_1 \in IS(AF)$, *and*
- *for each* $2 \leq i \leq n$*:* $S_i \in IS(AF^{E_{i-1}})$, *where* $E_{i-1} := S_1 \cup \cdots \cup S_{i-1}$.

The set $E := S_1 \cup \cdots \cup S_n$ *is called the* result *of the serialisation sequence* $\mathcal{S}$*. We write* $\Sigma(E)$ *to denote that there exists a serialisation sequence* $(S_1, \ldots, S_n)$ *whose result is E.*

Definition 2.12 translates, that a serialisation sequence is a way of breaking down a set of arguments, called an extension, into smaller, sequential sets (namely initial sets) based on a given AF.

Remark 2.1 (Notation for partial results). *For a serialisation sequence* $\mathcal{S} = (S_1, \ldots, S_n)$*, we use* $E_i := S_1 \cup \cdots \cup S_i$ *to denote the union of the first i steps, with* $E_0 := \emptyset$ *by convention.*

A serialisation sequence $(S_1, \ldots, S_n)$ produces an admissible set $E = S_1 \cup \cdots \cup S_n$, and for every admissible set there exists at least one such sequence [5]. By considering the distinctions between initial set types from Definition 2.9, we can refine serialisation sequences to characterise a broad spectrum of admissibility-based semantics [14].

Theorem 2.1 (Serialisation Characterisation). *Let* $AF = (A, R) \in \mathfrak{F}_{\mathfrak{A}}$ *and* $E \subseteq A$*. Then E is a* σ*-extension iff there exists a serialisation sequence* $\mathcal{S} = (S_1, \ldots, S_n)$ *with result E satisfying*

$$
\begin{aligned}
\sigma = ad: &\quad \textit{no additional condition,}\\
\sigma = co: &\quad IS^{\not\leftarrow}(AF^E) = \emptyset,\\
\sigma = gr: &\quad \forall i \in \{1, \ldots, n\}: S_i \in IS^{\not\leftarrow}(AF^{E_{i-1}})\\
&\quad \textit{and } IS^{\not\leftarrow}(AF^E) = \emptyset,\\
\sigma = pr: &\quad IS(AF^E) = \emptyset,\\
\sigma = st: &\quad AF^E = (\emptyset, \emptyset),
\end{aligned}
$$

where $E_i = S_1 \cup \cdots \cup S_i$ *for* $i = 0, \ldots, n$ *(with* $E_0 = \emptyset$*).*

Proof. See [14] for the characterisations of {ad, co, gr, pr, st}. The conditions in the table above correspond to Thimm's termination criteria β_σ. While the selection and termination functions in [14] operate on the current state (AF, S) of the transition system, the characterisation above evaluates the equivalent property on the final E-reduct AF^E [5, 4]. □

We denote by $\Sigma = \{\mathrm{ad}, \mathrm{co}, \mathrm{gr}, \mathrm{pr}, \mathrm{st}\}$ the set of *serialisable semantics*. For an AF, we write $\mathfrak{S}_\sigma(AF)$ to denote the set of all serialisation sequences of AF wrt. the semantics $\sigma \in \Sigma$.

To formalise the construction of extensions through serialisation, we introduce the concepts of selection and termination functions [14].

Definition 2.13 (State). *A state T is a tuple $T = (AF, S)$ with $AF \in \mathfrak{F}_{\mathfrak{A}}$ and $S \subseteq \mathfrak{A}$.*

Definition 2.14 (Selection Function). *A selection function α is any function $\alpha : 2^{\mathfrak{A}} \times 2^{\mathfrak{A}} \times 2^{\mathfrak{A}} \to 2^{\mathfrak{A}}$ with $\alpha(X, Y, Z) \subseteq X \cup Y \cup Z$ for all $X, Y, Z \subseteq \mathfrak{A}$.*

We apply a selection function α in the form $\alpha(\mathrm{IS}^{\not\leftarrow}(\mathrm{AF}), \mathrm{IS}^{\not\leftrightarrow}(\mathrm{AF}), \mathrm{IS}^{\leftrightarrow}(\mathrm{AF}))$ to select a subset of initial sets as eligible for the construction process.

Definition 2.15 (Termination Function). *A termination function β is any function $\beta : \mathfrak{F}_{\mathfrak{A}} \times 2^{\mathfrak{A}} \to \{0, 1\}$.*

A termination function β indicates when a construction of an admissible set is finished, which occurs when $\beta(\mathrm{AF}, S) = 1$.

For some selection function α, consider the following transition rule:

$$(\mathrm{AF}, S) \xrightarrow{S' \in \alpha(\mathrm{IS}^{\not\leftarrow}(\mathrm{AF}), \mathrm{IS}^{\not\leftrightarrow}(\mathrm{AF}), \mathrm{IS}^{\leftrightarrow}(\mathrm{AF}))} (\mathrm{AF}^{S'}, S \cup S')$$

If (AF', S') can be reached from (AF, S) via a finite number of steps with the above rule, we write $(\mathrm{AF}, S) \rightsquigarrow_{\alpha} (\mathrm{AF}', S')$. If, in addition, the state (AF', S') satisfies the termination criterion of β, i.e. $\beta(\mathrm{AF}', S) = 1$, then we write $(\mathrm{AF}, S) \rightsquigarrow_{\alpha,\beta} (\mathrm{AF}', S')$.

Given concrete instances of α and β, let $E_{\alpha,\beta}(\mathrm{AF})$ be the set of all S with $(\mathrm{AF}, \emptyset) \rightsquigarrow_{\alpha,\beta} (\mathrm{AF}', S)$ for some AF'.

Definition 2.16 (Serialisability). *A semantics σ is* serialisable *with a selection function α and a termination function β if $\sigma(AF) = E_{\alpha,\beta}(AF)$ for all $AF \in \mathfrak{F}_{\mathfrak{A}}$.*

Theorem 2.2 (Serialisability of AF Classical Semantics [14, Thms. 2–7]). *The following pairs of selection and termination functions characterise the AF classical semantics from Definition 2.7:*

- ***Admissible:*** $\alpha_{\mathrm{ad}}(X, Y, Z) = X \cup Y \cup Z$ *and* $\beta_{\mathrm{ad}}(\mathrm{AF}, S) = 1$
- ***Complete:*** α_{ad} *and* $\beta_{\mathrm{co}}(\mathrm{AF}, S) = \begin{cases} 1 & \text{if } \mathrm{IS}^{\not\leftarrow}(\mathrm{AF}) = \emptyset \\ 0 & \text{otherwise} \end{cases}$
- ***Grounded:*** $\alpha_{\mathrm{gr}}(X, Y, Z) = X$ *and* β_{co}
- ***Preferred:*** α_{ad} *and* $\beta_{\mathrm{pr}}(\mathrm{AF}, S) = \begin{cases} 1 & \text{if } \mathrm{IS}(\mathrm{AF}) = \emptyset \\ 0 & \text{otherwise} \end{cases}$
- ***Stable:*** α_{ad} *and* $\beta_{\mathrm{st}}(\mathrm{AF}, S) = \begin{cases} 1 & \text{if } \mathrm{AF} = (\emptyset, \emptyset) \\ 0 & \text{otherwise} \end{cases}$

These functions define the construction rules and stopping criteria for serialisation sequences that capture each semantics. See [14] for exact proofs and technical details. For an application of these constructions to ranking semantics, see [5].

Example 2.5 (Serialisation in AFs). *Consider AF_0 in Figure 1 with arguments $A = \{a,b,c,d,e,f,g,h\}$ and attack relation*

$$R = \{(a,b),\ (b,c),\ (b,f),\ (c,g),\ (d,c),\ (d,h),\ (e,f),\ (f,e),\ (f,g),\ (g,c),\ (g,g),\ (h,d)\}.$$

The initial sets of AF_0 are $\{a\},\{d\},\{e\},\{h\}$.

We recall a serialisation sequence that yields the preferred extension $E = \{a,d,e\}$. Let $S_1 = \{a\}$, $S_2 = \{d\}$, $S_3 = \{e\}$ and define $E_i = \bigcup_{j=1}^{i} S_j$ (with $E_0 = \emptyset$).

1. *Choose $S_1 = \{a\} \in IS_{AF_0}$. Then $S_1^+ = \{b\}$ and*

$$A_1 = A \setminus S_1^{\oplus} = \{c,d,e,f,g,h\}$$

2. *In the reduct $AF_0^{E_1}$ choose $S_2 = \{d\} \in IS_{AF_0^{E_1}}$. Here $S_2^+ = \{c,h\}$ and*

$$A_2 = A_1 \setminus S_2^{\oplus} = \{e,f,g\}$$

3. *In the reduct $AF_0^{E_2}$ choose $S_3 = \{e\} \in IS_{AF_0^{E_2}}$. Here $S_3^+ = \{f\}$ and*

$$A_3 = A_2 \setminus S_3^{\oplus} = \{g\}$$

4. *Let $E = E_3 = \{a,d,e\}$. Then $IS_{AF_0^E} = \emptyset$. Hence $\mathcal{S} = (S_1,S_2,S_3)$ is a serialisation sequence with result E.*

For preferred semantics, maximal admissible sets can be constructed by freely choosing any initial set at each reduction step until all are exhausted.

For grounded semantics, only unattacked initial sets may be picked at each stage, potentially forcing the process to terminate early.

Example 2.6 (Serialisation to the Grounded Extension in AF_0). *We recall that the grounded extension of AF_0 is $\{a\}$.*

1. *Initial sets in AF_0 are $\{a\},\{d\},\{e\},\{h\}$. For grounded semantics, only unattacked initial sets can be chosen. Here, $\{a\}$ is the unique unattacked initial set ($IS^{\not\leftarrow}(AF_0) = \{\{a\}\}$). Select $S_1 = \{a\}$. Remove a and its attacked arguments: $S_1^+ = \{b\}$. $A_1 = A \setminus (\{a,b\}) = \{c,d,e,f,g,h\}$.*
2. *In $AF_0^{\{a\}}$, unattacked initial sets: none. Therefore, the sequence terminates. The grounded extension is $\{a\}$.*

Thus only a can be selected, and the construction terminates in the reduct $AF_0^{\{a\}}$ because no unattacked initial set remains, i.e. $IS(AF_0^{\{a\}}) = \emptyset$.

In AF_0 from Figure 1, the serialisations for the complete and preferred semantics coincide for the extension a, d, e, as maximal admissible sets contain only unchallenged or unattacked initial sets at each step.

For most frameworks, the serialisability structure is reflected in all classical semantics, but specific AFs may have essential differences for complete, stable, or other extensions depending on attack cycles or singleton SCCs.

2.5 Bipolar Argumentation Frameworks (BAFs)

Bipolar Argumentation Frameworks (BAFs) extend Dung's Abstract Argumentation Frameworks by introducing an additional positive *support relation* alongside the traditional negative *attack relation*. This enables a richer and more expressive representation of both conflicting and collaborative argumentative dynamics [12, 16, 1, 7]. While the term "bipolar" suggests a simple positive/negative polarity distinction, BAFs actually allow for much richer structures through the interplay of attacks and supports. The addition of support relations enables BAFs to capture not only direct opposition but also cooperative reinforcement, making them particularly suitable for domains such as legal reasoning, decision support systems, and structured debate.

Definition 2.17 (Bipolar Argumentation Framework (BAF))**.** *A* bipolar argumentation framework *(BAF) is a triple*

$$B = (A, R, U)$$

where $A \subseteq \mathfrak{A}$ is a finite set of arguments*, $R \subseteq A \times A$ is the* attack relation*, and $U \subseteq A \times A$ is the* support relation*, such that $R \cap U = \emptyset$. The set of all bipolar argumentation frameworks (over $\mathfrak{A}$) is denoted by $\mathfrak{B}_{\mathfrak{A}}$.*

Notation 2.2 (Attack and Support Relations)**.** *Let $B = (A, R, U)$ be a bipolar argumentation framework.*

For $a, b \in A$:

$$a \xrightarrow{R} b \quad \Leftrightarrow \quad (a, b) \in R \quad \Leftrightarrow \quad aRb$$

$$a \xrightarrow{U} b \quad \Leftrightarrow \quad (a, b) \in U \quad \Leftrightarrow \quad aUb$$

Chains such as $a \xrightarrow{U} b \xrightarrow{U} c \xrightarrow{R} d$ are shorthand for the relational notation $(a, b) \in U$, $(b, c) \in U$, and $(c, d) \in R$.

The arrow notation is used in diagrams and informal paths to clarify relation types. In formal definitions and proofs, we use the relational notation.

Since a BAF includes both support and attack relations, Cayrol et al. [7, 6] define the concepts of *supported attack* and *secondary attack*, capturing interactions that combine chains of supports with a direct attack. Following [6], we define supported

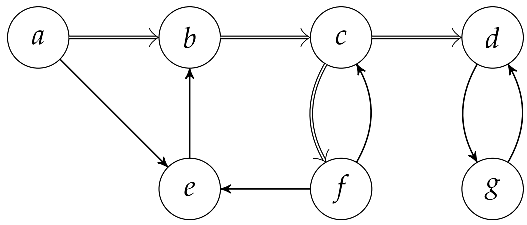

Figure 2: Example BAF B_0

and secondary attacks using relational composition (Definition 2.20). We express these definitions using both relational composition and explicit sequential paths for conciseness and clarity.

Remark 2.2 (Interpretation of Support)**.** *Following Cayrol et al. [7, 6], the support relation U in a BAF can be given different semantic interpretations:*

- ***Deductive support****: aUb means acceptance of a implies acceptance of b (if a is accepted, then b must be accepted);*
- ***Necessary support****: aUb means acceptance of a is necessary for b to be accepted (if b is accepted, then a must be accepted);*
- ***Evidential support****: Prima-facie arguments (self-supported) vs. standard arguments (requiring evidential grounding).*

Unless stated otherwise, this work adopts the deductive support *interpretation, which is standard in bipolar argumentation semantics [7, 9].*

Definition 2.18 (Relational Composition [13])**.** *Let $\rho_1, \rho_2 \subseteq A \times A$ be binary relations. The* composition *$\rho_1 \circ \rho_2$ is defined as:*

$$\rho_1 \circ \rho_2 := \{(a, c) \in A \times A \mid \exists b \in A : (a, b) \in \rho_1 \text{ and } (b, c) \in \rho_2\}$$

Intuitively, $(a, c) \in \rho_1 \circ \rho_2$ means there exists an intermediate argument b such that a is related to b via ρ_1 and b is related to c via ρ_2.

Definition 2.19 (Reflexive–Transitive Closure [13])**.** *Let $\rho \subseteq A \times A$ be a binary relation.*

- *The* transitive closure *ρ^+ is the smallest transitive relation containing ρ.*
- *The* reflexive–transitive closure *ρ^* is the smallest reflexive and transitive relation containing ρ.*

Equivalently, $(a, b) \in \rho^*$ iff either $a = b$ or there exists a sequence $a = a_1, \ldots, a_n = b$ with $n \geq 2$ such that $(a_i, a_{i+1}) \in \rho$ for all $1 \leq i \leq n - 1$.

Applying these closure operators to the support relation U, we obtain U^+ and U^*. We say that there is a (possibly empty) support chain from S to b iff there exists $s \in S$ such that $(s, b) \in U^*$, and a non-empty support chain iff there exists $s \in S$ such that $(s, b) \in U^+$.

Definition 2.20 (Supported and Secondary Attack [7, 6]). *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and $a, b \in A$.*

An argument a support-attacks *b iff there exist arguments $a_1, \ldots, a_n \in A$ with $n \geq 0$ such that*

(i) $aUa_1, a_1Ua_2, \ldots, a_{n-1}Ua_n$ (empty if $n = 0$), and

(ii) a_nRb (with $a_n := a$ when $n = 0$).

The set of all supported attacks *is*

$$R_u^B := U^* \circ R$$

Note that $R \subseteq R_u^B$ (direct attacks are support-attacks with $n = 0$).

An argument a secondarily attacks *argument b iff there exist arguments $a_1, \ldots, a_n \in A$ ($n \geq 1$) such that*

(i) aRa_1, and

(ii) $a_1Ua_2, \ldots, a_{n-1}Ua_n, a_nUb$.

The set of all secondary attacks *is denoted*

$$R_m^B := R \circ U^+$$

Note that $R \cap R_m^B = \emptyset$ (direct attacks are excluded).

Intuitively, supported attacks are (possibly empty) chains of supports ending in an attack, while secondary attacks are attacks followed by non-empty chains of supports.

Remark 2.3 (Union of Derived Attacks and Historical Context). *We often work with the union $R_u^B \cup R_m^B$ of supported and secondary attacks, since many notions depend only on the induced derived attack relation. Nevertheless, the compositional forms $R_u^B = U^* \circ R$ and $R_m^B = R \circ U^+$ are convenient when computing derived attacks in restricted frameworks.*

Supported and secondary attacks were introduced by Cayrol and Lagasquie-Schiex [7]. In our notation, u refers to support-based attacks and m to secondary attacks, historically so-called mediated attacks; see [6] for a survey.

Notation 2.3 (Derived Attacks and BAF Attack/Defence Sets). *For a BAF $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$, we write*

$$R_{\mathrm{der}}^B := R_u^B \cup R_m^B$$

for the union of all derived attacks (supported and secondary).

For $S \subseteq A$, we lift the AF notation from Section 2.1 to BAFs by replacing the attack relation R with the derived attack relation R_{der}^B:

- $S_B^+ = \{a \in A \mid \exists s \in S : (s, a) \in R_{\mathrm{der}}^B\}$ *denotes the set of arguments attacked by S;*
- $S_B^- = \{a \in A \mid \exists s \in S : (a, s) \in R_{\mathrm{der}}^B\}$ *denotes the set of arguments attacking S.*

For a single argument $a \in A$, we write a_B^+ and a_B^- for $\{a\}_B^+$ and $\{a\}_B^-$, respectively. When the framework B is clear from context, we may omit the subscript and write S^+ and S^-.

2.6 Semantics Based on Strong Conflict-Freeness and Coherence

We now define semantics for bipolar argumentation frameworks based on the foundational work of Cayrol and Lagasquie-Schiex [7]. The key idea is to extend Dung's classical semantics by accounting for both the attack and support relations through the derived notions of supported and secondary attacks.

Definition 2.21 (Strong Conflict-Freeness [7]). *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and $S \subseteq A$. The set S is* strongly conflict-free *iff there are no arguments $a, b \in S$ such that $(a, b) \in R^B_{\text{der}}$. We denote by b-cf$(B)$ the set of all strongly conflict-free sets in B.*

Remark 2.4 (Relationship to Classical Conflict-Freeness). *Strong conflict-freeness generalises classical conflict-freeness from Dung's frameworks. For any abstract argumentation framework $AF = (A, R)$, a set $S \subseteq A$ is conflict-free in AF if and only if S is strongly conflict-free in $(A, R, \emptyset)$. However, in general bipolar frameworks, strong conflict-freeness is a strictly stronger requirement than classical conflict-freeness (i.e. requiring only no direct attacks between arguments in S). To differentiate semantic concepts for BAFs from those for AFs, we denote a "b-" in the prefix.*

Cayrol et al. [7] define internal and external coherence. While the former is depicted by strong conflict-freeness, the latter is captured by the notion of safe sets and ensures that the considered set cannot simultaneously defeat and support a same argument. We summarise coherence as follows:

Definition 2.22 (Coherence [7]). *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and $S \subseteq A$. The set S is* coherent *iff there exists no argument $b \in A$ such that both*

- *(i) $(a, b) \in R^B_{\text{der}}$ for some $a \in S$, and*
- *(ii) Either $b \in S$, or there exist arguments $a_1, \dots, a_n \in A$ ($n \geq 1$) with $a_1 \in S$, $a_1 U a_2, \dots, a_{n-1} U a_n$, and $a_n U b$.*

We denote by b-coh(B) the set of all coherent sets in B.

Intuitively, a set S is coherent if there is no argument $b \in A$ such that S attacks b via the derived attacks R^B_{der} and either $b \in S$ or there exists $s \in S$ such that $(s, b) \in U^+$. This captures the idea that an acceptable set should not simultaneously attack and (directly or indirectly) support the same argument.

Proposition 2.4 (Coherence Implies Strong Conflict-Freeness). *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and $S \subseteq A$. If S is coherent, then S is strongly conflict-free.*

Proof. Suppose S is coherent. Assume for contradiction that S is not strongly conflict-free. Then there exist $a, b \in S$ with $(a, b) \in R^B_{\text{der}}$. Taking the case where $b \in S$ in Definition 2.22 (ii), we have that b is both attacked by S (via a) and is in S, violating coherence. Contradiction. □

Definition 2.23 (Strong Defence [7]). *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$, $S \subseteq A$, and $a \in A$. The set S* strongly defends *argument a iff for all $b \in A$ with $(b, a) \in R^B_{\text{der}}$, there exists $c \in S$ such that $(c, b) \in R^B_{\text{der}}$.*

Definition 2.24 (Coherent Admissibility [7]). *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and $S \subseteq A$. The set S is* coherently admissible *iff*

(i) S is coherent, and

(ii) S strongly defends all its elements.

We denote by b-ad(B) the set of all coherently admissible sets in B. This means, S is coherently admissible, if S has no internal or external conflicts and every element from S is defended by S, while defence may be direct or indirect.

Proposition 2.5 (Basic Properties of Coherent Admissibility). *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$. Then*

1. *$\emptyset \in$ b-ad(B) (the empty set is always coherently admissible).*
2. *If $U = \emptyset$ and $AF = (A, R)$, then b-ad$(B) =$ ad(AF) (coherent admissibility reduces to classical admissibility for standard AFs).*

Proof. (1): The empty set attacks and supports no arguments, hence is trivially coherent. It defends all its elements vacuously.

(2): When $U = \emptyset$, we have $R^B_u = R$ and $R^B_m = \emptyset$, so $R^B_{\text{der}} = R$. Strong conflict-freeness becomes classical conflict-freeness, coherence reduces to conflict-freeness (no support chains exist), and strong defence becomes classical defence. □

Definition 2.25 (Characteristic Function for BAFs). *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$. The characteristic function $\tau_B : 2^A \to 2^A$ is defined, for $S \subseteq A$, by*

$$\tau_B(S) = \{a \in A \mid a^-_B \subseteq S^+_B\}$$

i.e. $\tau_B(S)$ contains exactly those arguments that are strongly defended by S wrt. the derived attacks R^B_{der}.

The following bipolar semantics mirror the classical AF semantics of Definition 2.7, obtained by replacing admissibility with coherent admissibility and direct attacks R with the derived attacks R^B_{der}.

Definition 2.26 (BAF Classical Semantics)**.** *Let* $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$*. A coherently admissible set* $E \subseteq A$ *is called*

- coherently complete *(b-co) iff* $E \in$ *b-ad*(B) *and* $E = \tau_B(E)$*,*
  i.e. E includes all arguments it strongly defends;
- coherently grounded *(b-gr) iff E is the* $\subseteq$*-minimal element of b-co*(B)*.*
  Equivalently, $E = lfp(\tau_B)$*;*
- coherently preferred *(b-pr) iff E is* $\subseteq$*-maximal in b-ad*(B)*;*
- coherently stable *(b-st) iff* $E \in$ *b-ad*(B) *and* $E \cup E_B^+ = A$*,*
  leaving no non-member of E unattacked via derived attacks.

For a semantics $\sigma \in \{b\text{-}ad, b\text{-}co, b\text{-}gr, b\text{-}pr, b\text{-}st\}$*, we denote by* $\sigma(B)$ *the set of all* σ*-extensions of B.*

Proposition 2.6 (Reduction to Classical Semantics)**.** *Let* $B = (A, R, \emptyset) \in \mathfrak{B}_{\mathfrak{A}}$ *be a BAF with empty support relation, and let* $AF = (A, R)$ *be the corresponding abstract argumentation framework. Then for all semantics* $\sigma \in \{ad, co, gr, pr, st\}$*:*

$$b\text{-}\sigma(B) = \sigma(AF)$$

That is, BAF semantics coincide with classical AF semantics when no support relations are present, i.e. BAFs generalise AFs without altering their behaviour.

Proof. When $U = \emptyset$, we have $R_u^B = R$ and $R_m^B = \emptyset$, thus $R_{\text{der}}^B = R$. Therefore, coherence reduces to conflict-freeness, strong defence reduces to classical defence, and $\tau_B = \tau_{AF}$. The result follows from the definitions. □

Example 2.7 (Coherence-Based Semantics in BAFs)**.** *For BAF* B_0 *in Figure 2, the main coherence-based semantics as in Definition 2.26 are*

$\boldsymbol{b\text{-}cf}(B_0) = \{\, \emptyset, \{a\}, \{b\}, \{d\}, \{e\}, \{g\},$
$\{a,b\}, \{a,d\}, \{b,d\}, \{e,g\}, \{a,b,d\} \,\}$

(All subsets $S \subseteq A$ *such that for every* $a, b \in S$*, a does not attack b via derived attacks (supported or secondary).)*

$\boldsymbol{b\text{-}coh}(B_0) = \{\, \emptyset, \{d\}, \{e\}, \{g\}, \{e,g\} \,\}$

(No argument in these sets is both attacked and supported by the set itself, directly or via support chains.)

$\boldsymbol{b\text{-}ad}(B_0) = \{\emptyset\}$

(For this graph, only the empty set is coherently admissible. No non-empty subset defends itself against all attackers.)

3 BAF Serialisation

This section presents our contributions. We extend the concept of initial sets from Section 2.3 to BAFs by replacing classical admissibility with coherent admissibility from Section 2.6, and we lift the notion of serialisability from Section 2.4 to this setting. Moreover, we extend the reduct-based view underlying serialisability in AFs [14, 5] to BAFs, taking into account that removing arguments can disrupt support chains. We first introduce initial sets and reducts for BAFs, then study their properties, and finally develop a serialisation-based characterisation of coherent admissibility and related semantics, including computational aspects and the relationship to classical AFs.

3.1 Initial Sets in BAFs

The fundamental building blocks of serialisation are initial sets and reducts. In this subsection, we generalise initial sets from AFs to BAFs by replacing classical admissibility with coherent admissibility.

Definition 3.1 (Initial Sets in BAFs). *For $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$, a nonempty subset $S \subseteq A$ is called an* initial set *of B iff*

(i) S is coherently admissible in B ($S \in b\text{-}ad(B)$), and

(ii) there is no nonempty coherently admissible $S' \subsetneq S$.

Formally, iff

$$S \in b\text{-}ad(B) \quad \text{and} \quad \nexists S' \subsetneq S \left(S' \neq \emptyset \wedge S' \in b\text{-}ad(B)\right)$$

We write $IS(B)$ for the set of all initial sets of B.

Remark 3.1 (Design Choice). *An alternative characterisation of initial sets using only R and U directly (without explicit reference to R^B_{der} first) was considered. However, coherence verification requires tracking support chains regardless, offering no simplification. The R^B_{der}-based approach aligns with standard definitions [6] and enables direct adaptation of Thimm's AF framework [14] via the systematic translation pattern established in Table 2.*

Example 3.1 (Initial Sets in a BAF). *Consider $B_1 = (A, R, U)$ with $A = \{a, b, c, d\}$, $R = \{(b, c), (d, c)\}$ and $U = \{(a, b)\}$, as shown in Figure 3. We have*

- *$R^{B_1}_u = \{(a, c), (b, c), (d, c)\}$ (includes supported attack from a to c via b)*
- *$R^{B_1}_m = \emptyset$ (no attacks followed by supports)*
- *Thus, $R^{B_1}_{\text{der}} = R^{B_1}_u \cup R^{B_1}_m = R^{B_1}_u$*

The initial sets are $IS(B_1) = \{\{a\}, \{b\}, \{d\}\}$. Each of the unattacked singleton $\{a\}, \{b\}, \{d\}$ is coherently admissible. They have no attackers (so strong defence is vacuously satisfied) *and they cannot violate coherence in this framework. Minimality is immediate for singletons. In contrast,* $\{c\}$ *is not coherently admissible because it is attacked by* a, b, *and* d *and cannot be defended.*

Note that although a *supports* b, *this does not prevent* $\{b\}$ *from being an initial set. Coherence (Definition 2.22) is not a support-closure condition. It only forbids sets that simultaneously (derived-)attack and non-trivially support the same argument. Here,* $\{b\}$ *has no non-empty support chain to any argument it attacks via* $R^{B_1}_{\mathrm{der}}$, *so it is coherent.*

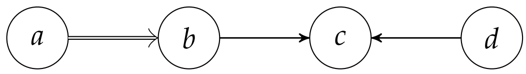

Figure 3: BAF B_1 from Example 3.1.

To refine the constructive process, we distinguish three types of initial sets based on their relationship with other minimal coherently admissible sets.

Definition 3.2 (Classification of Initial Sets in BAFs)**.** *For* $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ *and* $S \in IS(B)$, *let*

$$S^-_B = \{a \in A \mid \exists s \in S : (a, s) \in R^B_{\mathrm{der}}\}$$

denote the set of arguments that attack S *via derived attacks. We say that* S *is*

1. unattacked *iff* $S^-_B = \emptyset$;
2. unchallenged *iff* $S^-_B \neq \emptyset$ *and* $\nexists S' \in IS(B) \setminus \{S\}$ *with* $S' \cap S^-_B \neq \emptyset$, *i.e. none of the attackers of* S *occur in any other initial set;*
3. challenged *iff* $\exists S' \in IS(B) \setminus \{S\}$ *with* $S' \cap S^-_B \neq \emptyset$, *i.e. some other initial set contains an argument that attacks* S *via* R^B_{der}.

Equivalently, $S' \cap S^-_B \neq \emptyset$ *iff* $\exists s' \in S'\ \exists s \in S :\ (s', s) \in R^B_{\mathrm{der}}$.

The notions of unattacked, unchallenged, and challenged initial sets are mutually exclusive and exhaustive, forming a partition of the set of all initial sets $\mathrm{IS}(B)$. In the following, we denote by $\mathrm{IS}^{\not\leftarrow}(B)$, $\mathrm{IS}^{\not\leftrightarrow}(B)$, and $\mathrm{IS}^{\leftrightarrow}(B)$ the set of unattacked, unchallenged, and challenged initial sets in B, respectively. Thus, we have the partition

$$\mathrm{IS}(B) = \mathrm{IS}^{\not\leftarrow}(B)\ \cup\ \mathrm{IS}^{\not\leftrightarrow}(B)\ \cup\ \mathrm{IS}^{\leftrightarrow}(B)$$

where the union is pairwise disjoint.

Remark 3.2 (Unattacked Singletons as Initial Sets)**.** *If* $a \in A$ *is unattacked in* B *and* $\{a\}$ *is coherently admissible, then* $\{a\} \in IS^{\not\leftarrow}(B)$. *Indeed,* $\{a\}$ *is non-empty, coherently admissible, unattacked, and any proper subset is empty and thus not an initial set.*

Example 3.2 (Classification of Initial Sets in BAFs)**.** *Consider* $B_2 = (A, R, U)$ *with* $A = \{a, b, c, d, e\}$, $R = \{(b, c), (d, e)\}$, *and* $U = \{(a, b), (c, d)\}$, *as shown in Figure 4.*

The derived relations are $R_u^{B_2} = \{(a,c),(b,c),(c,e),(d,e)\}$ *and* $R_m^{B_2} = \{(b,d)\}$. *For each* $S \in IS(B_2)$*, the attacker set* $S_{B_2}^-$ *is computed wrt. the derived attacks* $R_{\text{der}}^{B_2}$.

The initial sets $IS(B_2) = \{\{a\},\{c\},\{d\}\}$ *are classified as follows:*

- $\{a\}$ *is* unattacked *since* $a_{B_2}^- = \emptyset$.
- $\{c\}$ *is* challenged *since* $c_{B_2}^- = \{a,b\}$ *and the initial set* $\{a\}$ *attacks it, i.e.* $\exists S' \in IS(B_2) \setminus \{\{c\}\}$ *with* $S' \cap c_{B_2}^- \neq \emptyset$.
- $\{d\}$ *is* unchallenged *since it is attacked by b, but b is not contained in any initial set, i.e.* $d_{B_2}^- = \{b\}$ *(via* $R_m^{B_2}$*) and* $\nexists S' \in IS(B_2) \setminus \{\{d\}\}$ *with* $S' \cap d_{B_2}^- \neq \emptyset$.

Thus

$$IS^{\not\leftarrow}(B_2) = \{\{a\}\}, \quad IS^{\not\leftrightarrow}(B_2) = \{\{d\}\}, \quad IS^{\leftrightarrow}(B_2) = \{\{c\}\}.$$

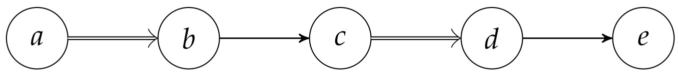

Figure 4: BAF B_2 from Example 3.2.

3.2 Reducts and Projections in BAFs

A key ingredient for extending serialisability from AFs to BAFs [14, 5] is an appropriate notion of reduct. In the bipolar setting, removing arguments may break support chains, so the BAF reduct must take supported and secondary attacks into account. We now define this reduct and illustrate its effect on derived attacks.

Definition 3.3 (Reduct for BAFs). *For* $B = (A,R,U) \in \mathfrak{B}_{\mathfrak{A}}$ *and* $S \subseteq A$*, the* S-reduct B^S *is defined as*

$$B^S = B|_{A \setminus S_B^{\oplus}}$$

where

$$S_B^{\oplus} := S \cup S_B^+$$

denotes the range *of S in B. The restriction or projection* $B|_X$ *for* $X \subseteq A$ *is defined as*

$$B|_X = (X, R \cap (X \times X), U \cap (X \times X))$$

Example 3.3 (Computing a BAF Reduct). *Consider again the BAF* $B_1 = (A,R,U)$ *from Figure 3, with* $A = \{a,b,c,d\}$, $R = \{(b,c),(d,c)\}$ *and* $U = \{(a,b)\}$.

(1)*: The derived attacks are*

$$R_u^{B_1} = \{(a,c),(b,c),(d,c)\}, \quad R_m^{B_1} = \emptyset$$

(2)*: Let* $S = \{a\}$*. Then the set of arguments attacked by S is*

$$S_{B_1}^+ = \{x \in A \mid (a,x) \in R_{\text{der}}^{B_1}\} = \{c\}$$

(3)*: The range is* $S^{\oplus}_{B_1} = \{a, c\}$*. We form the reduct by projecting to*

$$A' = A \setminus \{a, c\} = \{b, d\}$$

Since all relations in B_1 *involved either a or c, the reduct relations become empty*

$$B_1^{\{a\}} = (\{b, d\}, \emptyset, \emptyset)$$

The support (a, b) *and the attacks* (b, c) *and* (d, c) *are all removed because at least one endpoint is no longer in* A'*. Only the arguments b and d remain, now completely isolated.*

Remark 3.3 (Reduct Interpretation and Relation to AFs)**.** *The reduct* B^S *can be read as an intermediate state in which the arguments in S are accepted while arguments in* S^+_B *are rejected. Accordingly,* B^S *removes all arguments in* $S^{\oplus}_B$ *(i.e. S and all arguments attacked by S via derived attacks* R^B_{der}*), and then restricts* both *relations R and U to the remaining arguments. For readability, we write* $A(B^S)$ *to denote the carrier of* B^S*, i.e.*

$$A(B^S) = A \setminus S^{\oplus}_B$$

If $U = \emptyset$*, then the BAF reduct coincides with the classical AF reduct. Derived attacks collapse to direct attacks (i.e.* $R^B_u = R$ *and* $R^B_m = \emptyset$*, hence* $R^B_{\text{der}} = R$*), so* S^+_B *is exactly the set of arguments directly attacked by S. Moreover, Definitions 2.22 to 2.24 reduce to the standard AF notions (conflict-freeness and defence) since no support chains exist. Therefore b-ad*(B) *and IS*(B) *coincide with ad*$((A, R))$ *and IS*$((A, R))$ *when* $U = \emptyset$ *(cf. Proposition 2.6).*

While Definition 3.3 uses R^B_u *and* R^B_m*, the set* S^+_B *can equivalently be defined directly in terms of R and U as* $S^+_B = S^{+u}_B \cup S^{+m}_B$*, where*

Supported attacks *(chains through U ending in R):*

$$S^{+u}_B = \left\{ x \in A \;\middle|\; \begin{array}{l} \exists k \geq 0,\ a_0, \ldots, a_k \in A : \\ a_0 \in S,\ (a_0, a_1) \in U,\ \ldots,\ (a_{k-1}, a_k) \in U,\ (a_k, x) \in R \end{array} \right\}$$

Secondary attacks *(attack through R followed by chains through U):*

$$S^{+m}_B = \left\{ x \in A \;\middle|\; \begin{array}{l} \exists \ell \geq 1,\ b_0, \ldots, b_\ell \in A : \\ b_0 \in S,\ (b_0, b_1) \in R,\ (b_1, b_2) \in U,\ \ldots,\ (b_{\ell-1}, b_\ell) \in U,\ b_\ell = x \end{array} \right\}$$

This makes the chain structure explicit. Computationally, it is equivalent to computing R^B_u *and* R^B_m *via the transitive closure over U, so the derived relations are kept for notational convenience and for alignment with the literature [7].*

Example 3.4 (Indirect Attacks via Support Chains)**.** *Consider the two BAFs in Figure 5:*

$$B_3 = (A, R_3, U_3) \text{ with } A = \{a, b, c\},\ U_3 = \{(a, b)\},\ R_3 = \{(b, c)\}$$

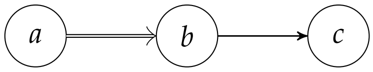

BAF B_3 (supported attack)

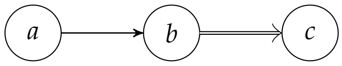

BAF B_4 (secondary attack)

Figure 5: BAFs B_3 and B_4 depicting supported and secondary attack. In both frameworks, argument c is indirectly attacked from a, but not by a single R- or U-edge.

$$B_4 = (A, R_4, U_4) \textit{ with } A = \{a, b, c\},\ R_4 = \{(a, b)\},\ U_4 = \{(b, c)\}$$

In B_3, a does not directly attack any argument, but it supports b and b attacks c, so there is a supported attack from a to c (i.e. $(a, c) \in R_u^{B_3}$). Hence c must be removed in the reduct $B_3^{\{a\}}$, although $(a, c) \notin R_3 \cup U_3$.

In B_4, a attacks b and b supports c, which yields a secondary attack from a to c (i.e. $(a, c) \in R_m^{B_4}$). Thus c is also removed in the reduct $B_4^{\{a\}}$.

These two frameworks show that reducts must take into account supported and secondary attacks (captured by R_{der}^B), not only direct attacks in R.

Remark 3.4 (Computing Derived Relations in Reducts)**.** *For the reduct $B^S = (A', R', U')$ where $A' = A \setminus S_B^{\oplus}$, $R' = R \cap (A' \times A')$, and $U' = U \cap (A' \times A')$, the derived relations must be recomputed:*

$$R_u^{B^S} = U'^* \circ R', \qquad R_m^{B^S} = R' \circ U'^+$$

Note that $R_u^{B^S} \neq R_u^B \cap (A' \times A')$ in general, because support chains may be broken by the removal of arguments. Initial sets in B^S are computed based on the freshly derived relations in the restricted framework.

Table 1: Comparison of reducts in AFs and BAFs.

Reduct type	Description
Classical AF reduct	$AF^S = AF \mid_{A \setminus S^{\oplus}}$, where $S^{\oplus}$ contains all arguments directly attacked by S (via R only).
BAF reduct	$B^S = B \mid_{A \setminus S^{\oplus}_B}$, where $S^{\oplus}_B$ contains all arguments attacked by S via supported or secondary attacks (i.e. via R^B_{der}).
Direct characterisation	Possible via explicit path enumeration (Remark 3.3), computationally equivalent to using R^B_{der}.

Table 1 summarises the key differences between classical AF reducts and BAF reducts, highlighting that direct characterisation using only R and U is possible but computationally equivalent to using derived relations.

Having defined the central concepts for BAFs, we now formally study their properties and how they generalise classical argumentation frameworks.

3.3 Properties of BAF Reducts

Proposition 3.1 (BAF Generalises AF for Reducts)**.** *Let $AF = (A, R)$ and $B = (A, R, \emptyset)$ be the BAF obtained from AF by adding an empty support relation. For any $S \subseteq A$:*

1. *$S^+_B = S^+$ (same attacked sets),*
2. *$B^S = (A \setminus S^{\oplus}_B, R \cap ((A \setminus S^{\oplus}_B) \times (A \setminus S^{\oplus}_B)), \emptyset)$ corresponds to AF^S,*
3. *$b\text{-}ad(B^S) = ad(AF^S)$.*

Proof.

1. For $U = \emptyset$ we have $R^B_u = R$ and $R^B_m = \emptyset$, hence $R^B_{\text{der}} = R$. Thus

$$S^+_B = \{a \in A \mid \exists s \in S : (s, a) \in R\} = S^+$$

and therefore $S^{\oplus}_B = S \cup S^+_B = S \cup S^+ = S^{\oplus}$.

2. By definition of the BAF reduct,

$$B^S = B|_{A \setminus S^{\oplus}_B} = (A \setminus S^{\oplus}_B,\ R \cap ((A \setminus S^{\oplus}_B) \times (A \setminus S^{\oplus}_B)),\ \emptyset)$$

which is exactly $AF^S = (A \setminus S^{\oplus}, R')$ where $R' = R \cap ((A \setminus S^{\oplus}) \times (A \setminus S^{\oplus}))$.

3. The last claim follows from Proposition 2.6 applied to B^S, since $U = \emptyset$ implies that coherent admissibility in B^S coincides with classical admissibility in AF^S.

□

Example 3.5 (BAF Reduct Computation). *Consider the BAF $B_5 = (A, R, U)$ shown in Figure 6, where $A = \{a, b, c, d, e\}$, $R = \{(b, c), (d, e)\}$, and $U = \{(a, b), (e, c)\}$.*

Using Definition 2.20, we obtain the derived relations

$$R_u^{B_5} = \{(a, c), (b, c), (d, e)\} \quad \textit{and} \quad R_m^{B_5} = \{(d, c)\}$$

Indeed, $U^+ = U$ and $U^* = U \cup \{(x, x) \mid x \in A\}$, hence $R_u^{B_5} = U^* \circ R$ and $R_m^{B_5} = R \circ U^+$. For instance, the chain $a\,U\,b\,R\,c$ yields $(a, c) \in R_u^{B_5}$, and $d\,R\,e\,U\,c$ yields $(d, c) \in R_m^{B_5}$.

Let $S = \{a\}$. Then $S^+_{B_5} = \{c\}$ since $(a, c) \in R_u^{B_5}$, i.e.

$$S^+_{B_5} = \{x \in A \mid (a, x) \in R_u^{B_5} \cup R_m^{B_5}\} = \{c\}$$

Consequently, the reduct $B_5^S = (A', R', U')$ is given by

$$A' = A \setminus S^{\oplus}_{B_5} = \{b, d, e\}, \quad R' = R \cap (A' \times A') = \{(d, e)\}, \quad U' = U \cap (A' \times A') = \emptyset$$

Here, $(a, b) \notin U'$ since $a \notin A'$, and $(e, c) \notin U'$ since $c \notin A'$.
Since $U' = \emptyset$, we have $U'^+ = \emptyset$ and $U'^* = \{(x, x) \mid x \in A'\}$, and thus

$$R_u^{B_5^S} = U'^* \circ R' = R' = \{(d, e)\} \quad \textit{and} \quad R_m^{B_5^S} = R' \circ U'^+ = \emptyset$$

In particular, the supported attack $(a, c) \in R_u^{B_5}$ disappears in the reduct.

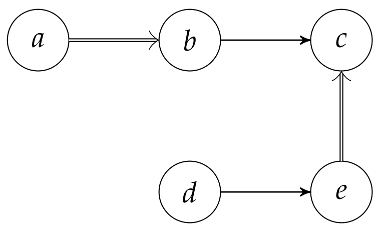

Figure 6: BAF B_5 used in Example 3.5.

Example 3.6 (Support-Chain Breakage). *We show why R_u^B and R_m^B must be recomputed after taking a reduct. Consider the BAF $B_6 = (A, R, U)$ from Figure 7, where $A = \{a, b, c, d\}$, $R = \{(c, d)\}$, and $U = \{(a, b), (b, c)\}$.*

Since $a \xrightarrow{U} b \xrightarrow{U} c \xrightarrow{R} d$, Definition 2.20 yields

$$R_u^{B_6} = U^* \circ R = \{(a, d), (b, d), (c, d)\} \neq \emptyset$$

Let $S = \{b\}$. Then $S^+_{B_6} = \{d\}$ (since $(b, d) \in R_u^{B_6}$), and

$$B_6^S = (\{a, c\}, \emptyset, \emptyset)$$

Hence $R_u^{B_6^S} = \emptyset$, since in B_6^S the support relation is empty and thus no support chain exists to induce a supported attack. Therefore, the derived relations (in particular R_u and R_m) must be recomputed *for the reduct. In general, derived relations in the reduct cannot be obtained by merely restricting $R_u^{B_6}$ and $R_m^{B_6}$ to the remaining arguments $A' = A \setminus (S \cup S_{B_6}^+)$. Note that in this example, restricting $R_u^{B_6}$ to A' (i.e. considering $R_u^{B_6} \cap (A' \times A')$) also yields $\emptyset$. However, this agreement is accidental and need not hold in general.*

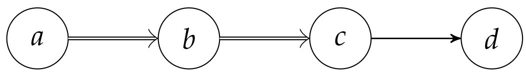

Figure 7: BAF B_6 from Example 3.6

Equipped with these foundational properties for BAF reducts, we can now construct and characterise coherent admissibility through serialisation procedures.

3.4 Characterising Coherent Admissibility through Serialisation

We extend the serialisation-based characterisations from abstract argumentation frameworks to bipolar argumentation frameworks, showing that coherently admissible sets can be constructed by iteratively selecting initial sets. Before instantiating Thimm's selection and termination framework for BAFs, we establish that coherent admissibility admits a reduct-based decomposition into initial sets, which is the key prerequisite for serialisability.

We use the BAF projection notation $B|_X$ for $X \subseteq A$ as introduced in Definition 3.3.

Lemma 3.1 (Projection does not introduce Derived Attacks)**.** *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and let $X \subseteq A$. Consider the projection $B|_X$ to X. Then*

$$R_u^{B|_X} \subseteq R_u^B \cap (X \times X) \qquad \text{and} \qquad R_m^{B|_X} \subseteq R_m^B \cap (X \times X)$$

Consequently,

$$R_{\mathrm{der}}^{B|_X} \subseteq R_{\mathrm{der}}^B \cap (X \times X)$$

Proof. We show that $R_u^{B|_X} \subseteq R_u^B \cap (X \times X)$. The proof for $R_m^{B|_X} \subseteq R_m^B \cap (X \times X)$ is analogous.

Let $(x, y) \in R_u^{B|_X}$. By Definition 2.20, there exist $k \geq 0$ and arguments $a_0, \dots, a_k \in X$ such that $a_0 = x$, $(a_i, a_{i+1}) \in U \cap (X \times X)$ for all $i = 0, \dots, k-1$, and $(a_k, y) \in R \cap (X \times X)$. All arguments $a_0, \dots, a_k$ lie in X by construction.

Since $U \cap (X \times X) \subseteq U$ and $R \cap (X \times X) \subseteq R$, the same sequence establishes $(x, y) \in R_u^B$. Moreover, since $x, y \in X$, we have $(x, y) \in R_u^B \cap (X \times X)$.

The inclusion for R_{der}^B follows from $R_{\mathrm{der}}^B = R_u^B \cup R_m^B$ (Notation 2.3). ☐

Remark 3.5 (Strictness of Projection)**.** *The inclusions in Lemma 3.1 can be strict, since a sequence establishing* $(x, y) \in R^B_{\text{der}}$ *may use intermediate arguments outside X that are no longer available in the projection.*

Notation 3.1 (Simplified Range Notation)**.** *For readability, for a fixed set* $S_1 \subseteq A$ *we write* $S^{\oplus}_{1,B}$ *instead of* $(S_1)^{\oplus}_B$*, and* $S^{+}_{1,B}$ *instead of* $(S_1)^{+}_B$*. Whenever the framework is not clear from context, we explicitly write the subscript (e.g.* S^+_B*,* $S^{\oplus}_B$*) to indicate that* $+$ *and* $\oplus$ *are taken wrt. the derived attacks of the current framework (and analogously in reducts such as* B^{S_1}*).*

Since S^+_B is defined via R^B_{der} (Notation 2.3), Lemma 3.1 immediately yields a corresponding one-direction preservation property for attacked arguments under projection, formalised in the following corollary.

Corollary 3.1 (Projection does not create New Attacked Arguments)**.** *Let* $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$*, let* $X \subseteq A$ *and let* $B|_X$ *be the projection to X as above. For every* $S \subseteq X$ *we have*

$$S^+_{B|_X} \subseteq S^+_B \cap X$$

Proof. By Notation 2.3, for any $B' = (A', R', U')$,

$$S^+_{B'} = \{a \in A' \mid \exists s \in S : (s, a) \in R^{B'}_{\text{der}}\}$$

In particular, for $B' = B|_X$ we obtain

$$S^+_{B|_X} = \{a \in X \mid \exists s \in S : (s, a) \in R^{B|_X}_{\text{der}}\}$$

By Lemma 3.1, $R^{B|_X}_{\text{der}} \subseteq R^B_{\text{der}} \cap (X \times X)$. Since $R^{B|_X}_{\text{der}} \subseteq (X \times X)$, every such target argument a lies in X. Moreover, as $(s, a) \in R^{B|_X}_{\text{der}} \subseteq R^B_{\text{der}}$ for some $s \in S$, we have $a \in S^+_B$ by the definition of S^+_B. Hence every such a lies in $S^+_B \cap X$. □

With the above auxiliary results, we can now prove the main serialisation characterisation for BAFs.

Theorem 3.1 (Constructive Characterisation of Coherent Admissibility)**.** *Let* $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ *be a BAF and* $S \subseteq A$*. Then S is coherently admissible iff either*

- $S = \emptyset$*, or*
- $S = S_1 \cup S_2$ *where* $S_1 \in IS(B)$ *and* S_2 *is coherently admissible in* B^{S_1}*.*

Proof. (⇒) Let S be coherently admissible in B and assume $S \neq \emptyset$. Since A is finite, also S is finite, and therefore there exists a non-empty subset $S_1 \subseteq S$ that is coherently admissible in B and $\subseteq$-minimal among the non-empty coherently admissible subsets of S. By Definition 3.1, we have $S_1 \in \text{IS}(B)$. Let $S_2 := S \setminus S_1$.

Here $S^{\oplus}_{1,B} = S_1 \cup S^{+}_{1,B}$ is taken wrt. the derived attacks of the original framework B, because it determines the carrier A' of the reduct.

Let $A' := A \setminus S^{\oplus}_{1,B}$. By Definition 3.3, the reduct is the projection $B^{S_1} = B|_{A'}$. Write $B^{S_1} = (A', R', U')$. In particular, $S_2 \subseteq A'$.

(1): S_2 is coherent in B^{S_1}. Assume for contradiction that S_2 is not coherent in B^{S_1}. Then by Definition 2.22 there exists $b \in A'$ such that both

(i) $(a, b) \in R^{B^{S_1}}_{\text{der}}$ for some $a \in S_2$, and

(ii) either $b \in S_2$, or there exist $a_1, \ldots, a_n \in A'$ ($n \geq 1$) with $a_1 \in S_2$ and $(a_i, a_{i+1}) \in U'$ for all $i = 1, \ldots, n-1$, and $(a_n, b) \in U'$.

Since $B^{S_1} = B|_{A'}$, Lemma 3.1 yields $R^{B^{S_1}}_{\text{der}} \subseteq R^{B}_{\text{der}} \cap (A' \times A')$, hence $(a, b) \in R^{B}_{\text{der}}$ with $a \in S_2 \subseteq S$.

Moreover, $U' = U \cap (A' \times A')$ as $B^{S_1} = B|_{A'}$. Thus in case (ii) there exist $a_1, \ldots, a_n \in A$ ($n \geq 1$) with $a_1 \in S_2 \subseteq S$ such that $(a_i, a_{i+1}) \in U$ for all $i = 1, \ldots, n-1$ and $(a_n, b) \in U$. Hence, in B, the set S both attacks b (via R^{B}_{der}) and satisfies the non-empty support-chain condition (ii) of Definition 2.22 for b, contradicting coherence of S in B and thus Definition 2.24.

(2): S_2 strongly defends all its elements in B^{S_1}. Let $a \in S_2$ and let $b \in A'$ with $(b, a) \in R^{B^{S_1}}_{\text{der}}$. We show that there exists $c \in S_2$ with $(c, b) \in R^{B^{S_1}}_{\text{der}}$.

Since $B^{S_1} = B|_{A'}$, Lemma 3.1 yields $R^{B^{S_1}}_{\text{der}} \subseteq R^{B}_{\text{der}} \cap (A' \times A')$, hence $(b, a) \in R^{B}_{\text{der}}$.

As S is coherently admissible in B, S strongly defends a in B (Definitions 2.23 and 2.24). Hence there exists $c \in S$ with $(c, b) \in R^{B}_{\text{der}}$.

We claim that $c \notin S_1$. Otherwise, $b \in S^{+}_{1,B}$ by the definition of $S^{+}_{1,B}$ (Notation 2.3), contradicting $b \in A' = A \setminus S^{\oplus}_{1,B}$, where $S^{\oplus}_{1,B} = S_1 \cup S^{+}_{1,B}$. Therefore $c \in S \setminus S_1 = S_2$.

It remains to show $(c, b) \in R^{B^{S_1}}_{\text{der}}$. Assume for contradiction that $(c, b) \notin R^{B^{S_1}}_{\text{der}}$.

Since $(c, b) \in R^{B}_{\text{der}}$, by Notation 2.3 we have either

(a) $(c, b) \in R^{B}_{u}$, i.e. there exist $k \geq 0$ and arguments $d_0, \ldots, d_k \in A$ with $d_0 = c$, $(d_i, d_{i+1}) \in U$ for $i = 0, \ldots, k-1$, and $(d_k, b) \in R$, or

(b) $(c, b) \in R^{B}_{m}$, i.e. there exist $k \geq 1$ and arguments $d_0, \ldots, d_k \in A$ with $d_0 = c$, $(d_0, d_1) \in R$, $(d_i, d_{i+1}) \in U$ for $i = 1, \ldots, k-1$, and $(d_k, b) \in U$.

For readability, we call a finite sequence $(v_0, \ldots, v_m)$ with $(v_i, v_{i+1}) \in U$ for all i a *U-sequence* (a non-empty U-sequence means $m \geq 1$).

Idea. We pick a sequence for $(c, b) \in R^{B}_{\text{der}}$ that minimises the number of intermediate arguments from $S^{\oplus}_{1,B}$. Any first occurrence of such an argument leads to a coherence contradiction, so a sequence can be chosen entirely within A'.

Claim 3.1 (Existence of a Suitable Sequence)**.** *There exists a sequence $d_0, \ldots, d_k$ as above establishing $(c, b) \in R^B_{\text{der}}$ such that no d_i lies in $S^{\oplus}_{1,B}$. In particular, $(c, b) \in R^{B|_{A'}}_{\text{der}}$.*

Proof of Claim 3.1. Among all sequences of the above form establishing $(c, b) \in R^B_{\text{der}}$, choose one that uses a minimal number of arguments from $S^{\oplus}_{1,B}$. Assume for contradiction that this sequence contains some argument from $S^{\oplus}_{1,B}$, and let x be the first such argument on the sequence (starting from $d_0 = c$), with predecessor $y \notin S^{\oplus}_{1,B}$.

$x \in S_1$: If $(y, x) \in U$, then the prefix yields arguments $c = a_0, a_1, \ldots, a_\ell = x$ ($\ell \geq 1$) such that $(a_i, a_{i+1}) \in U$ for all $i = 0, \ldots, \ell - 1$. Since $c \in S$ and $x \in S$, this shows that S satisfies the non-empty support-chain condition (ii) of Definition 2.22 for $b := x$, contradicting Definition 2.22 (case $b \in S$). Otherwise, we have $(y, x) \in R$, hence $(c, x) \in R^B_{\text{der}}$ with $c, x \in S$, also contradicting Definition 2.22 (case $b \in S$).

$x \in S^+_{1,B}$: Fix $s \in S_1 \subseteq S$ with $(s, x) \in R^B_{\text{der}}$ (definition of $S^+_{1,B}$).

Supported attack: suppose $(c, b) \in R^B_u$. Then all edges before the final attack are U-edges, hence $(y, x) \in U$ and there exist arguments $c = a_0, a_1, \ldots, a_\ell = y$ ($\ell \geq 0$) such that $(a_i, a_{i+1}) \in U$ for all $i = 0, \ldots, \ell - 1$. Appending $(y, x) \in U$ yields arguments $p_0, \ldots, p_m \in A$ with $m \geq 1$ such that $p_0 = c$, $p_m = x$, and $(p_i, p_{i+1}) \in U$ for all $i = 0, \ldots, m - 1$. In particular, the side-condition $n \geq 1$ in Definition 2.22(ii) is satisfied for $b := x$. Concatenating a sequence establishing $(s, x) \in R^B_{\text{der}}$ with this U-sequence from x to b yields $(s, b) \in R^B_{\text{der}}$, hence $b \in S^+_{1,B}$, contradicting $b \in A'$.

Secondary attack: suppose $(c, b) \in R^B_m$. By definition of R_m, we have $(d_0, d_1) \in R$ and there exist arguments $e_0, \ldots, e_t \in A$ with $t \geq 1$ such that $e_0 = d_1$, $(e_i, e_{i+1}) \in U$ for all $i = 0, \ldots, t - 1$, and $e_t = b$.

$x = d_1$: Then $(c, x) \in R$. Moreover, by the definition of the sequence $e_0, \ldots, e_t$ fixed above, we have $e_0 = x$.

If $x \in S_1 \subseteq S$, then $(c, x) \in R \subseteq R^B_{\text{der}}$ with $c, x \in S$, contradicting Definition 2.22 (case $b \in S$). So $x \in S^+_{1,B}$ (since $x \in S^{\oplus}_{1,B} = S_1 \cup S^+_{1,B}$).

Appending a sequence establishing $(s, x) \in R^B_{\text{der}}$ with the above U-sequence from x to b. This yields $(s, b) \in R^B_{\text{der}}$, hence $b \in S^+_{1,B}$, contradicting $b \in A'$.

$x \neq d_1$: Since x occurs among the arguments $e_0, \ldots, e_t$ with $(e_i, e_{i+1}) \in U$ for all $i = 0, \ldots, t - 1$, we have $x = e_j$ for some $j \in \{0, \ldots, t\}$. Moreover, in this subcase we have $x \neq d_1 = e_0$, and we also have $x \neq b = e_t$ (since $b \in A'$ but $x \in S^{\oplus}_{1,B}$). Hence there exists an index $j \in \{1, \ldots, t - 1\}$ with $e_j = x$. Hence, for the arguments $e_j = x, e_{j+1}, \ldots, e_t = b$ we have $(e_i, e_{i+1}) \in U$ for all $i = j, \ldots, t - 1$, and in particular there exists $i \in \{j, \ldots, t - 1\}$, hence at least one pair $(e_i, e_{i+1}) \in U$ occurs.

Extending a sequence establishing $(s, x) \in R^B_{\text{der}}$ with the arguments $e_j = x, e_{j+1}, \ldots, e_t = b$ (where $(e_i, e_{i+1}) \in U$ for all $i = j, \ldots, t - 1$) yields $(s, b) \in R^B_{\text{der}}$. Hence $b \in S^+_{1,B} \subseteq S^{\oplus}_{1,B}$, contradicting $b \in A'$.

Therefore, the chosen minimal sequence contains no element from $S^{\oplus}_{1,B}$. □

By the claim, there exists a sequence $d_0, \ldots, d_k$ establishing $(c, b) \in R^B_{\text{der}}$ with all $d_i \in A'$. Since $B^{S_1} = B|_{A'}$, we have $R' = R \cap (A' \times A')$ and $U' = U \cap (A' \times A')$, so the same sequence establishes $(c, b) \in R^{B^{S_1}}_{\text{der}}$, contradiction.

Therefore, $(c, b) \in R^{B^{S_1}}_{\text{der}}$ for some $c \in S_2$, and S_2 strongly defends a in B^{S_1}.

Thus S_2 is coherent and strongly defends all its elements in B^{S_1}, hence S_2 is coherently admissible in B^{S_1} by Definition 2.24.

($\Leftarrow$) Assume $S = \emptyset$. Then S is coherently admissible by Proposition 2.5. Now assume $S = S_1 \cup S_2$ with $S_1 \in \text{IS}(B)$ and S_2 coherently admissible in B^{S_1}.

(1): Coherence of S in B. Assume for contradiction that S is not coherent in B. Then there exists $c \in A$ such that both

(i) $(a, c) \in R^B_{\text{der}}$ for some $a \in S$, and

(ii) either $c \in S$, or there exist $a_1, \ldots, a_n \in A$ ($n \geq 1$) with $a_1 \in S$ and $(a_i, a_{i+1}) \in U$ for all $i = 1, \ldots, n-1$, and $(a_n, c) \in U$.

If $c \in A'$, let $a \in S$ be such that $(a, c) \in R^B_{\text{der}}$ (from (i)). Then $a \notin S_1$, since otherwise $c \in S^+_{1,B} \subseteq S^{\oplus}_{1,B}$, contradicting $c \in A' = A \setminus S^{\oplus}_{1,B}$. Hence $a \in S_2$. Then the argument of ($\Rightarrow$), part *(1)* (applied to S_2 in $B^{S_1} = B|_{A'}$) yields that S_2 is not coherent in B^{S_1}, contradiction.

If $c \in S^{\oplus}_{1,B}$, then $c \in S_1$ or $c \in S^+_{1,B}$. If $c \in S_1$, then $c \in S$ and we directly obtain a violation of coherence of S by Definition 2.22 (case $c \in S$), contradiction. If $c \in S^+_{1,B}$, then by definition of $S^+_{1,B}$ there exists $s \in S_1 \subseteq S$ such that $(s, c) \in R^B_{\text{der}}$, i.e. S attacks c. Together with (ii) above this violates Definition 2.22, contradiction. Hence S is coherent.

(2): Strong defence of S in B. Let $a \in S$ and $b \in A$ with $(b, a) \in R^B_{\text{der}}$.

If $a \in S_1$, then since S_1 is coherently admissible, there exists $c \in S_1 \subseteq S$ with $(c, b) \in R^B_{\text{der}}$.

Now assume $a \in S_2$. Since $A' = A \setminus S^{\oplus}_{1,B}$, we have either $b \in A'$ or $b \in S^{\oplus}_{1,B} = S_1 \cup S^+_{1,B}$. We distinguish cases accordingly:

- If $b \in S_1$, then since S_1 is coherently admissible, there exists $c \in S_1 \subseteq S$ with $(c, b) \in R^B_{\text{der}}$.
- If $b \in S^+_{1,B}$, then by definition of $S^+_{1,B}$ there exists $c \in S_1 \subseteq S$ with $(c, b) \in R^B_{\text{der}}$.
- If $b \in A'$, then since $B^{S_1} = B|_{A'}$, Lemma 3.1 yields $R^{B^{S_1}}_{\text{der}} \subseteq R^B_{\text{der}} \cap (A' \times A')$, hence $(b, a) \in R^{B^{S_1}}_{\text{der}}$. Since S_2 is coherently admissible in B^{S_1}, there exists $c \in S_2 \subseteq S$ with $(c, b) \in R^{B^{S_1}}_{\text{der}} \subseteq R^B_{\text{der}}$ (again by Lemma 3.1).

Thus S strongly defends all its elements. Therefore, S is coherently admissible in B by Definition 2.24. □

An immediate consequence of Theorem 3.1 is that every coherently admissible set can be decomposed into a sequence of initial sets.

Corollary 3.2 (Recursive Decomposition)**.** *Every non-empty coherently admissible set S can be written as $S = S_1 \cup \ldots \cup S_n$ with pairwise disjoint S_i ($i = 1, \ldots, n$), where S_1 is an initial set of B and every S_i ($i = 2, \ldots, n$) is an initial set of $B^{S_1 \cup \ldots \cup S_{i-1}}$. Furthermore, only non-empty coherently admissible sets can be written in such a fashion.*

Proof. This follows by iterative application of Theorem 3.1, and the finiteness of S guarantees termination after finitely many steps. □

Remark 3.6 (Constructive Reading)**.** *Corollary 3.2 can be read constructively. Starting with a non-empty coherently admissible set S, Theorem 3.1 yields an initial set $S_1 \in IS(B)$ with $S_1 \subseteq S$ such that $S \setminus S_1$ is coherently admissible in the reduct B^{S_1}. Repeating this step in successive reducts produces a finite sequence of pairwise disjoint initial sets $S_1, \ldots, S_n$ whose union is S, and finiteness of S guarantees termination after finitely many iterations.*

Example 3.7 (Serialisation Sequence for BAFs)**.** *We give a purely schematic illustration of the recursive decomposition from Corollary 3.2. Let B be a BAF and let $E = \{a, b, c, d\} \neq \emptyset$ be coherently admissible in B.*

Assume that there exist sets S_1, S_2, S_3 such that

- $S_1 \in IS(B)$,
- $S_2 \in IS(B^{S_1})$,
- $S_3 \in IS(B^{S_1 \cup S_2})$,

and $E = S_1 \cup S_2 \cup S_3$. Then $\mathcal{S} = (S_1, S_2, S_3)$ is a serialisation sequence with result E.

We now give a concrete instance including the computation of derived attacks and reducts in Example 3.8.

Example 3.8 (Concrete BAF Serialisation)**.** *Consider the BAF $B_7 = (A, R, U)$ from Figure 8, where $A = \{a, b, c, d, e\}$, $R = \{(b, a), (d, c), (e, d)\}$, and $U = \{(a, c), (c, e)\}$.*

Using Definition 2.20, the derived attack relations are

- $R_u^{B_7} = \{(b, a), (d, c), (e, d), (a, d), (c, d)\}$,
- $R_m^{B_7} = \{(b, c), (b, e), (d, e)\}$.

In particular, $\{a\}, \{b\}, \{d\}, \{e\} \in IS(B_7)$.

For $S_1 = \{b\}$ we obtain

$$\{b\}^+_{B_7} = \{a, c, e\}$$

since $(b, a) \in R_u^{B_7}$ and $(b, c), (b, e) \in R_m^{B_7}$.

Hence the reduct is

$$B_7^{\{b\}} = (\{d\}, \emptyset, \emptyset)$$

because we remove $\{b\} \cup \{b\}^+_{B_7} = \{a, b, c, e\}$ *from* A *and all incident attacks and supports vanish.*

In $B_7^{\{b\}}$*, the set* $\{d\}$ *is coherently admissible and (since no proper non-empty subset exists) it is an initial set. Therefore,* $\mathcal{S} = (\{b\}, \{d\})$ *is a serialisation sequence with result* $E = \{b, d\}$*.*

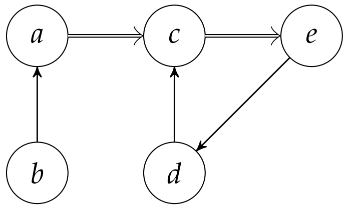

Figure 8: BAF B_7 from Example 3.8

Table 2: Translation from AF concepts to BAF equivalents for serialisation.

Concept	**AF**	**BAF Equivalent**
Admissible sets	$\text{ad}(AF)$	$\text{b-ad}(B)$ (uses coherence and strong defence)
Initial sets	$\text{IS}(AF)$	$\text{IS}(B)$ (computed using coherent admissibility)
Attack relation	R (direct attacks only)	R^B_{der} (includes supported and secondary attacks)
Reduct	$AF^S = AF\|_{A \setminus S^\oplus}$	$B^S = B\|_{A \setminus S^\oplus_B}$ (cf. Definition 3.3)
Attacked set	$S^+ = \{a \mid \exists s \in S : sRa\}$	$S^+_B = \{a \mid \exists s \in S : (s, a) \in R^B_{\text{der}}\}$

Proposition 3.2 (Serialisation Sequences in the AF Case)**.** *Let* $AF = (A, R) \in \mathfrak{F}_\mathfrak{A}$ *and* $B_{AF} = (A, R, \emptyset) \in \mathfrak{B}_\mathfrak{A}$*. Then* $IS(AF) = IS(B_{AF})$*, and for every* $S \subseteq A$ *the AF reduct* AF^S *coincides with the BAF reduct* $(B_{AF})^S$ *when interpreted as an AF.*

Proof. Initial sets. A set S is an initial set in AF iff it is nonempty, admissible, and minimal. Since admissibility coincides with coherent admissibility when $U = \emptyset$ (Proposition 2.6), minimality is also preserved. Thus $\text{IS}(AF) = \text{IS}(B_{AF})$.

Reducts. The AF reduct removes $S^\oplus$ where $S^+ = \{a \mid \exists s \in S : sRa\}$. The BAF reduct removes $S^\oplus_B$ where $S^+_B = \{a \mid \exists s \in S : (s, a) \in R^B_{\text{der}}\}$. Since $U = \emptyset$, we have $R^B_{\text{der}} = R$ and hence $S^+_B = S^+$ and $S^\oplus_B = S^\oplus$. Both reduct constructions then restrict R to the same carrier, so they coincide. □

Remark 3.7 (Conservative Extension)**.** *Table 2 and Proposition 3.2 show that Thimm's AF serialisability framework [14] conservatively extends to BAFs. When* $U = \emptyset$*, all BAF* *notions collapse to their AF counterparts (cf. Proposition 2.6 and Remark 3.3). In the bipolar setting, the interaction between support and attack is captured via R_u^B (supported attacks) and R_m^B (secondary attacks), which collapse to direct attacks when $U = \emptyset$.*

Having established the existence of serialisation sequences for b-ad via Theorem 3.1 and Corollary 3.2, we now refine this construction by adding selection and termination constraints, mirroring Section 2.4.

3.5 Serialisability in BAFs

We now extend the concept of serialisability from Section 2.4 to BAFs, adapting the framework of Thimm [14] to account for the interplay between attack and support relations (cf. Table 2 for the systematic translation).

Following Thimm's account of serialisability in AFs, we treat extension construction as a step-wise process over states (B, S). At each step we select an initial set of the current framework and move to the corresponding reduct, accumulating accepted arguments. Different semantics are obtained by restricting (i) which types of initial sets may be selected and (ii) when the process terminates, captured by suitable selection and termination functions (α, β).

3.5.1 Serialisation Sequences and Auxiliary Results

We first formalise serialisation sequences for BAFs and collect auxiliary results that will be used in the subsequent semantic characterisations.

Definition 3.4 (Serialisation Sequence for BAFs)**.** *A* serialisation sequence *for $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ is a sequence $\mathcal{S} = (S_1, \ldots, S_n)$ such that*

- *$S_1 \in IS(B)$, and*
- *for each $2 \leq i \leq n$: $S_i \in IS(B^{E_{i-1}})$, where $E_{i-1} := S_1 \cup \cdots \cup S_{i-1}$.*

The set $E := S_1 \cup \cdots \cup S_n$ is called the result *of the serialisation sequence $\mathcal{S}$.*

Notation 3.2 (Partial results)**.** *For a serialisation sequence $\mathcal{S} = (S_1, \ldots, S_n)$ in a BAF, we use $E_i := S_1 \cup \cdots \cup S_i$ to denote the union of the first i steps, with $E_0 := \emptyset$ by convention.*

The following lemmata establish key properties needed for the main serialisability theorem.

Lemma 3.2 (Coherence Preservation Under Projection)**.** *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$, let $A' \subseteq A$, and let $S \subseteq A'$. If S is coherent in B, then S is coherent in the projection $B|_{A'}$.*

Proof. Assume that S is coherent in B. Suppose S is not coherent in $B|_{A'}$. Then, by Definition 2.22, there exists $b \in A'$ such that (i) S attacks b via $R_{\text{der}}^{B|_{A'}}$, and (ii) S

supports b in $B|_{A'}$ (either $b \in S$ or there is a non-empty U-chain from some $s \in S$ to b).

By Lemma 3.1, derived attacks in $B|_{A'}$ are also derived attacks in B. Since $U|_{A'} \subseteq U$, support chains are preserved. Hence S would simultaneously attack and support b in B, contradicting coherence in B. Therefore S is coherent in $B|_{A'}$. $\square$

The following lemmata establish key properties needed for the main serialisability results of this subsection.

Lemma 3.3 (Coherent Union Property)**.** *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and let $S_1, S_2 \subseteq A$. If S_1 and S_2 are coherently admissible in B and $S_1 \cup S_2$ is coherent in B, then $S_1 \cup S_2$ is coherently admissible in B.*

Proof. Since $S_1 \cup S_2 \in$ b-coh(B) by assumption, it suffices by Definition 2.24 to show that $S_1 \cup S_2$ strongly defends all its elements.

Let $a \in S_1 \cup S_2$ and let $(b, a) \in R^B_{\text{der}}$. If $a \in S_1$, then $S_1 \in$ b-ad(B) yields some $c \in S_1 \subseteq S_1 \cup S_2$ with $(c, b) \in R^B_{\text{der}}$. If $a \in S_2$, the argument is identical using $S_2 \in$ b-ad(B).

Hence $S_1 \cup S_2$ strongly defends all its elements. Together with $S_1 \cup S_2 \in$ b-coh(B), Definition 2.24 yields $S_1 \cup S_2 \in$ b-ad(B). $\square$

In what follows, we will use Corollary 3.2, which follows from Theorem 3.1 by iterative application.

Lemma 3.4 (Serial Extension and Lifting)**.** *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and let $E, S \subseteq A$. If $E \in$ b-ad(B) and $S \in$ b-ad(B^E), then $E \cup S \in$ b-ad(B).*

Proof. Since $S \in$ b-ad(B^E), we have

$$S \subseteq A(B^E) \;=\; A \setminus E^{\oplus}_B \qquad \text{(cf. Remark 3.3)}$$

By Corollary 3.2, fix a serialisation $(E_1, \ldots, E_k)$ of E and define

$$F_i := E_1 \cup \cdots \cup E_i \quad (i \in \{0, \ldots, k\}), \qquad F_0 := \emptyset, \;\; F_k = E.$$

We prove by backward induction on $i = k, k-1, \ldots, 0$ that

$$(E \setminus F_i) \cup S \in \text{b-ad}(B^{F_i})$$

$i = k$: Here $F_k = E$, hence $(E \setminus F_k) \cup S = S$. Thus the claim is $S \in$ b-ad(B^E), which holds by assumption.

$i \to i-1$: Assume $(E \setminus F_i) \cup S \in \text{b-ad}(B^{F_i})$. Work in the framework $B^{F_{i-1}}$. By serialisation, $E_i \in \text{IS}(B^{F_{i-1}})$ and $F_i = F_{i-1} \cup E_i$. Moreover, iterated reducts compose, i.e.

$$(B^{F_{i-1}})^{E_i} = B^{F_i}$$

Indeed, by Definition 3.3, reducts are obtained by projecting to the corresponding carrier and restricting R and U accordingly. Unfolding the definitions and using $F_i = F_{i-1} \cup E_i$ yields the claim. Set

$$S_1 := E_i, \qquad S_2 := (E \setminus F_i) \cup S$$

By the induction hypothesis and $(B^{F_{i-1}})^{S_1} = (B^{F_{i-1}})^{E_i} = B^{F_i}$ we obtain

$$S_2 \in \text{b-ad}\big((B^{F_{i-1}})^{S_1}\big)$$

Since $S_1 = E_i \in \text{IS}(B^{F_{i-1}})$ and $S_2 \in \text{b-ad}\big((B^{F_{i-1}})^{S_1}\big)$, we can apply the $(\Leftarrow)$-direction of Theorem 3.1 in $B^{F_{i-1}}$ to the decomposition

$$(E \setminus F_{i-1}) \cup S \;=\; S_1 \cup S_2 \;=\; E_i \cup (E \setminus F_i) \cup S$$

and obtain

$$(E \setminus F_{i-1}) \cup S \in \text{b-ad}(B^{F_{i-1}})$$

This completes the induction. For $i = 0$ we obtain

$$(E \setminus F_0) \cup S = E \cup S \in \text{b-ad}(B^{F_0}) = \text{b-ad}(B)$$

i.e. $E \cup S \in \text{b-ad}(B)$. □

We can now state the main serialisability result for coherent admissibility. By Corollary 3.2, every non-empty $E \in \text{b-ad}(B)$ admits a finite decomposition into initial sets of successive reducts, i.e. a serialisation sequence in the sense of Definition 3.4.

3.5.2 Semantic Characterisations via Serialisation

We now lift the serialisation sequences for b-ad to semantic characterisations for the individual coherence-based semantics from Definition 2.26.

Theorem 3.2 (Serialisability of Coherently Admissible Semantics). *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and $E \subseteq A$. Then $E \in$ b-ad(B) iff there exists a serialisation sequence $\mathcal{S} = (S_1, \ldots, S_n)$ with result E.*

Proof. $(\Rightarrow)$ If $E = \emptyset$, choose the empty sequence. If $E \neq \emptyset$ and $E \in \text{b-ad}(B)$, then Corollary 3.2 yields a finite sequence of initial sets in successive reducts whose union is E, hence a serialisation sequence with result E in the sense of Definition 3.4. $(\Leftarrow)$

Let $\mathcal{S} = (S_1, \ldots, S_n)$ be a serialisation sequence with result E. We show by induction on n that $E \in \text{b-ad}(B)$.

$n = 0$: Then $E = \emptyset \in \text{b-ad}(B)$.

$n \geq 1$: Let $E_{n-1} := S_1 \cup \cdots \cup S_{n-1}$ be the prefix result. By the induction hypothesis, $E_{n-1} \in \text{b-ad}(B)$. Moreover, $S_n \in \text{IS}(B^{E_{n-1}})$ implies $S_n \in \text{b-ad}(B^{E_{n-1}})$. By Lemma 3.4, we obtain $E_{n-1} \cup S_n \in \text{b-ad}(B)$, i.e. $E \in \text{b-ad}(B)$. □

Before characterising further coherence-based semantics via serialisation, we recall that the reduct notation B^E and its carrier $A(B^E)$ are defined in Definition 3.3 and Remark 3.3.

Theorem 3.3 (Serialisability of Coherently Complete Semantics)**.** *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and $E \subseteq A$. Then $E \in$ b-co(B) iff there exists a serialisation sequence $\mathcal{S} = (S_1, \ldots, S_n)$ with result E and $IS^{\not\leftarrow}(B^E) = \emptyset$.*

Proof. We use the following auxiliary claim, which is the bipolar counterpart of the standard AF argument (cf. [14] for the AF case).

Claim 3.2 (From Strong Defence to Unattacked Initial Sets in B^E)**.** *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and $E \subseteq A$. For every $a \in A(B^E) = A \setminus E^{\oplus}_B$ we have:*

$$a \text{ is strongly defended by } E \text{ in } B \quad \Longrightarrow \quad \{a\} \in IS^{\not\leftarrow}(B^E)$$

Proof of Claim 3.2. Let $a \in A(B^E)$ be strongly defended by E in B.

Towards a contradiction, assume there exists $x \in A(B^E)$ with $(x, a) \in R^{B^E}_{\text{der}}$. By Lemma 3.1, this implies $(x, a) \in R^B_{\text{der}}$. Since E strongly defends a in B, there exists $c \in E$ with $(c, x) \in R^B_{\text{der}}$. Hence $x \in E^+_B \subseteq E^{\oplus}_B$, contradicting $x \in A(B^E) = A \setminus E^{\oplus}_B$. Therefore a is unattacked in B^E.

Since $\{a\}$ is coherent (singletons are trivially coherent) and strongly defends its only element (as a is unattacked in B^E), we have $\{a\} \in \text{b-ad}(B^E)$ by Definition 2.24. Minimality is immediate for singletons, hence $\{a\} \in \text{IS}^{\not\leftarrow}(B^E)$. □

($\Rightarrow$) Let $E \in \text{b-co}(B)$.

Assume for contradiction that $\text{IS}^{\not\leftarrow}(B^E) \neq \emptyset$. Pick $\{a\} \in \text{IS}^{\not\leftarrow}(B^E)$. Then $a \in A(B^E)$ and a is unattacked in B^E.

We show that a is strongly defended by E in B. Let $b \in A$ with $(b, a) \in R^B_{\text{der}}$.

If $b \notin E^{\oplus}_B$, then $b \in A(B^E)$ and, by Lemma 3.1, $(b, a) \in R^{B^E}_{\text{der}}$, contradicting that a is unattacked in B^E. Hence $b \in E^{\oplus}_B = E \cup E^+_B$.

Moreover, $b \notin E$: if $b \in E$ and $(b, a) \in R^B_{\text{der}}$, then $a \in E^+_B$, hence $a \in E^{\oplus}_B$, contradicting $a \in A(B^E) = A \setminus E^{\oplus}_B$. Therefore $b \in E^+_B$, and by definition of E^+_B there exists $c \in E$ with $(c, b) \in R^B_{\text{der}}$.

Thus every derived attacker b of a is counter-attacked by some $c \in E$, i.e. E strongly defends a in B.

Since E is coherent-complete, strong defence implies $a \in E$ (cf. Definition 2.26). But $a \in A(B^E) = A \setminus E_B^{\oplus}$ and $E \subseteq E_B^{\oplus}$, hence $A(B^E)$ is disjoint from E. Contradiction. Thus $\mathrm{IS}^{\not\leftarrow}(B^E) = \emptyset$.

($\Leftarrow$) Let $\mathcal{S}$ be a serialisation sequence with result E such that $\mathrm{IS}^{\not\leftarrow}(B^E) = \emptyset$. By Theorem 3.2, we have $E \in \text{b-ad}(B)$.

Let $a \in A$ be strongly defended by E in B and assume towards a contradiction that $a \notin E$.

We first show that $a \notin E_B^+$. Assume towards a contradiction that $a \in E_B^+$. Then there exists $e \in E$ with $(e, a) \in R_{\text{der}}^B$. Since E strongly defends a (by assumption), there exists $c \in E$ with $(c, e) \in R_{\text{der}}^B$. Thus E contains both e and c with $(c, e) \in R_{\text{der}}^B$, contradicting strong conflict-freeness of E (since $E \in \text{b-ad}(B)$ implies coherence and hence strong conflict-freeness, cf. Proposition 2.4). Therefore $a \notin E_B^+$.

Consequently, $a \notin E_B^{\oplus} = E \cup E_B^+$ and thus $a \in A(B^E)$.

By Claim 3.2, $\{a\} \in \mathrm{IS}^{\not\leftarrow}(B^E)$, contradicting $\mathrm{IS}^{\not\leftarrow}(B^E) = \emptyset$. Therefore $a \in E$, and E is complete. □

Corollary 3.3 (Coherently Complete Characterisation)**.** *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and $S \subseteq A$. S is coherently complete iff either*

- *$S = \emptyset$ and $IS^{\not\leftarrow}(B) = \emptyset$, or*
- *$S = S_1 \cup S_2$ where $S_1 \in IS(B)$ and S_2 is coherently complete in B^{S_1}.*

Proof. Follows directly from Theorem 3.3 by recursive application. □

Theorem 3.4 (Serialisability of Coherently Grounded Semantics)**.** *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and $E \subseteq A$. Then $E \in$ b-gr(B) iff there exists a serialisation sequence $\mathcal{S} = (S_1, \ldots, S_n)$ with result E such that $\forall i \in \{1, \ldots, n\} : S_i \in IS^{\not\leftarrow}(B^{E_{i-1}})$ and $IS^{\not\leftarrow}(B^E) = \emptyset$.*

Proof. ($\Rightarrow$) Let $E \in \text{b-gr}(B)$. Then E is the $\subseteq$-minimal element of $\text{b-co}(B)$ (cf. Definition 2.26), hence $E \in \text{b-co}(B)$. By Theorem 3.3, there exists a serialisation sequence $\mathcal{S} = (S_1, \ldots, S_n)$ with result E such that $\mathrm{IS}^{\not\leftarrow}(B^E) = \emptyset$.

It remains to show that $\mathcal{S}$ can be chosen so that $S_i \in \mathrm{IS}^{\not\leftarrow}(B^{E_{i-1}})$ for all $i \in \{1, \ldots, n\}$. Suppose, for contradiction, that this fails for some index j, i.e. $S_j \in \mathrm{IS}(B^{E_{j-1}})$ is attacked in $B^{E_{j-1}}$. By Theorem 3.2, the prefix $(S_1, \ldots, S_{j-1})$ yields $E_{j-1} \in \text{b-ad}(B)$.

We show $E_{j-1} \in \text{b-co}(B)$, which contradicts the $\subseteq$-minimality of E (since $j \geq 1$ implies $E_{j-1} \subsetneq E$). Let $y \in A$ be strongly defended by E_{j-1} in B; assume towards a contradiction that $y \notin E_{j-1}$.

We first show $y \notin (E_{j-1})^{\oplus}_{B}$. If $y \in (E_{j-1})^{+}_{B}$, there exists $e \in E_{j-1}$ with $(e, y) \in R^{B}_{\mathrm{der}}$. Since E_{j-1} strongly defends y, there exists $c \in E_{j-1}$ with $(c, e) \in R^{B}_{\mathrm{der}}$, contradicting strong conflict-freeness of E_{j-1} (which holds because $E_{j-1} \in$ b-ad(B) implies coherence, cf. Proposition 2.4). Hence $y \notin (E_{j-1})^{\oplus}_{B}$, so $y \in A(B^{E_{j-1}}) = A \setminus (E_{j-1})^{\oplus}_{B}$.

By Claim 3.2 (with E_{j-1} in place of E), $\{y\} \in \mathrm{IS}^{\not\leftarrow}(B^{E_{j-1}})$. By Theorem 3.3 ($\Leftarrow$), applied to E_{j-1} and the serialisation $(S_1, \ldots, S_{j-1})$ extended by the single step $\{y\}$ (which is unattacked and hence in $\mathrm{IS}^{\not\leftarrow}(B^{E_{j-1}})$), we see that $E_{j-1} \cup \{y\}$ is coherently admissible. Moreover, $\{y\}$ is strongly defended by E_{j-1}, so every strongly defended argument of E_{j-1} is also strongly defended by E_{j-1}, and since $y \notin E_{j-1}$ is strongly defended, E_{j-1} fails to contain all strongly defended arguments. Thus $y \in E_{j-1}$, a contradiction.

Therefore every argument strongly defended by E_{j-1} in B belongs to E_{j-1}, so $E_{j-1} \in$ b-co(B) (cf. Definition 2.26). Since $j \geq 1$ we have $E_{j-1} \subsetneq E$, contradicting $\subseteq$-minimality of E in b-co(B). Hence every step satisfies $S_i \in \mathrm{IS}^{\not\leftarrow}(B^{E_{i-1}})$, and together with $\mathrm{IS}^{\not\leftarrow}(B^{E}) = \emptyset$ this proves ($\Rightarrow$).

($\Leftarrow$) Let $\mathcal{S} = (S_1, \ldots, S_n)$ be a serialisation sequence with result E such that $\forall i : S_i \in \mathrm{IS}^{\not\leftarrow}(B^{E_{i-1}})$ and $\mathrm{IS}^{\not\leftarrow}(B^{E}) = \emptyset$. By Theorem 3.3, it follows that $E \in$ b-co(B).

To show that E is grounded, let $E' \in$ b-co(B) be arbitrary. We prove by induction on $i \in \{0, \ldots, n\}$ that $E_i \subseteq E'$.

($i = 0$): $E_0 = \emptyset \subseteq E'$.

($i \geq 1$): Assume $E_{i-1} \subseteq E'$. Let $a \in S_i$. Since $S_i \in \mathrm{IS}^{\not\leftarrow}(B^{E_{i-1}})$, a is unattacked in $B^{E_{i-1}}$. Hence every derived attacker of a in B lies in $(E_{i-1})^{\oplus}_{B}$, so E_{i-1} strongly defends a in B. Since $E_{i-1} \subseteq E'$ and E' is coherently complete, E' also strongly defends a, hence $a \in E'$. Thus $S_i \subseteq E'$ and $E_i = E_{i-1} \cup S_i \subseteq E'$.

Hence $E = E_n \subseteq E'$ for all $E' \in$ b-co(B), so E is the $\subseteq$-minimal element of b-co(B), i.e. $E \in$ b-gr(B). □

Remark 3.8 (Unattacked Initial Sets in AFs and BAFs). *In the classical AF setting (i.e. $U = \emptyset$), unattacked initial sets are always singletons [14]. In our BAF setting, unattacked initial sets $IS^{\not\leftarrow}(B)$ may in principle be larger, as coherence is defined relative to derived attacks and support chains. However, the serialisation results for b-gr below only rely on the selection of unattacked initial sets and do not require them to be singletons. In particular, whenever an unattacked singleton $\{a\}$ is coherently admissible, it is itself an element of $IS^{\not\leftarrow}(B)$ (cf. Remark 3.2). Thus, the constructive process for coherently grounded semantics remains well-defined even without a general singleton property for $IS^{\not\leftarrow}(B)$.*

Corollary 3.4 (Coherently Grounded Characterisation). *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and $E \subseteq A$. $E \in$ b-gr(B) iff either $E = \emptyset$ and $IS^{\not\leftarrow}(B) = \emptyset$, or $E = S_1 \cup S_2$ where $S_1 \in IS^{\not\leftarrow}(B)$ and $S_2 \in$ b-gr(B^{S_1}).*

Proof. Follows from Theorem 3.4 by recursive application. □

Theorem 3.5 (Serialisability of Coherently Preferred Semantics). *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and $E \subseteq A$. Then $E \in$ b-pr(B) iff there exists a serialisation sequence $\mathcal{S} = (S_1, \ldots, S_n)$ with result E and IS$(B^E) = \emptyset$.*

Proof. ($\Rightarrow$) Let $E \in$ b-pr(B). Then E is $\subseteq$-maximal in b-ad(B). By Theorem 3.2, there exists a serialisation sequence $(S_1, \ldots, S_n)$ with result E.

Assume for contradiction that IS$(B^E) \neq \emptyset$. Pick some non-empty $S \in$ IS(B^E). By definition of initial sets, $S \in$ b-ad(B^E). Since $E \in$ b-ad(B) and $S \in$ b-ad(B^E), Lemma 3.4 yields $E \cup S \in$ b-ad(B). Moreover, $S \subseteq A(B^E) = A \setminus E^{\oplus}_B$, hence $S \cap E = \emptyset$ and thus $E \subsetneq E \cup S$, contradicting the $\subseteq$-maximality of E in b-ad(B). Therefore IS$(B^E) = \emptyset$.

($\Leftarrow$) Let $(S_1, \ldots, S_n)$ be a serialisation sequence with result E such that IS$(B^E) = \emptyset$. By Theorem 3.2, $E \in$ b-ad(B).

Assume for contradiction that E is not $\subseteq$-maximal, i.e. there exists $E' \in$ b-ad(B) with $E \subsetneq E'$. Define $S' := E' \cap A(B^E)$. Then $S' \neq \emptyset$ and $S' \subseteq A(B^E) = A \setminus E^{\oplus}_B$. Moreover, $S' \in$ b-ad(B^E). Hence there exists an initial set $S \subseteq S'$ in B^E with $S \in$ IS(B^E), contradicting IS$(B^E) = \emptyset$.

Therefore E is $\subseteq$-maximal, i.e. $E \in$ b-pr(B). □

Corollary 3.5 (Coherently Preferred Characterisation). *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and $E \subseteq A$. $E \in$ b-pr(B) iff either $E = \emptyset$ and IS$(B) = \emptyset$, or $E = S_1 \cup S_2$ where $S_1 \in$ IS(B) and $S_2 \in$ b-pr(B^{S_1}).*

Proof. Follows from Theorem 3.5 by recursive application. □

We recall the corresponding definitions from Definition 2.26. For the stable case, note that the condition $E \in$ b-ad(B) and $E^{\oplus}_B = A$ (cf. Definition 2.26) is equivalent to $B^E = (\emptyset, \emptyset, \emptyset)$, since $A(B^E) = A \setminus E^{\oplus}_B$ by Definition 3.3.

Theorem 3.6 (Serialisability of Coherently Stable Semantics). *Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and $E \subseteq A$. Then $E \in$ b-st(B) iff $E \in$ b-ad(B) and $E^{\oplus}_B = A$ (cf. Definition 2.26). By Definition 3.3 we have $A(B^E) = A \setminus E^{\oplus}_B$. Hence $E^{\oplus}_B = A$ is equivalent to $A(B^E) = \emptyset$, i.e. $B^E = (\emptyset, \emptyset, \emptyset)$.*

Proof. ($\Rightarrow$) Let $E \in$ b-st(B). By Theorem 3.2, there exists a serialisation sequence $\mathcal{S}$ with result E, and by the above equivalence we have $B^E = (\emptyset, \emptyset, \emptyset)$.

($\Leftarrow$) Let $\mathcal{S}$ be a serialisation sequence with result E such that $B^E = (\emptyset, \emptyset, \emptyset)$. By Theorem 3.2 we obtain $E \in$ b-ad(B). Again using the above equivalence, $B^E = (\emptyset, \emptyset, \emptyset)$ implies $E^{\oplus}_B = A$, hence $E \in$ b-st(B). □

Corollary 3.6 (Coherently Stable Characterisation). *Let* $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ *and* $E \subseteq A$. $E \in$ *b-st*(B) *iff either* $E = \emptyset$ *and* $A = \emptyset$, *or* $E = S_1 \cup S_2$ *where* $S_1 \in IS(B)$ *and* $S_2 \in$ *b-st*(B^{S_1}).

Proof. Follows from Theorem 3.6 by recursive application. □

This mirrors Thimm's AF results for grounded and preferred semantics (cf. [14, Theorems 4–6]), with defence replaced by strong defence and initial sets interpreted in the BAF sense.

3.5.3 Operational View and Selection Functions

Theorem 3.7 (Serialisation Characterisation for BAFs). *Let* $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$, $E \subseteq A$, *and* $\sigma \in \{\text{b-ad}, \text{b-co}, \text{b-gr}, \text{b-pr}, \text{b-st}\}$. *Then* $E \in \sigma(B)$ *(cf. Definition 2.26) iff there exists a serialisation sequence* $\mathcal{S} = (S_1, \ldots, S_n)$ *with result E such that, for* $E_i := S_1 \cup \cdots \cup S_i$ $(i = 1, \ldots, n)$ *and* $E_0 := \emptyset$, *the following holds:*

$$
\begin{aligned}
\sigma = \text{b-ad}: &\quad \text{no additional condition,} \\
\sigma = \text{b-co}: &\quad IS^{\not\leftarrow}(B^E) = \emptyset, \\
\sigma = \text{b-gr}: &\quad \forall i \in \{1, \ldots, n\} : \; S_i \in IS^{\not\leftarrow}(B^{E_{i-1}}) \;\; \text{and} \;\; IS^{\not\leftarrow}(B^E) = \emptyset, \\
\sigma = \text{b-pr}: &\quad IS(B^E) = \emptyset, \\
\sigma = \text{b-st}: &\quad B^E = (\emptyset, \emptyset, \emptyset).
\end{aligned}
$$

Proof of Theorem 3.7. The case σ = b-ad follows from Theorem 3.2. The cases $\sigma \in \{\text{b-co}, \text{b-gr}, \text{b-pr}, \text{b-st}\}$ follow from Theorems 3.3–3.6 respectively, by unfolding Definition 2.26 and rewriting the conditions in terms of the final reduct B^E and the partial results E_i. □

Remark 3.9 (Serialisation Characterisation for BAFs and AFs). *Each case corresponds to the AF characterisation in Theorem 2.1 by replacing classical concepts with their coherent counterparts (defence* $\rightsquigarrow$ *strong defence,* $R \rightsquigarrow R^B_{\text{der}}$*, ad* $\rightsquigarrow$ *b-ad) and by using BAF-initial sets* $IS(B)$ *in place of* $IS(AF)$.

Having established the serialisation conditions for all coherent semantics, we now turn to concrete examples that instantiate these conditions.

Example 3.9 (Serialisation to a Coherently Preferred Extension). *A coherently preferred extension is obtained by iterated initial-set selection until no initial sets remain, i.e. until* $IS(B^E) = \emptyset$.

Consider the bipolar argumentation framework $B_8 = (A, R, U)$ *where* $A = \{a, b, c, d, e, f\}$, $R = \{(b, c), (d, c), (e, f)\}$, *and* $U = \{(a, b), (a, d)\}$.

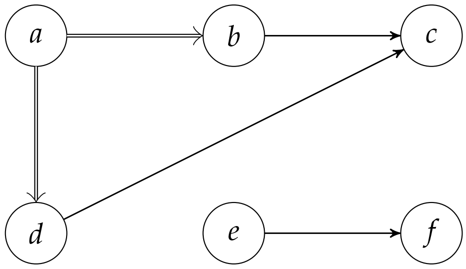

Figure 9: BAF B_8 for Example 3.9

The derived relations are

$$R_u^{B_8} = (U \circ R) \cup R = \{(a,c),(b,c),(d,c),(e,f)\},$$
$$R_m^{B_8} = R \circ U = \emptyset.$$

More generally, supported attacks are induced by (possibly empty) support chains ending in an attack, i.e. by the construction in Definition 2.20. In the present example, this yields exactly $(U \circ R) \cup R$.

In B_8*, the initial sets are* $IS(B_8) = \{\{a\},\{e\}\}$*, both unattacked.*

(1): *Choose* $S_1 = \{a\}$. *Then* $\{a\}^{\oplus}_{B_8} = \{a\} \cup \{c\} = \{a,c\}$*, and the reduct is* $B_8^{\{a\}} = B_8|_{A\setminus\{a,c\}} = (A_1, R_1, U_1)$ *with* $A_1 = \{b,d,e,f\}$, $R_1 = \{(e,f)\}$, $U_1 = \emptyset$. *Removing a breaks the support chains* (a,b) *and* (a,d)*, hence the supported attack* (a,c) *vanishes in the reduct. The derived relations in* $B_8^{\{a\}}$ *are* $R_u^{B_8^{\{a\}}} = \{(e,f)\}$ *and* $R_m^{B_8^{\{a\}}} = \emptyset$.

(2): *In* $B_8^{\{a\}}$*, the initial sets are* $IS(B_8^{\{a\}}) = \{\{b\},\{d\},\{e\}\}$*, all unattacked. Since* $IS(B_8^{\{a\}}) \neq \emptyset$*, we continue. Choose* $S_2 = \{e\}$. *Then* $\{e\}^{\oplus}_{B_8^{\{a\}}} = \{e,f\}$*, and the reduct is* $B_8^{\{a,e\}} = B_8|_{A\setminus\{a,c,e,f\}} = (A_2, R_2, U_2)$ *with* $A_2 = \{b,d\}$, $R_2 = \emptyset$, $U_2 = \emptyset$. *Hence* $R_u^{B_8^{\{a,e\}}} = R_m^{B_8^{\{a,e\}}} = \emptyset$.

(3): *In* $B_8^{\{a,e\}}$*, the initial sets are* $IS(B_8^{\{a,e\}}) = \{\{b\},\{d\}\}$. *Choose* $S_3 = \{b\}$. *Then* $B_8^{\{a,e,b\}} = (A_3, R_3, U_3)$ *with* $A_3 = \{d\}$ *and* $R_3 = U_3 = \emptyset$*, hence* $IS(B_8^{\{a,e,b\}}) = \{\{d\}\}$. *Choose* $S_4 = \{d\}$. *Then* $B_8^{\{a,e,b,d\}} = (\emptyset, \emptyset, \emptyset)$*, hence the serialisation terminates.*

The serialisation sequence $\mathcal{S} = (\{a\},\{e\},\{b\},\{d\})$ *yields* $E = \{a,e,b,d\}$. *Since* $IS(B_8^E) = \emptyset$*, the termination condition for b-pr-serialisations is satisfied (cf. Theorem 3.7), hence* $E \in$ *b-pr*(B_8).

Verification: adding c violates coherence (since $(a,c),(b,c),(d,c) \in R_u^{B_8}$*), and adding f violates coherence (since* $(e,f) \in R_u^{B_8}$*).*

Example 3.10 (Coherently Grounded vs. Coherently Preferred Serialisation)**.** *The coherently grounded and coherently preferred semantics differ only in their selection and ter-*

mination conditions. In particular, coherently grounded serialisations restrict selections to unattacked initial sets, whereas preferred serialisations continue until no initial set remains. Thus, semantics can be characterised by their selection and termination criteria [14].

Consider the bipolar argumentation framework $B_9 = (A, R, U)$ *where* $A = \{a, b, c\}$, $R = \{(b, c), (c, b)\}$, *and* $U = \{(a, b)\}$.

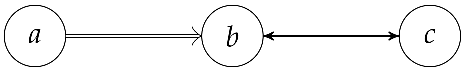

Figure 10: BAF B_9 from Example 3.10

In B_9, *the initial sets are* $IS(B_9) = \{\{a\}, \{b\}, \{c\}\}$. *Classification:* $\{a\}$ *is* unattacked. $\{b\}$ *is* challenged *(attacked by c via* $(c, b) \in R_u^{B_9}$, *and* $\{c\} \in IS(B_9)$*).* $\{c\}$ *is* challenged *(attacked by a and b via* $(a, c), (b, c) \in R_u^{B_9}$*). Thus* $IS^{\not\leftarrow}(B_9) = \{\{a\}\}$ *and* $IS^{\leftrightarrow}(B_9) = \{\{b\}, \{c\}\}$. *In particular,* $\{b\}$ *and* $\{c\}$ *challenge each other, as each attacks the other via* $R_u^{B_9}$.

Grounded serialisation (b-gr). *For grounded semantics, only* unattacked *initial sets may be selected at each step, i.e.* $\alpha_{b\text{-}gr}(X, Y, Z) = X$ *where* $X = IS^{\not\leftarrow}(B)$.

(1): *In* B_9, *the only unattacked initial set is* $\{a\}$. *Select* $S_1 = \{a\}$. *Then* $\{a\}^{\oplus}_{B_9} = \{a, c\}$ *(since* $(a, c) \in R_u^{B_9}$*), and the reduct is* $B_9^{\{a\}} = B_9|_{A\setminus\{a,c\}} = (A_1, R_1, U_1)$ *where* $A_1 = \{b\}$, $R_1 = \emptyset$, $U_1 = \emptyset$ *(support* (a, b) *removed). The derived relations in* $B_9^{\{a\}}$ *are* $R_u^{B_9^{\{a\}}} = R_m^{B_9^{\{a\}}} = \emptyset$.

(2): *In* $B_9^{\{a\}}$, *the initial sets are* $IS^{\not\leftarrow}(B_9^{\{a\}}) = \{\{b\}\}$. *Select* $S_2 = \{b\}$. *Then* $B_9^{\{a,b\}} = B_9|_{\emptyset} = (\emptyset, \emptyset, \emptyset)$, *hence* $IS^{\not\leftarrow}(B_9^{\{a,b\}}) = \emptyset$ *and the serialisation terminates.*

The grounded serialisation is $\mathcal{S}_{gr} = (\{a\}, \{b\})$ *with result* $E_{gr} = \{a, b\}$. *To verify that this coincides with the least fixed point characterisation of grounded semantics (Definition 2.26), we compute* τ_{B_9} *from Definition 2.25. The equalities can be checked by inspecting* $R_{\text{der}}^{B_9}$ *and applying strong defence (Definition 2.23).*

(i) $\tau_{B_9}(\emptyset) = \{a\}$,

(ii) $\tau_{B_9}(\{a\}) = \{a, b\}$,

(iii) $\tau_{B_9}(\{a, b\}) = \{a, b\}$.

Therefore $E_{gr} = \{a, b\}$ *is the least fixed point of* τ_{B_9}.

Here, the equalities can be checked by inspecting the derived attackers in $R_{\text{der}}^{B_9}$ *and applying the strong-defence condition from Definition 2.23.*

Preferred serialisation (b-pr). *For preferred semantics, we may select* any *initial set at each step:* $\alpha_{b\text{-}pr}(X, Y, Z) = X \cup Y \cup Z$. *We continue until no initial sets remain:* $\beta_{b\text{-}pr}(B, S) = 1$ *iff* $IS(B) = \emptyset$.

(1)*: Select* $S_1 = \{a\}$, *then* $S_2 = \{b\}$, *yielding* $E_1 = \{a, b\}$ *(same as grounded).*

(2)*: Select $S_1 = \{c\}$ (a challenged initial set). Then $\{c\}^{\oplus}_{B_9} = \{b, c\}$ (since $(c, b) \in R_u^{B_9}$), and the reduct is $B_9^{\{c\}} = B_9|_{A\setminus\{b,c\}} = (\{a\}, \emptyset, \emptyset)$. In $B_9^{\{c\}}$, the initial set is $IS(B_9^{\{c\}}) = \{\{a\}\}$. Select $S_2 = \{a\}$. Then $B_9^{\{c,a\}} = B_9|_{\emptyset} = (\emptyset, \emptyset, \emptyset)$, hence $IS(B_9^{\{c,a\}}) = \emptyset$ and the serialisation terminates.*

Thus we have two possible preferred serialisations:

$$\mathcal{S}_{pr}^{(1)} = \big(\{a\}, \{b\}\big) \quad \text{with } E_1 = \{a, b\}, \qquad \mathcal{S}_{pr}^{(2)} = \big(\{c\}, \{a\}\big) \quad \text{with } E_2 = \{a, c\}$$

Therefore b-pr$(B_9) = \{\{a, b\}, \{a, c\}\}$ (challenged selections allowed).

Both extensions are ⊆-maximal: $\{a, b\}$ cannot be extended (adding c violates coherence, as c is attacked via $(a, c), (b, c) \in R_u^{B_9}$). $\{a, c\}$ cannot be extended (adding b violates coherence, as b is attacked via $(c, b) \in R_u^{B_9}$).

The grounded extension $\{a, b\}$ is the ⊆-minimal complete extension (cf. Definition 2.26).

Definition 3.5 (BAF States and Transitions)**.** *A* state *is a pair $T = (B, S)$ where $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and $S \subseteq A$. A transition step from $T = (B, S)$ to $T' = (B', S')$ is written as*

$$(B, S) \xrightarrow{S_0} \big(B^{S_0},\, S \cup S_0\big) = (B', S')$$

where $S_0 \in IS(B)$.

Thus, in each step we select an initial set S_0 of the current framework and move to the reduct B^{S_0} (cf. Definition 3.3), while accumulating the chosen arguments in the second component.

To align our operational view with Thimm's reachability-based presentation, we introduce an explicit reachability notation on states. This is purely notational and will be used to abbreviate finite transition sequences.

Definition 3.6 (Reachability of states)**.** *For states $T = (B, S)$ and $T' = (B', S')$, we write $T \Rightarrow T'$ iff there exists some $S_0 \in IS(B)$ such that*

$$(B, S) \xrightarrow{S_0} \big(B^{S_0},\, S \cup S_0\big) = (B', S')$$

Let α be a selection function such that $\alpha(B, S) \subseteq IS(B)$ for all states (B, S). We write $T \Rightarrow_\alpha T'$ iff $T \Rightarrow T'$ via some $S_0 \in \alpha(B, S)$. We write $\Rightarrow^*$ and $\Rightarrow^*_\alpha$ for the reflexive–transitive closures (cf. Definition 2.19). Let β be a termination function on states. We write $T \Rightarrow^*_{\alpha,\beta} T'$ iff $T \Rightarrow^*_\alpha T'$ and $\beta(T') = 1$.

Finally, we formalise the selection and termination functions for BAF semantics.

Theorem 3.8 (BAF Semantics via Selection and Termination Functions)**.** *The following semantics are serialisable with the given selection and termination functions, extending the framework of Theorem 2.2. Let (X, Y, Z) be the partition $(IS^{\not\leftarrow}(B), IS^{\not\leftrightarrow}(B), IS^{\leftrightarrow}(B))$ from Definition 3.2. For brevity, we write $\alpha_\sigma(B, S) := \alpha_\sigma(X, Y, Z)$, where (X, Y, Z) is the partition of $IS(B)$ from Definition 3.2(hence $\alpha_\sigma(B, S) \subseteq IS(B)$):*

- *b-ad:* $\alpha_{b\text{-}ad}(X, Y, Z) = X \cup Y \cup Z$, *and* $\beta_{b\text{-}ad}(B, S) = 1$
- *b-co:* $\alpha_{b\text{-}co} = \alpha_{b\text{-}ad}$ *and* $\beta_{b\text{-}co}(B, S) = \begin{cases} 1 & \text{if } IS^{\not\leftarrow}(B) = \emptyset \\ 0 & \text{otherwise} \end{cases}$
- *b-gr:* $\alpha_{b\text{-}gr}(X, Y, Z) = X$ *and* $\beta_{b\text{-}gr}(B, S) = \beta_{b\text{-}co}(B, S)$
- *b-pr:* $\alpha_{b\text{-}pr} = \alpha_{b\text{-}ad}$ *and* $\beta_{b\text{-}pr}(B, S) = \begin{cases} 1 & \text{if } IS(B) = \emptyset \\ 0 & \text{otherwise} \end{cases}$
- *b-st:* $\alpha_{b\text{-}st} = \alpha_{b\text{-}ad}$ *and* $\beta_{b\text{-}st}(B, S) = \begin{cases} 1 & \text{if } B = (\emptyset, \emptyset, \emptyset) \\ 0 & \text{otherwise} \end{cases}$

In other words, $\alpha_{b\text{-}ad}$ *and* $\alpha_{b\text{-}pr}$ *permit the selection of* arbitrary *initial sets, i.e. any element of* $IS(B)$ *regardless of whether it is unattacked, unchallenged, or challenged, while* $\alpha_{b\text{-}gr}$ *restricts selections to unattacked initial sets* $IS^{\not\leftarrow}(B)$ *only.*

Proof. By Theorem 3.7, we have that $E \in \sigma(B)$ iff there exists a serialisation sequence $\mathcal{S}$ with result E whose intermediate steps and final reduct B^E satisfy the corresponding condition in the characterisation.

For each σ, the functions $(\alpha_\sigma, \beta_\sigma)$ enforce exactly these constraints:

- For b-ad, no additional projection is imposed, matching the results from Theorem 3.2.
- For b-co, the termination function $\beta_{\text{b-co}}$ stops the construction exactly in those terminal states (B, S) with $\text{IS}^{\not\leftarrow}(B) = \emptyset$. In particular, for (B^E, E) this yields $\text{IS}^{\not\leftarrow}(B^E) = \emptyset$, as required by Theorem 3.7.
- For b-gr, the selection function $\alpha_{\text{b-gr}}$ restricts choices to unattacked initial sets at each stage, while $\beta_{\text{b-gr}}$ ensures the grounded extension property.
- For b-pr, $\beta_{\text{b-pr}}$ stops the construction exactly in those terminal states (B, S) with $\text{IS}(B) = \emptyset$. In particular, for the terminal state (B^E, E) this yields $\text{IS}(B^E) = \emptyset$ as required by Theorem 3.7.
- For b-st, $\beta_{\text{b-st}}$ enforces that the final reduct is empty, mirroring the condition $B^E = (\emptyset, \emptyset, \emptyset)$ which is equivalent to $E^{\oplus}_B = A$.

Hence, the serialisable semantics induced by $(\alpha_\sigma, \beta_\sigma)$ coincide with the coherent-σ semantics for each σ. □

To conclude this section, we investigate how the selection and termination functions (α, β) defined in Theorem 3.8 govern the step-by-step construction of extensions. This example specifically highlights the mechanism of support-chain breakage.

For the following example, we use the notation $\mathcal{S}_\sigma$ to denote a serialisation sequence for the respective semantics and their selection and termination conditions.

Example 3.11 (Operational Selection and Support-Chain Breakage)**.** *Consider $B_{10} = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ shown in Figure 11 with $A = \{a, b, c, d, e\}$, $R = \{(c, b)\}$, and $U = \{(e, d), (d, c), (a, c)\}$. Here, c is the sole attacker of b, while c is independently supported by a and the support chain $e \xrightarrow{U} d$.*

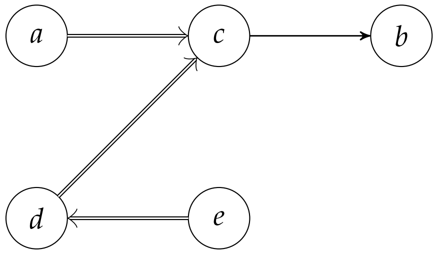

Figure 11: BAF B_{10} from Example 3.11.

The derived attack relations are $R_u^{B_{10}} = \{(c, b), (d, b), (e, b), (a, b)\}$ and $R_m^{B_{10}} = \emptyset$. Note that b is attacked by every other argument via supported attacks. The initial sets are $IS(B_{10}) = \{\{a\}, \{c\}, \{d\}, \{e\}\}$, all of which are unattacked.

We now trace the construction of a coherently preferred extension using the transition system $(B, S) \xrightarrow{\alpha} (B', S \cup S')$. Starting from the initial state $T_0 = (B_{10}, \emptyset)$:

1. *The selection function $\alpha_{b\text{-}pr}$ permits picking any initial set. We choose $S_1 = \{a\}$.*

$$(B_{10}, \emptyset) \xrightarrow{S_1=\{a\}} (B_{10}^{\{a\}}, \{a\})$$

The reduct $B_{10}^{\{a\}}$ removes a and its target b (since $(a, b) \in R_{\text{der}}^{B_{10}}$). The remaining arguments are $\{c, d, e\}$ with $U' = \{(e, d), (d, c)\}$. Since the attack on b is removed, no derived attacks remain in this reduct ($R_{\text{der}}^{B_{10}^{\{a\}}} = \emptyset$).

2. *In $B_{10}^{\{a\}}$, all remaining singletons are unattacked initial sets. We pick $S_2 = \{e\}$.*

$$(B_{10}^{\{a\}}, \{a\}) \xrightarrow{S_2=\{e\}} (B_{10}^{\{a,e\}}, \{a, e\})$$

The reduct $B_{10}^{\{a,e\}}$ consists of $\{c, d\}$. No attacks exist.

3. *Finally, we have $(B_{10}^{\{a,e\}}, \{a, e\})$ with $IS(B_{10}^{\{a,e\}}) = \{\{c\}, \{d\}\} \neq \emptyset$. According to the termination function for preferred semantics (cf. Theorem 3.8), we have $\beta_{b\text{-}pr}(B_{10}^{\{a,e\}}, \{a, e\}) = 0$ as long as $IS(B_{10}^{\{a,e\}}) \neq \emptyset$. Thus, the serialisation must continue. After two further steps selecting $\{c\}$ and $\{d\}$, we reach the state $((\emptyset, \emptyset, \emptyset), \{a, e, c, d\})$, where $IS(B_{10}^{\{a,e,c,d\}}) = \emptyset$ and hence $\beta_{b\text{-}pr}(B_{10}^{\{a,e,c,d\}}, \{a, e, c, d\}) = 1$.*

Let $T_1 := (B_{10}^{\{a\}}, \{a\})$ and $T_2 := (B_{10}^{\{a,e\}}, \{a, e\})$. After two further steps we reach $T_4 := ((\emptyset, \emptyset, \emptyset), \{a, e, c, d\})$. Hence $T_0 \Rightarrow^*_{\alpha_{\text{b-pr}}} T_4$ and $\beta_{\text{b-pr}}(T_4) = 1$, i.e. $T_0 \Rightarrow^*_{\alpha_{\text{b-pr}}, \beta_{\text{b-pr}}} T_4$.

The resulting extension $E = \{a, e, c, d\}$ is the unique coherently preferred extension. Alternatively, a serialisation starting with $S_1 = \{e\}$ first removes e and b, then proceeds analogously by selecting remaining initial sets until all arguments are exhausted, again yielding the preferred extension $E = \{a, e, c, d\}$.

In both serialisations, b is removed in the first reduction step because $b \in S^{+}_{1,B_{10}}$, i.e. b lies in the range of the first selected singleton wrt. the derived attacks of B_{10}. More precisely, selecting $S_1 = \{a\}$ yields $(a, b) \in R_u^{B_{10}} = U^* \circ R$ via the support chain $a \xrightarrow{U} c$ followed by $c \xrightarrow{R} b$, and selecting $S_1 = \{e\}$ yields $(e, b) \in R_u^{B_{10}}$ via $e \xrightarrow{U} d \xrightarrow{U} c \xrightarrow{R} b$. After taking the reduct, the support relation is restricted to the remaining carrier, so support chains may disappear and the derived relations R_u and R_m must be recomputed in the reduct (cf. Remark 3.4). Thus, b is eliminated in the first reduction step regardless of whether $S_1 = \{a\}$ or $S_1 = \{e\}$ is selected. This example shows that the α/β-construction operates on $S_B^{\oplus}$ computed from the current derived relations, which may change after each reduct due to support-chain breakage. Thus, derived relations must be recomputed after each reduct (cf. Remark 3.4).

The functions α and β provide the formal language to track these subtle dependencies that distinguish BAFs from classical AFs. This operational construction mirrors Thimm's serialisation view for AFs. Extensions arise as sets accumulated along finite state transitions to a β-terminal state, with the key difference that reducts in BAFs require recomputing derived attacks after support-chain breakage.

Remark 3.10 (Comparison to AF characterisation)**.** *Theorem 3.8 generalises Thimm's AF results for admissibility-based semantics [14] to the bipolar setting. In the AF case, each semantics is proved separately via selection and termination functions (α, β), and the corresponding reduct-based characterisations are then derived as corollaries. In contrast, we first establish a unified serialisation characterisation for BAFs (via the conditions in Theorem 3.7) and then give dedicated proofs for the coherent semantics (b-co, b-gr, b-pr, b-st) on the semantic level (coherent admissibility, completeness, groundedness, etc.). Theorem 3.8 subsequently extracts the explicit (α, β)-functions from this characterisation, providing a direct bridge to Thimm's construction principle and supporting algorithmic implementations.*

3.6 Computational Complexity

This subsection focuses on the decision problems needed later for our reduct- and initial-set-based constructions, namely coherent admissibility verification and the existence and uniqueness of initial sets. Upper bounds are obtained by transferring the corresponding AF results along the embedding $AF \mapsto (A, R, \emptyset)$, while hardness carries over via the same embedding.

Throughout this section, we use the polynomial-time embedding of AFs into BAFs given by

$$AF = (A, R) \longmapsto B_{AF} = (A, R, \emptyset)$$

By Proposition 2.6, the corresponding AF notions are preserved when $U = \emptyset$. Hence, every lower bound for the AF case transfers to the BAF case via this mapping.

Proposition 3.3 (Complexity of Coherent Admissibility Verification)**.** *Deciding whether $S \in b\text{-}ad(B)$ is in* **P**.

Proof. Let $B = (A, R, U) \in \mathfrak{B}_{\mathfrak{A}}$ and $S \subseteq A$. Compute the reachability relation $U^+ \subseteq A \times A$ (non-empty U-paths) by graph reachability in (A, U), e.g. by running BFS/DFS from each node. Let $I := \{(a, a) \mid a \in A\}$ and set $U^* := I \cup U^+$.

Define the supported and secondary attack relations as $R_u^B := U^* \circ R$ and $R_m^B := R \circ U^+$. Then the derived attack relation $R_{\text{der}}^B := R_u^B \cup R_m^B$ is computable in polynomial time.

Given R_{der}^B, strong conflict-freeness and coherence of S can be checked in polynomial time, and strong defence can be verified by checking for every $(b, a) \in R_{\text{der}}^B$ with $a \in S$ whether there exists $c \in S$ with $(c, b) \in R_{\text{der}}^B$ (Definitions 2.22 and 2.23). □

Proposition 3.4 (Complexity of Initial Set Existence)**.** *Deciding whether $IS(B) \neq \emptyset$ is* **NP***-hard. Moreover, restricted to embedded instances $B_{AF} = (A, R, \emptyset)$, the problem is* **NP***-complete.*

Proof. **NP***-hardness.* **NP**-hardness follows from the AF case [14, Prop. 6] via the embedding $AF \mapsto (A, R, \emptyset)$ and Proposition 2.6, which preserves initial sets.

Membership in **NP** *(embedded instances).* We show membership by transferring the **NP**-certificate from the AF case. Let $AF = (A, R)$ and consider the embedded instance $B_{AF} = (A, R, \emptyset)$. By [14, Prop. 6], the AF problem IS(AF) $\neq \emptyset$ admits an **NP**-certificate whose correctness can be verified in polynomial time. By Proposition 2.6, admissibility and initial sets in B_{AF} coincide with the AF notions, and all admissibility checks required by the verifier are decidable in polynomial time by Proposition 3.3. Hence the problem is in **NP** on embedded instances. □

Remark 3.11 (Uniqueness)**.** *We write* $|IS(B)| = 1$ *to denote that there exists* exactly one *initial set in* B.

Proposition 3.5 (Complexity of Initial Set Uniqueness)**.** *Deciding whether* $|IS(B)| = 1$ *is* **DP***-hard. Moreover, restricted to embedded instances* $B_{AF} = (A, R, \emptyset)$*, the problem is* **DP***-complete.*

Proof. Membership in **DP** *(embedded instances).* We have $|\text{IS}(B_{AF})| = 1$ iff (i) $\text{IS}(B_{AF}) \neq \emptyset$ and (ii) $|\text{IS}(B_{AF})| \leq 1$. Condition (i) is in **NP** by Proposition 3.4 (restricted to embedded instances). For (ii), consider the complement $|\text{IS}(B_{AF})| \geq 2$, which is in **NP**. guess two distinct sets $S_1 \neq S_2$ and verify $S_1, S_2 \in \text{IS}(B_{AF})$. For embedded instances $B_{AF} = (A, R, \emptyset)$, initial-set verification is polynomial-time decidable in AFs (cf. [14, Prop. 5]), and the admissibility checks required there are decidable in polynomial time by Proposition 3.3. Thus (ii) is in **coNP**, and the conjunction is in **DP**.

DP*-hardness.* Hardness transfers from the AF case [14, Prop. 7] via the embedding $AF \mapsto (A, R, \emptyset)$ and Proposition 2.6. □

Remark 3.12 (Acceptance under Preferred Semantics (AF benchmark))**.** *In the AF setting, skeptical acceptance wrt. preferred semantics is stated as* Π_2^P*-complete in [15] (citing Dunne and Bench-Capon [11]). We mention this only as an AF benchmark. A dedicated complexity analysis of acceptance in BAFs is left for future work.*

Table 3: Complexity of core decision problems related to coherent admissibility and initial sets in BAFs.

Problem	**Complexity**	**Reference**
Verify $S \in \text{b-ad}(B)$	**P**	Proposition 3.3
$\text{IS}(B) \neq \emptyset$?	**NP**-hard; **NP**-complete for B_{AF}	Proposition 3.4
Uniqueness: $|\text{IS}(B)| = 1$	**DP**-hard; **DP**-complete for B_{AF}	Proposition 3.5

Proposition 3.6 (Complexity Transfer for Initial-Set Problems)**.** *For the decision problems in Proposition 3.3, Proposition 3.4, and Proposition 3.5, all AF lower bounds transfer to BAFs via the embedding* $AF \mapsto (A, R, \emptyset)$*. Moreover, for embedded instances* $B_{AF} = (A, R, \emptyset)$ *the corresponding AF and BAF notions coincide, and hence the AF upper bounds apply on that sub-class.*

Proof. The embedding $AF \mapsto (A, R, \emptyset)$ is polynomial-time, and by Proposition 2.6 it preserves the relevant notions when $U = \emptyset$. Therefore, every AF hardness result yields the same hardness for the corresponding BAF problem. For embedded instances $B_{AF} = (A, R, \emptyset)$, the problems coincide with the AF versions, and the upper bounds from [14, Prop. 5–7] apply. Additionally, coherent admissibility verification for general BAFs is in **P** by Proposition 3.3. □

In summary, verifying coherent admissibility is polynomial-time decidable for general BAFs. For initial-set existence and uniqueness, we obtain AF-tight classifications on the embedded sub-class $B_{AF} = (A, R, \emptyset)$ and inherit the corresponding AF hardness for general BAFs via the embedding. A full membership classification of these initial-set problems for arbitrary BAFs is beyond the scope of this thesis.

3.7 Relationship to Classical Argumentation Frameworks

Our serialisation framework represents a genuine generalisation of Thimm's AF approach [14]. This relationship is characterised by three central observations regarding compatibility, structural dynamics, and complexity (cf. Remark 3.10).

First, the framework is *backwards compatible*. As established in Proposition 2.6 and Proposition 3.2, for any BAF where the support relation $U = \emptyset$, all bipolar notions (coherent admissibility, initial sets, reducts, and resulting extensions) collapse to their classical AF counterparts, and the serialisation structure is preserved. In particular, admissibility coincides with coherent admissibility (via Proposition 2.5), initial sets coincide, and AF reducts agree with the corresponding BAF reducts when interpreted as AFs. Thus, any serialisation sequence valid in an AF remains valid when interpreted as a BAF without supports.

Second, the introduction of support relations leads to a *non-monotonic disruption* of attack paths, a phenomenon absent in classical AFs. In AF serialisation, removing an argument in a reduct only eliminates attacks. In the BAF setting, however, removing an argument can break a support chain, which in turn causes a mediated attack (supported or secondary) to vanish. This structural effect, shown in Examples 3.9 and 3.11, necessitates the re-computation of derived relations R^B_{der} at each stage of the construction (cf. Definitions 2.20 and 2.22 and Remark 3.3).

Finally, while the interaction of support and attack significantly enriches the semantic landscape, the fundamental *computational complexity* of the core decision problems we consider remains invariant. As analysed in Section 3.6 and Proposition 3.6, the lower bounds for AF decision problems transfer to BAFs via the embedding $AF \mapsto (A, R, \emptyset)$, and coherent admissibility verification is still polynomial-time decidable (cf. Proposition 3.3). The combinatorial difficulty of finding and distinguishing extensions is thus still rooted in the underlying conflict structure, rather than in the mere presence of support.

Remark 3.13 (SCC structure in BAFs)**.** *In AFs, initial sets admit a structural analysis via SCCs [14]. For BAFs, however, SCC-based decompositions are less immediate, since SCCs could be defined wrt. R, $R \cup U$, or R^B_{der}, leading to different graph structures. We leave a systematic SCC-based analysis for future work.*

4 Conclusion and Future Work

4.1 Summary of Contributions

This thesis developed a serialisation-based characterisation of coherence-based semantics for bipolar argumentation frameworks (BAFs) under the deductive interpretation of support. In the classical AF setting, serialisation provides a constructive view on admissibility-based semantics. Extensions can be built stepwise by repeatedly selecting inclusion-minimal admissible building blocks (initial sets) and continuing the construction in a suitable reduct. Our main goal was to lift this perspective to BAFs, where the additional support relation interacts with attack and thereby changes how acceptability is determined.

The transfer from AFs to BAFs is not a syntactic rewrite, because in BAFs the relevant notion of attack is no longer given by the base relation R alone. Instead, supported or secondary attacks induce *derived attacks*, which we capture by the relation R^B_{der} (cf. Definition 2.20 and Notation 2.3). As a consequence, the AF-style iterative construction via initial-set selection and reduct formation is technically more involved in the bipolar setting. Removing arguments may break support chains, and therefore the derived attack relation in the next reduct is *not* simply the restriction of the old one. This support-chain breakage is a genuinely bipolar phenomenon and it forces a careful choice of the construction primitives, namely the initial sets, and of the reduct operation itself. In particular, for a reduct one has to restrict not only the carrier but *both* relations R and U, and then recompute the derived relations in the restricted framework (cf. Remarks 3.3 and 3.4).

The core idea of the thesis is a systematic translation pattern from AF concepts to bipolar counterparts that preserves the serialisation viewpoint while remaining faithful to coherence. Concretely, we replace the base attack relation by the derived attack relation, and classical admissibility by *coherent admissibility* (coherence plus strong defence, cf. Definitions 2.22 to 2.24). On that basis, we define bipolar initial sets as the non-empty $\subseteq$-minimal coherently admissible sets (Definition 3.1) and bipolar reducts by removing the range, denoted by $S^{\oplus}_B$, i.e. S together with all arguments attacked by S via supported or secondary attacks, and then projecting the framework to the remaining carrier (Definition 3.3). A natural alternative would have been to characterise these notions directly in terms of R and U without explicitly forming R^B_{der} first. However, coherence verification already requires tracking support chains, so such a reformulation would not simplify the technical core while it would obscure the relationship to the standard definitions and to Thimm's AF serialisation framework (Remark 3.1). At the same time, the resulting bipolar definitions form a conservative extension of the AF case. When $U = \emptyset$, derived attacks collapse to direct attacks and the bipolar reduct coincides with the classical AF reduct, so all bipolar notions reduce to their AF counterparts (cf. Propositions 2.6 and 3.2 and Remark 3.3).

With bipolar initial sets and reducts in place, we established the key prerequisite for serialisation. Coherent admissibility admits a reduct-based decomposition into initial sets. A crucial result provides a reduct-based decomposition of coherent admissibility (Theorem 3.1). Every non-empty coherently admissible set can be obtained by first choosing an initial set and then continuing with a coherently admissible remainder in the corresponding reduct. Iterating this one-step decomposition yields a finite recursive decomposition into a sequence of initial sets in successive reducts (Corollary 3.2), providing an explicit constructive reading of coherent admissibility. Technically, a recurring theme in these proofs is that projections and reducts must not introduce spurious derived attacks. This is captured by the monotonicity properties of derived relations under projection (Lemma 3.1 and Corollary 3.1) and the accompanying strictness discussion (Remark 3.5).

Building on the decomposition results, we lifted Thimm's selection and termination perspective to the bipolar setting.[2] Operationally, this means that extensions can be understood as results of finite transition sequences that iteratively select eligible initial sets and compute reducts, terminating exactly when the respective termination condition for the target semantics is met. The worked constructions also highlight a point absent in AF serialisation. After each reduct, derived attack relation must be recomputed, since reduct steps may break support chains that previously established supported or secondary attacks (Remark 3.4).

Overall, we establish uniform serialisation characterisations for the coherence-based semantics studied in this thesis (coherently admissible, complete, grounded, preferred, and stable semantics) via serialisation sequences, reduct conditions, and corresponding selection and termination functions (α, β) (Section 3.5). On the computational side, we show that verification of coherent admissibility is decidable in polynomial time (Proposition 3.3). Moreover, initial-set existence and uniqueness remain hard, with completeness results already on embedded AF instances $B_{AF} = (A, R, \emptyset)$ (Proposition 3.6). Thus, deductive support enriches the modelling setting without increasing the asymptotic complexity of the core problems considered here.

[2]See in particular the serialisation characterisation results and the corresponding (α, β) selection and termination functions in Section 3.5.

4.2 Future Research

Several directions remain open for research.

Alternative characterisations and structure. While we intentionally based our development on derived attacks (R^B_{der}) to align with standard bipolar definitions and to enable a clean serialisation transfer, it remains interesting to explore alternative characterisations that work directly with R and U while exposing additional structure, for instance by identifying fragments where coherence checking and initial-set identification admit simpler sufficient conditions (Remark 3.1). Likewise, the SCC-based (strongly connected component, Definition 2.10) analysis of initial sets is well developed for AFs. Extending and exploiting a comparable structural theory in BAFs could enable decomposition principles and parallel computation schemes.

Algorithms and implementations. The (α, β) characterisations suggest concrete algorithms that construct extensions by iteratively selecting initial sets and computing reducts. A next step is to implement these procedures for large BAFs and evaluate them empirically, including optimisations such as incremental maintenance of derived attacks across reduct steps, SCC-based decomposition, and heuristics for initial-set selection (cf. Example 3.11 for the operational relevance of support-chain breakage).

Complexity beyond the core problems. This thesis established polynomial-time verification for coherent admissibility and transferred hardness results for initial-set existence and uniqueness via the AF embedding (Propositions 3.3 and 3.6). A more complete complexity landscape for BAFs remains to be investigated, in particular for acceptance problems under coherently preferred semantics and for parameterised complexity with respect to structural measures of (A, R, U) (cf. Remark 3.12 for the intended scope distinction).

Alternative support interpretations. Finally, the present development focuses on the deductive reading of support. Adapting serialisation to other standard support interpretations, such as necessary or evidential support, would clarify which parts of the translation pattern are robust and which parts are tied to the deductive interaction between support and attack (cf. the discussion around support interpretations in Section 2.6).

References

[1] Leila Amgoud, Claudette Cayrol, and Marie-Christine Lagasquie-Schiex. On the bipolarity in argumentation frameworks. In *Proceedings of the 10th International Workshop on Non-Monotonic Reasoning (NMR 2004)*, pages 1–9, Whistler, Canada, 2004. University of British Columbia.

[2] Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert van der Torre. Abstract argumentation frameworks and their semantics. In *Handbook of Formal Argumentation*, volume 1, pages 159–236. College Publications, 2018.

[3] Ringo Baumann, Gerhard Brewka, and Markus Ulbricht. Revisiting the foundations of abstract argumentation – semantics based on weak admissibility and weak defense. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI'20)*, pages 2742–2749, New York, NY, USA, 2020. AAAI Press.

[4] Lars Bengel, Julian Sander, and Matthias Thimm. Characterising serialisation equivalence for abstract argumentation. In *Proceedings of the European Conference on Artificial Intelligence (ECAI 2024)*, volume 392 of *Frontiers in Artificial Intelligence and Applications*, pages 3340–3347. IOS Press, 2024.

[5] Lydia Blümel and Matthias Thimm. A ranking semantics for abstract argumentation based on serialisability. In H. Prakken and N. Green, editors, *Proceedings of the 8th International Conference on Computational Models of Argumentation (COMMA 2022)*, volume 353 of *Frontiers in Artificial Intelligence and Applications*, pages 80–91. IOS Press, 2022.

[6] Claudette Cayrol, Andrea Cohen, and Marie-Christine Lagasquie-Schiex. Higher-order interactions (bipolar or not) in abstract argumentation: A state of the art. In Dov Gabbay, Massimiliano Giacomin, Guillermo R. Simari, and Matthias Thimm, editors, *Handbook of Formal Argumentation*, volume 2, chapter 1. College Publications, 2021.

[7] Claudette Cayrol and Marie-Christine Lagasquie-Schiex. On the acceptability of arguments in bipolar argumentation frameworks. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty, 8th European Conference, ECSQARU 2005*, pages 378–389, Barcelona, Spain, July 2005. Springer.

[8] Claudette Cayrol and Marie-Christine Lagasquie-Schiex. Bipolar abstract argumentation systems. In *Argumentation in Artificial Intelligence*, pages 65–84. Springer, 2009.

[9] Claudette Cayrol and Marie-Christine Lagasquie-Schiex. Bipolarity in argumentation graphs: Towards a better understanding. *International Journal of Approximate Reasoning*, 54(7):876–899, 2013. Special issue: Uncertainty in Artificial Intelligence and Databases.

[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):321–358, 1995.

[11] Paul E. Dunne and Trevor J. M. Bench-Capon. Coherence in finite argument systems. *Artificial Intelligence*, 141(1–2):187–203, 2002.

[12] Nikos Karacapilidis and Dimitris Papadias. Computer supported argumentation and collaborative decision making: the hermes system. *Information Systems*, 26(4):259–277, 2001.

[13] Kenneth H. Rosen. *Discrete Mathematics and Its Applications*. McGraw-Hill Education, New York, 7th edition, 2012.

[14] Matthias Thimm. Revisiting initial sets in abstract argumentation. *Argument & Computation*, 13(3):325–360, 2022.

[15] Matthias Thimm, Federico Cerutti, and Mauro Vallati. Skeptical reasoning with preferred semantics in abstract argumentation without computing preferred extensions. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI 2021)*, pages 2069–2075. International Joint Conferences on Artificial Intelligence Organization, 2021.

[16] Bart Verheij. Deflog: on the logical interpretation of prima facie justified assumptions. *Journal of Logic and Computation*, 13(3):319–346, 2003.

[17] Yuming Xu and Claudette Cayrol. Initial sets in abstract argumentation frameworks. *Journal of Applied Non-Classical Logics*, 28(2–3):260–279, 2018.

[18] Yuming Xu, Lidong Xu, and Claudette Cayrol. On structural analysis of extension-based argumentation semantics. In *Proceedings of the Second Chinese Conference on Logic and Argumentation (CLAR'18)*, Lecture Notes in Computer Science. Springer, 2018.

[19] Yunshen Xu and Claudette Cayrol. Initial sets in abstract argumentation frameworks. In *Proceedings of the 1st Chinese Conference on Logic and Argumentation (CLAR '16)*, volume 1811 of *Lecture Notes in Computer Science*, pages 72–85. Springer, 2016.