Rewriting Rules for the Computation of Goal-Oriented Changes in an Argumentation System

Dionysios Kontarinis[1], Elise Bonzon[1], Nicolas Maudet[2],
Alan Perotti[3], Leendert van der Torre[4], Serena Villata[5]

[1] LIPADE, Université Paris Descartes,
[2] LIP6, Université Pierre et Marie Curie,
[3] Turin University,
[4] University of Luxembourg,
[5] INRIA, Sophia Antipolis

16 september, 2013

A type of multi-agent debate

- Every agent has a private argumentation system (AS).
- The debate has the form of a single argumentation system, **the Gameboard (GB)**, modified whenever agents insert new arguments and attacks, or state their opinions on arguments and attacks.
- Agents focus on the status of a single argument, called the **issue**.
- In order to satisfy an argumentative goal, an agent must know how to modify the GB.
- The **minimal change** which satisfies a goal is worth studying because:
  - It requires the least amount of effort (eg. resources, time).
  - It is simple and highlights the key points of the debate.

Research questions

- Define an argumentation framework where it is possible to reason efficiently about **change**.
- Study the properties of **minimal change** achieving a given goal.
- Define and implement a **procedure computing the minimal change** for a given goal.
- Ensure some wished properties of the above procedure.

General Framework

In this work we are using:

- Abstract argumentation (the structure of arguments is not specified).
- Extension-based acceptability (admissible, preferred, complete and grounded extensions).
- Both credulous and sceptical acceptability.
- Attack semantics.

Acceptability - Extensions

Example

Let $AS = \langle A, R \rangle$ be a system such that:
$A = \{a, b, c, d, e\}$, $R = \{(a, b), (b, a), (c, d), (e, d)\}$.

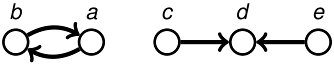

Admissible extensions (defend themselves against attacks):
eg. $\{b\}$, $\{a\}$, $\{b, c\}$, $\{b, c, e\}$
Preferred extensions (admissible and $\subseteq$-maximal):
$\{a, c, e\}$, $\{b, c, e\}$

Acceptability - Extensions

Example

Let $AS = \langle A, R \rangle$ be a system such that:
$A = \{a, b, c, d, e\}$, $R = \{(a, b), (b, a), (c, d), (e, d)\}$.

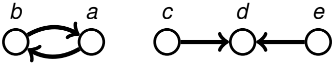

Complete extensions (admissible and contain all the arguments they defend):
$\{c, e\}$, $\{a, c, e\}$, $\{b, c, e\}$
Grounded extension (intersection of all complete extensions):
$\{c, e\}$

Acceptability - Extensions

Example

Let $AS = \langle A, R \rangle$ be a system such that:
$A = \{a, b, c, d, e\}$, $R = \{(a, b), (b, a), (c, d), (e, d)\}$.

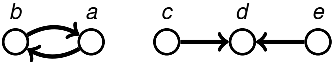

Preferred extensions: $\{a, c, e\}$, $\{b, c, e\}$

Under preferred semantics, argument *a* is:

- **Credulously** accepted (as $\exists$ preferred extension containing *a*).
- **Not sceptically** accepted (as it does not hold that "$\forall$ preferred extension, that extension contains *a*").

Acceptability - Extensions

Example

According to the **attack semantics**, in the previous example, there are 3 possible labellings:

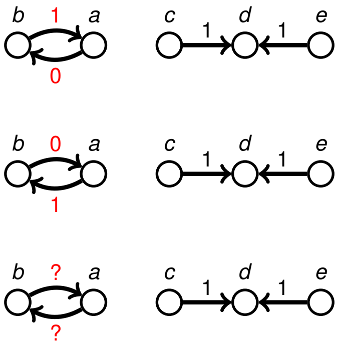

Disagreement over the existence of attacks

- In a debate, it is possible for two agents to have conflicting opinions over the existence of an attack.
- Example:
  - **Argument A**: For these reasons (...), **the measures cut public spending**.
  - **Argument B**: For these reasons (...), **the deficit will remain unchanged**.
  - The conflict between the two arguments seems to be debated!
  - An agent may believe that argument A attacks B.
  - Another agent may believe that argument A does not attack B.
- Such a disagreement in a debate can be resolved through voting.
- We will focus on these debated attacks.

The multi-agent debate setting

We consider debates with the following characteristics:

- A number of agents, each one with a private argumentation system (AS).
- A common argumentation system, called **Gameboard (GB)**, updated through the agents' locutions.
- Agents focus on the single **issue** of the debate, and try to make its status coincide, on the GB and on their private systems.

A modifiable Argumentation System

- $GB = \langle A, R, R^{+}, R^{-} \rangle$ is the debate's gameboard.
- Set of arguments A (abstract and fixed).
- Set of attacks R which are in the GB.
- Set of attacks which can be added (R^{+}) and removed (R^{-}) from the GB.

Example

Let $GB = \langle A, R, R^{+}, R^{-} \rangle$ be a system such that: $A = \{a, b, c, d, e\}$,
$R = \{(a,b),(b,a),(c,d),(e,d)\}$, $R^{+} = \{(a,c),(b,c)\}$, $R^{-} = \{(c,d),(e,d)\}$.

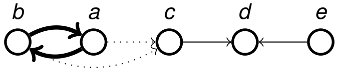

The notion of atom

Definition

Let $GB = \langle A, R, R^+, R^- \rangle$ be a system, and let $d \in A$, $x \in R$.

Types of Atoms	**Meaning**
$\top$ (resp. $\bot$)	success (resp. failure)
$PRO(d)$ (resp. $CON(d)$)	accept (resp. reject d)
$(x, +, \#)$ (resp. $(x, -, \#)$)	add (resp. remove) x
$(x, 1, \#)$ (resp. $(x, 0, \#)$, $(x, ?, \#)$)	x is '1' (resp. '0', '?')
$(x, 1, *)$ (resp. $(x, 0, *)$, $(x, ?, **)$, $(x, ?, *)$)	x must be '1' (resp. '0', '?', '?')

Move on a system

A **move** consists on adding and removing some attacks.

The set of atoms $m = \{(x, y, \#) \mid x \in R \text{ and } y \in \{+, -\}\}$ is a **move on** *GB* iff:

1. $\forall (x, +, \#) \in m$, $x \in R^+$,
2. $\forall (x, -, \#) \in m$, $x \in R^-$.

Example, cont. The move $m = \{(ed, -, \#), (ac, +, \#)\}$ on the system

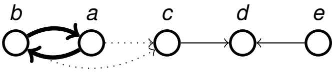

leads to the following modified system, denoted $\Delta(GB, m)$:

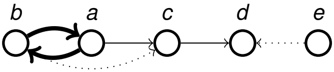

Argumentative Goals

Given:

- A system $GB = \langle A, R, R^+, R^- \rangle$.
- The issue $d \in A$.
- A semantics $S \in \{Adm, Pref, Comp\}$.
- A type of acceptability $X \in \{\exists, \forall\}$ ($\exists$: credulous, $\forall$: sceptical).

We define:

- The **positive goal** $S_X(d)$, which is satisfied iff
  d est is accepted under S and the type of acceptability X.
- The **negative goal** $\neg S_X(d)$, which is satisfied iff
  d is not accepted under S and the type of acceptability X.

Successful Moves and Target Sets

- The move m on GB is **successful for the goal** g iff g is satisfied on $\Delta(GB, m)$.
- The move m on GB is a **target set for the goal** g iff m is $\subseteq$-minimal among the successful moves for g.

For the positive goal S_X (resp. negative goal $\neg S_X$):

- The set of successful moves is denoted $\mathbb{M}^S_X$ (resp. $\mathbb{M}^S_{\neg X}$).
- The set of target sets is denoted $\mathbb{T}^S_X$ (resp. $\mathbb{T}^S_{\neg X}$).

Example, cont.

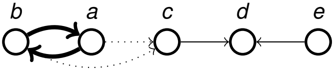

The move $m = \{(ed, -, \#), (ac, +, \#)\}$ is a target set for the goal $Adm_\exists(d)$, but m is not a successful move for the goals $Pref_\forall(d)$ and $Comp_\forall(d)$.

Intuition behind target set computation

- We start from the issue d, and we navigate the argumentation graph, going backwards.
- By applying a set of rewriting rules, we add and remove attacks, thus deciding the status of attacks relevant to the issue.
- The above is done in the form of a term-rewriting procedure.

The Maude Rewriting System

- Maude is a system based on rewriting logic.
- It is also a declarative programming language.
- It can model complex systems (eg. social systems, biological systems, programming languages).

Rewriting Rules

- We define two types of rewriting rules:
  1. **Expansion rules** which replace an atom by some new atoms.
  2. **Simplification rules** which replace two atoms by a single atom.
- An argumentation system (together with a goal) is encoded as a conjunction of atoms.
- The rewriting rules operate on the above encoding.

Rewriting rules - atom $(x, 1, *)$

Atom $(x, 1, *)$ means that attack x must be labelled ‘1’.
The following expansion rules help us find all the ways to achieve this “task”:

1. $(x, 1, *)$ and (y hits x) => $(x, 1, *)$ and $(y, 0, *)$
2. $(x, 1, *)$ and (y hits x) => $(x, 1, *)$ and $(y, -, \#)$ and $(y, 0, \#)$, if the attack y is removable.

Rewriting rules - atom $PRO(d)$

Atom $PRO(d)$ means that d must be accepted.
The following expansion rules help us find all the ways to achieve this “task”:

1. $PRO(d)$ and (y hitsArg d) => $PRO(d)$ and $(y, 0, *)$
2. $PRO(d)$ and (y hitsArg d) => $PRO(d)$ and $(y, -, \#)$ and $(y, 0, \#)$,
   if the attack y is removable.

Simplification Rules

- Fired when two “conflicting” atoms appear in the same conjunct.
- Ensure the termination of the procedure we define next.

Rewriting Rrocedure (RP)

- We start by expanding the atom $PRO(d)$ or $CON(d)$.
- An expansion is followed by one or more simplifications.
- The procedure finishes when no conjunct contains an expandable atom.
- For every computed conjunction, we only take its $(x, +, \#)$ and $(x, -, \#)$ atoms, thus obtaining a set of moves $\mathcal{M}_d^{PRO}$ or $\mathcal{M}_d^{CON}$.

Example, cont.

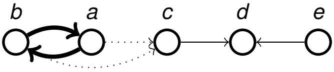

The procedure, starting with $PRO(d)$, computes 3 moves:

$$\mathcal{M}_d^{PRO} = \{\{(cd, -, \#), (ed, -, \#)\},\\ \{(ac, +, \#), (ed, -, \#)\},\\ \{(bc, +, \#), (ed, -, \#)\}\}$$

Properties of RP

- **Termination**:
  RP always terminates after a finite number of rule applications.
- **Determinism**:
  The result of *RP* does not depend on the order of the atoms forming the input conjunct.

Properties of RP - Case of *PRO*(*d*)

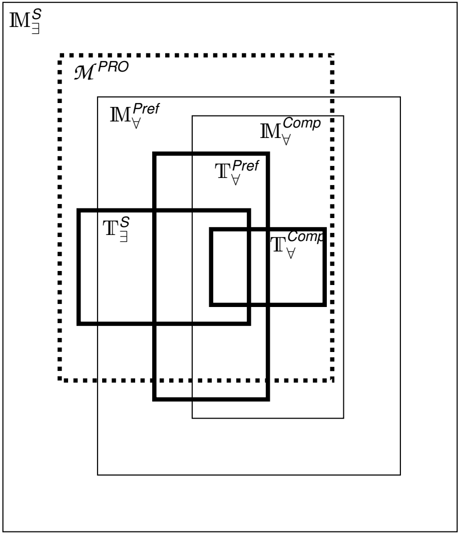

Properties of RP - Case of $CON(d)$

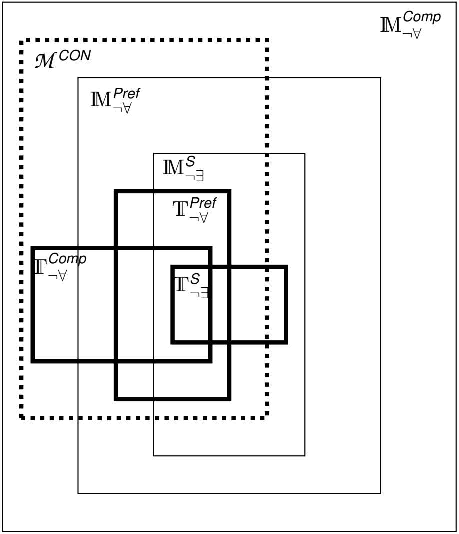

Summing-up

- In this work we have used:
  1. Abstract argumentation.
  2. Extension-based acceptability (both credulous and sceptical).
  3. Attack semantics.
  4. Modifiable argumentation systems where attacks can be added and removed.
- We have provided:
  1. Some properties of the change (mainly of the **minimal change**) needed in order to achieve a goal.
  2. A set of rewriting rules to be used with the Maude system, defining a procedure which computes target sets for different argumentative goals.
  3. Some interesting properties of that procedure.

Extensions of this work

- Further study the properties of minimal change and target sets.
- Tweak the rewriting procedure in order to be able to compute target sets for different types of goals (eg. more semantics, multi-issue goals).
- Define multi-agent debate protocols based on target sets.
- Study the nature of target sets in different systems (eg. systems retrieved from on-line discussions, randomly generated systems).