Adaptive Strategies for Practical Argument-Based Negotiation

Michael Rovatsos
The University of Edinburgh
Iyad Rahwan
The British University in Dubai
Felix Fischer, Gerhard Weiss
Technical University of Munich

ArgMAS Workshop, Utrecht, 26th July 2005

Introduction

- Work on argumentation frameworks is progressing rapidly . . .

Introduction

- Work on argumentation frameworks is progressing rapidly ...

Introduction

- Work on argumentation frameworks is progressing rapidly ...
  ... but how about building agents that can use them?
- We concentrate on two aspects:

Introduction

- Work on argumentation frameworks is progressing rapidly . . .
  . . . but how about building agents that can use them?
- We concentrate on two aspects:
  1. dealing with complex argumentation protocols from an agent's point of view

Introduction

- Work on argumentation frameworks is progressing rapidly ...
  ... but how about building agents that can use them?
- We concentrate on two aspects:
  1. dealing with complex argumentation protocols from an agent's point of view
  2. learning optimal argumentation strategies

Introduction

- Work on argumentation frameworks is progressing rapidly . . .
  . . . but how about building agents that can use them?
- We concentrate on two aspects:
  1. dealing with complex argumentation protocols from an agent's point of view
  2. learning optimal argumentation strategies
- Learning is necessitated by open systems view

Introduction

- Work on argumentation frameworks is progressing rapidly . . .
  . . . but how about building agents that can use them?
- We concentrate on two aspects:
  1. dealing with complex argumentation protocols from an agent's point of view
  2. learning optimal argumentation strategies
- Learning is necessitated by open systems view
- Agents have to find out

Introduction

- Work on argumentation frameworks is progressing rapidly . . .
  . . . but how about building agents that can use them?
- We concentrate on two aspects:
  1. dealing with complex argumentation protocols from an agent's point of view
  2. learning optimal argumentation strategies
- Learning is necessitated by open systems view
- Agents have to find out
  - which argumentation strategies are useful in a given social context

Introduction

- Work on argumentation frameworks is progressing rapidly . . .
  . . . but how about building agents that can use them?
- We concentrate on two aspects:
  1. dealing with complex argumentation protocols from an agent's point of view
  2. learning optimal argumentation strategies
- Learning is necessitated by open systems view
- Agents have to find out
  - which argumentation strategies are useful in a given social context
  - whether and how other agents stick to the provided argumentation mechanism (protocols, constraints)

Strategic Interaction – Agent View

- For an agent, communication protocols impose restrictions on communication while (hopefully) providing information to anticipate future (joint) action in return

Strategic Interaction – Agent View

- For an agent, communication protocols impose restrictions on communication while (hopefully) providing information to anticipate future (joint) action in return
- Given (a set of) interaction protocol(s), i.e. restrictions on surface structure of conversations + constraints regarding their use, how should an agent use them?

Strategic Interaction – Agent View

- For an agent, communication protocols impose restrictions on communication while (hopefully) providing information to anticipate future (joint) action in return
- Given (a set of) interaction protocol(s), i.e. restrictions on surface structure of conversations + constraints regarding their use, how should an agent use them?
- Ultimate goal: influence others' actions while preserving one's own autonomy

Strategic Interaction – Agent View

- For an agent, communication protocols impose restrictions on communication while (hopefully) providing information to anticipate future (joint) action in return
- Given (a set of) interaction protocol(s), i.e. restrictions on surface structure of conversations + constraints regarding their use, how should an agent use them?
- Ultimate goal: influence others' actions while preserving one's own autonomy
- The reactions of others depend on their previous experience with the agent (and vice versa) via expectations

Strategic Interaction – Agent View

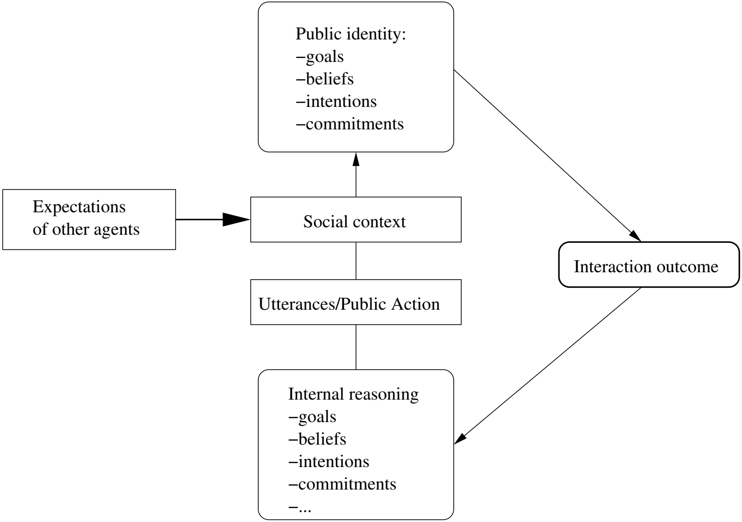

Argumentation-Based Negotiation

- In argumentation-based negotiation (ABN), agents exchange arguments to reach beneficial agreement

Argumentation-Based Negotiation

- In argumentation-based negotiation (ABN), agents exchange arguments to reach beneficial agreement
- Agents exchange arguments concerning propositions about the world (and corresponding proofs)

Argumentation-Based Negotiation

- In argumentation-based negotiation (ABN), agents exchange arguments to reach beneficial agreement
- Agents exchange arguments concerning propositions about the world (and corresponding proofs)
- The "world" may thereby include social commitments, mental states of agents, etc.

Argumentation-Based Negotiation

- In argumentation-based negotiation (ABN), agents exchange arguments to reach beneficial agreement
- Agents exchange arguments concerning propositions about the world (and corresponding proofs)
- The "world" may thereby include social commitments, mental states of agents, etc.
- In contrast to proposal-based negotiation (PBN), ABN uses highly expressive content languages and complex protocols

Argumentation-Based Negotiation

- In argumentation-based negotiation (ABN), agents exchange arguments to reach beneficial agreement
- Agents exchange arguments concerning propositions about the world (and corresponding proofs)
- The "world" may thereby include social commitments, mental states of agents, etc.
- In contrast to proposal-based negotiation (PBN), ABN uses highly expressive content languages and complex protocols
- Allows for exploiting the reasoning capabilities of knowledge-based agents with deductive reasoning capabilities

Research Question

Given a set of argumentation patterns tied to constraints regarding (among other things) the participants' ostensible internal structure, how can we design an agent capable of employing these patterns in order to optimise her own long-term profit?

Outline

Outline

The Interaction Frames Approach

- Goal: learn patterns of agent conversations from experience and apply them strategically in one's own interactions

The Interaction Frames Approach

- Goal: learn patterns of agent conversations from experience and apply them strategically in one's own interactions
- Abstract framework InFFrA (see AAMAS-02), notion of empirical semantics (see AAMAS-03), here only m^2InFFrA as an instance based on probabilistic models of conversations

The Interaction Frames Approach

- Goal: learn patterns of agent conversations from experience and apply them strategically in one's own interactions
- Abstract framework InFFrA (see AAMAS-02), notion of empirical semantics (see AAMAS-03), here only m^2InFFrA as an instance based on probabilistic models of conversations
- In m^2InFFrA, each conversation pattern (interaction frame) consists of

The Interaction Frames Approach

- Goal: learn patterns of agent conversations from experience and apply them strategically in one's own interactions
- Abstract framework InFFrA (see AAMAS-02), notion of empirical semantics (see AAMAS-03), here only m^2InFFrA as an instance based on probabilistic models of conversations
- In m^2InFFrA, each conversation pattern (interaction frame) consists of
  - a sequence of message patterns (speech-act like, augmented with variables)

The Interaction Frames Approach

- Goal: learn patterns of agent conversations from experience and apply them strategically in one's own interactions
- Abstract framework InFFrA (see AAMAS-02), notion of empirical semantics (see AAMAS-03), here only m^2InFFrA as an instance based on probabilistic models of conversations
- In m^2InFFrA, each conversation pattern (interaction frame) consists of
  - a sequence of message patterns (speech-act like, augmented with variables)
  - pairs of logical conditions and variable substitutions

The Interaction Frames Approach

- Goal: learn patterns of agent conversations from experience and apply them strategically in one's own interactions
- Abstract framework InFFrA (see AAMAS-02), notion of empirical semantics (see AAMAS-03), here only m^2InFFrA as an instance based on probabilistic models of conversations
- In m^2InFFrA, each conversation pattern (interaction frame) consists of
  - a sequence of message patterns (speech-act like, augmented with variables)
  - pairs of logical conditions and variable substitutions
  - occurrence counters representing previous enactments

The Interaction Frames Approach

- Goal: learn patterns of agent conversations from experience and apply them strategically in one's own interactions
- Abstract framework InFFrA (see AAMAS-02), notion of empirical semantics (see AAMAS-03), here only m^2InFFrA as an instance based on probabilistic models of conversations
- In m^2InFFrA, each conversation pattern (interaction frame) consists of
  - a sequence of message patterns (speech-act like, augmented with variables)
  - pairs of logical conditions and variable substitutions
  - occurrence counters representing previous enactments
- The architecture combines hierarchical reinforcement learning methods, case-based reasoning and clustering techniques to learn "framing", i.e. strategic use of frames

Example

$$
\begin{aligned}
F = \Big\langle & \Big\langle \xrightarrow{5} \mathtt{request}(A_1, A_2, X) \xrightarrow{3} \mathtt{accept}(A_2, A_1, X) \\
& \quad \xrightarrow{2} \mathtt{confirm}(A_1, A_2, X) \xrightarrow{2} \mathtt{do}(A_2, X)\Big\rangle, \\
& \Big\langle \{\mathit{self}(A_1), \mathit{other}(A_2), \mathit{can}(A_1, \mathtt{do}(A_1, X))\}, \\
& \quad \{\mathit{agent}(A_1), \mathit{agent}(A_2), \mathit{action}(X)\}\Big\rangle, \\
& \Big\langle \xrightarrow{4} \langle [A_1/\mathtt{agent\_1}], [A_2/\mathtt{agent\_2}]\rangle, \\
& \quad \xrightarrow{1} \langle [A_1/\mathtt{agent\_3}], [A_2/\mathtt{agent\_1}], [X/\mathtt{deliver\_goods}]\rangle\Big\rangle\Big\rangle
\end{aligned}
$$

Frame Semantics

- Given a conversation prefix w and a knowledge base KB, a set $\mathcal{F} = \{F_1, \ldots, F_n\}$ of frames induces a continuation probability

$$P(w'|w) = \sum_{F \in \mathcal{F}} P(w'|F, w)P(F|w) = \sum_{F \in \mathcal{F}, ww' = T(F)\vartheta} P(\vartheta|F, w)P(F|w)$$

Frame Semantics

- Given a conversation prefix w and a knowledge base KB, a set $\mathcal{F} = \{F_1, \ldots, F_n\}$ of frames induces a continuation probability

$$P(w'|w) = \sum_{F \in \mathcal{F}} P(w'|F,w)P(F|w) = \sum_{F \in \mathcal{F}, ww' = T(F)\vartheta} P(\vartheta|F,w)P(F|w)$$

- Define probability of ϑ proportional to its *similarity* to F:

$$P(\vartheta|F,w) \propto \sigma(\vartheta, F) = \sum_{i=1}^{|\Theta(F)|} \overbrace{\sigma(T(F)\vartheta, T(F)\Theta(F)[i])}^{\text{similarity}} \overbrace{h_\Theta(F)[i]}^{\text{frequency}} \overbrace{c_i(F, \vartheta, KB)}^{\text{relevance}}$$

The Framing Process

- Frames represent classes of interactions

The Framing Process

- Frames represent classes of interactions
- Proposed hierarchical decision-making approach:
  1. Select the appropriate frame for a given situation (i.e. classify the situation)
  2. Optimise within the selected frame while disregarding other frames

The Framing Process

- Frames represent classes of interactions
- Proposed hierarchical decision-making approach:
  1. Select the appropriate frame for a given situation (i.e. classify the situation)
  2. Optimise within the selected frame while disregarding other frames
- Apply hierarchical reinforcement learning methods to learn usefulness of frames in a given communication situation
  - Start with an initial set of pre-defined frames ("social rules")
  - Adapt frame models according to observed behaviour (or oneself and of others)

The Framing Process

- Frames represent classes of interactions
- Proposed hierarchical decision-making approach:
  1. Select the appropriate frame for a given situation (i.e. classify the situation)
  2. Optimise within the selected frame while disregarding other frames
- Apply hierarchical reinforcement learning methods to learn usefulness of frames in a given communication situation
  - Start with an initial set of pre-defined frames ("social rules")
  - Adapt frame models according to observed behaviour (or oneself and of others)
- Important: Architecture allows deviation from existing frames on all sides

Outline

Outline

Proposal-Based Negotiation Frames

$$F_1 = \Big\langle \langle \xrightarrow{0} \texttt{request}(A, B, X) \xrightarrow{0} \texttt{accept}(B, A, X) \xrightarrow{0} \texttt{confirm}(A, B, X) \xrightarrow{0} \texttt{do}(B, X) \rangle,$$
$$\langle \mathit{can}(B, X)@3, \mathit{effects}(X)@4\} \rangle$$
$$\langle \xrightarrow{0} \langle\rangle \rangle \Big\rangle$$

$$F_2 = \Big\langle \langle \xrightarrow{0} \texttt{request}(A, B, X) \xrightarrow{0} \texttt{propose}(B, A, Y) \xrightarrow{0} \texttt{accept}(A, B, Y) \xrightarrow{0} \texttt{do}(B, Y) \rangle,$$
$$\langle \{\mathit{can}(B, Y)@3, \mathit{effects}(Y)@4\} \rangle$$
$$\langle \xrightarrow{0} \langle\rangle \rangle \Big\rangle$$

$$F_3 = \Big\langle \langle \xrightarrow{0} \texttt{request}(A, B, X) \xrightarrow{0} \texttt{propose}-\texttt{also}(B, A, Y) \xrightarrow{0} \texttt{accept}(A, B, Y)$$
$$\xrightarrow{0} \texttt{do}(B, X) \xrightarrow{0} \texttt{do}(A, Y) \rangle,$$
$$\langle \{\mathit{can}(B, X)@3, \mathit{effects}(X)@4, \mathit{can}(A, Y)@4, \mathit{effects}(Y)@5\} \rangle$$
$$\langle \xrightarrow{0} \langle\rangle \rangle \Big\rangle$$

Interest-Based Negotiation (IBN)

- A specific framework for ABN

Interest-Based Negotiation (IBN)

- A specific framework for ABN
- As opposed to PBN, IBN allows agents to

Interest-Based Negotiation (IBN)

- A specific framework for ABN
- As opposed to PBN, IBN allows agents to
  - obtain information about others' beliefs and goals

Interest-Based Negotiation (IBN)

- A specific framework for ABN
- As opposed to PBN, IBN allows agents to
  - obtain information about others' beliefs and goals
  - point at others' misconceptions

Interest-Based Negotiation (IBN)

- A specific framework for ABN
- As opposed to PBN, IBN allows agents to
  - obtain information about others' beliefs and goals
  - point at others' misconceptions
  - identify/suggest alternatives

Interest-Based Negotiation (IBN)

- A specific framework for ABN
- As opposed to PBN, IBN allows agents to
  - obtain information about others' beliefs and goals
  - point at others' misconceptions
  - identify/suggest alternatives
- Approach due to Rahwan et al.

Interest-Based Negotiation (IBN)

- A specific framework for ABN
- As opposed to PBN, IBN allows agents to
  - obtain information about others' beliefs and goals
  - point at others' misconceptions
  - identify/suggest alternatives
- Approach due to Rahwan et al.
- Our goal: not performance improvement, but coping with more complex communication "regime"

Interest-Based Negotiation (IBN)

- A specific framework for ABN
- As opposed to PBN, IBN allows agents to
  - obtain information about others' beliefs and goals
  - point at others' misconceptions
  - identify/suggest alternatives
- Approach due to Rahwan et al.
- Our goal: not performance improvement, but coping with more complex communication "regime"
- In experiments, to complicate things further we disallow "breaking" frames

IBN – Dialogue Model

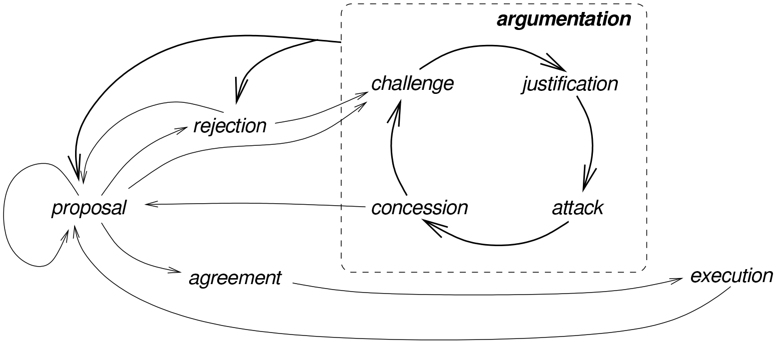

IBN Frames – Example

$$F_{AGM} = \Big\langle \big\langle \xrightarrow{0} \texttt{request}(A, B, X) \xrightarrow{0} \texttt{ask-reason}(B, A, \texttt{request}(X)) \xrightarrow{0} \texttt{inform-goal}(A, B, G) \xrightarrow{0} \texttt{attack-goal}(B, A, \mathit{alternative\text{-}action}(Y)) \xrightarrow{0} \texttt{concede}(A, B, Y) \xrightarrow{0} \texttt{do}(B, Y) \big\rangle, \big\langle \{ can(B, X), goal(A, G), achieves(X, G), achieves(Y, G), X \neq Y, can(B, Y)@5, effects(Y)@6 \} \big\rangle, \big\langle \xrightarrow{0} \langle\rangle \big\rangle \Big\rangle$$

Engineering IBN Frames

- A total of 11 frames, domain-independent

Engineering IBN Frames

- A total of 11 frames, domain-independent
- Interaction frames allow for instant modification and empirical evaluation of argumentation mechanism

Engineering IBN Frames

- A total of 11 frames, domain-independent
- Interaction frames allow for instant modification and empirical evaluation of argumentation mechanism
- Single-shot vs. iterative case (via planning-like concatenation of frames)

Engineering IBN Frames

- A total of 11 frames, domain-independent
- Interaction frames allow for instant modification and empirical evaluation of argumentation mechanism
- Single-shot vs. iterative case (via planning-like concatenation of frames)
- Trade-off between generality and specificity of frames

Engineering IBN Frames

- A total of 11 frames, domain-independent
- Interaction frames allow for instant modification and empirical evaluation of argumentation mechanism
- Single-shot vs. iterative case (via planning-like concatenation of frames)
- Trade-off between generality and specificity of frames
  - having only a few very general frames increases search space at level of utterance generation (many substitutions)

Engineering IBN Frames

- A total of 11 frames, domain-independent
- Interaction frames allow for instant modification and empirical evaluation of argumentation mechanism
- Single-shot vs. iterative case (via planning-like concatenation of frames)
- Trade-off between generality and specificity of frames
  - having only a few very general frames increases search space at level of utterance generation (many substitutions)
  - having many specific ones is not elegant and space-consuming

Outline

Outline

Link Exchange Negotiations

- Imagine agents representing Web sites are able to conduct inference about the content of other pages (e.g. using Semantic Web methods)

Link Exchange Negotiations

- Imagine agents representing Web sites are able to conduct inference about the content of other pages (e.g. using Semantic Web methods)
- Automated inspection of other sites + Knowledge about own preferences (i.e. those of one's owner) = Assessment of own stance of opinions expressed in other sites

Link Exchange Negotiations

- Imagine agents representing Web sites are able to conduct inference about the content of other pages (e.g. using Semantic Web methods)
- Automated inspection of other sites + Knowledge about own preferences (i.e. those of one's owner) = Assessment of own stance of opinions expressed in other sites
- Goal of each Web site owner (and his agent): Maximal dissemination of one's own opinion

Link Exchange Negotiations

- Imagine agents representing Web sites are able to conduct inference about the content of other pages (e.g. using Semantic Web methods)
- Automated inspection of other sites + Knowledge about own preferences (i.e. those of one's owner) = Assessment of own stance of opinions expressed in other sites
- Goal of each Web site owner (and his agent): Maximal dissemination of one's own opinion
- This can be achieved by:

Link Exchange Negotiations

- Imagine agents representing Web sites are able to conduct inference about the content of other pages (e.g. using Semantic Web methods)
- Automated inspection of other sites + Knowledge about own preferences (i.e. those of one's owner) = Assessment of own stance of opinions expressed in other sites
- Goal of each Web site owner (and his agent): Maximal dissemination of one's own opinion
- This can be achieved by:
  - Maximising the popularity of one's own site

Link Exchange Negotiations

- Imagine agents representing Web sites are able to conduct inference about the content of other pages (e.g. using Semantic Web methods)
- Automated inspection of other sites + Knowledge about own preferences (i.e. those of one's owner) = Assessment of own stance of opinions expressed in other sites
- Goal of each Web site owner (and his agent): Maximal dissemination of one's own opinion
- This can be achieved by:
  - Maximising the popularity of one's own site
  - Increasing the popularity of sites that express similar opinions

Link Exchange Negotiations

- Imagine agents representing Web sites are able to conduct inference about the content of other pages (e.g. using Semantic Web methods)
- Automated inspection of other sites + Knowledge about own preferences (i.e. those of one's owner) = Assessment of own stance of opinions expressed in other sites
- Goal of each Web site owner (and his agent): Maximal dissemination of one's own opinion
- This can be achieved by:
  - Maximising the popularity of one's own site
  - Increasing the popularity of sites that express similar opinions
  - Decrease the popularity of sites with unfavourable opinions

Link Exchange Negotiations

- Traffic provides a measure for popularity, and is affected by links between sites

Link Exchange Negotiations

- Traffic provides a measure for popularity, and is affected by links between sites
- Links are weighted with numerical “ratings” expressing opinion source site has of target site

Link Exchange Negotiations

- Traffic provides a measure for popularity, and is affected by links between sites
- Links are weighted with numerical "ratings" expressing opinion source site has of target site
  - In a more advanced system, these would correspond to comments such as *"Click here for my favourite site on topic X"*

Link Exchange Negotiations

- Traffic provides a measure for popularity, and is affected by links between sites
- Links are weighted with numerical "ratings" expressing opinion source site has of target site
  - In a more advanced system, these would correspond to comments such as *"Click here for my favourite site on topic X"*
- Of course, the displayed ratings (actual link weights) can differ from the (private) actual ratings

Link Exchange Negotiations

- Traffic provides a measure for popularity, and is affected by links between sites
- Links are weighted with numerical "ratings" expressing opinion source site has of target site
  - In a more advanced system, these would correspond to comments such as *"Click here for my favourite site on topic X"*
- Of course, the displayed ratings (actual link weights) can differ from the (private) actual ratings
- Agent goal: maximise opinion dissemination (in terms of some utility measure) through **negotiation** with other agents about link exchange

Link Exchange Negotiations

- Traffic provides a measure for popularity, and is affected by links between sites
- Links are weighted with numerical "ratings" expressing opinion source site has of target site
  - In a more advanced system, these would correspond to comments such as *"Click here for my favourite site on topic X"*
- Of course, the displayed ratings (actual link weights) can differ from the (private) actual ratings
- Agent goal: maximise opinion dissemination (in terms of some utility measure) through **negotiation** with other agents about link exchange
- System goal: increase linkage transparency on the WWW

The LIESON System

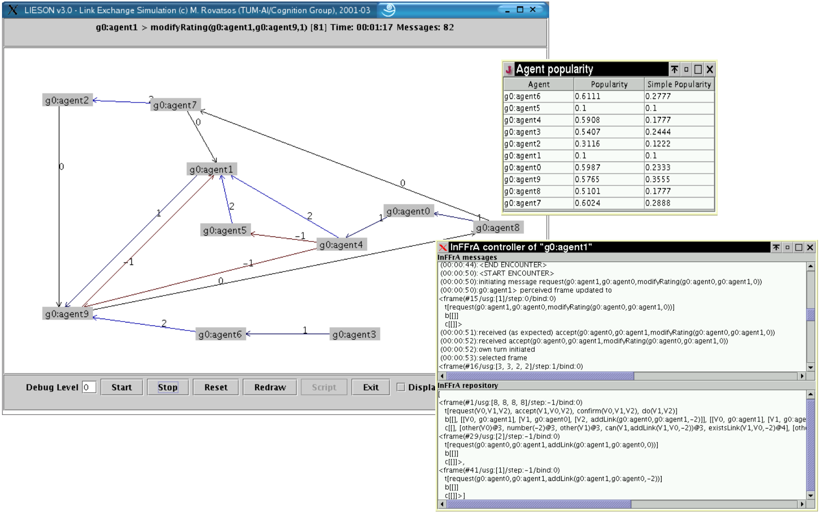

IBN – Goal graphs

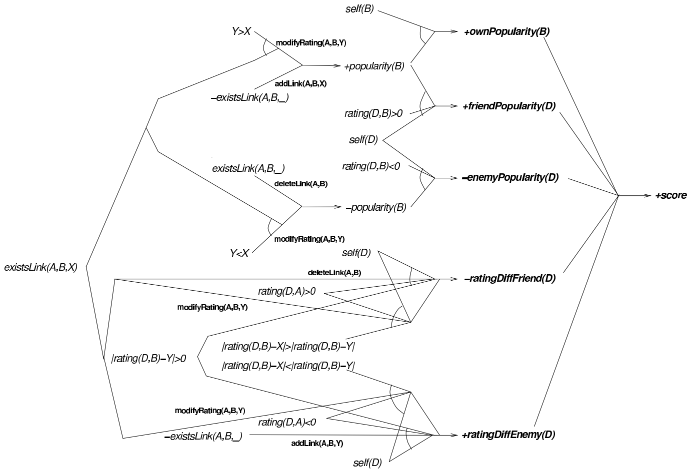

IBN – Goal graph (detail)

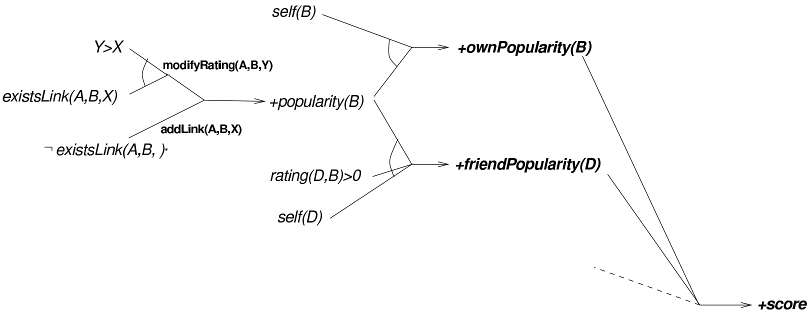

Outline

Outline

Proposal-Based Negotiation

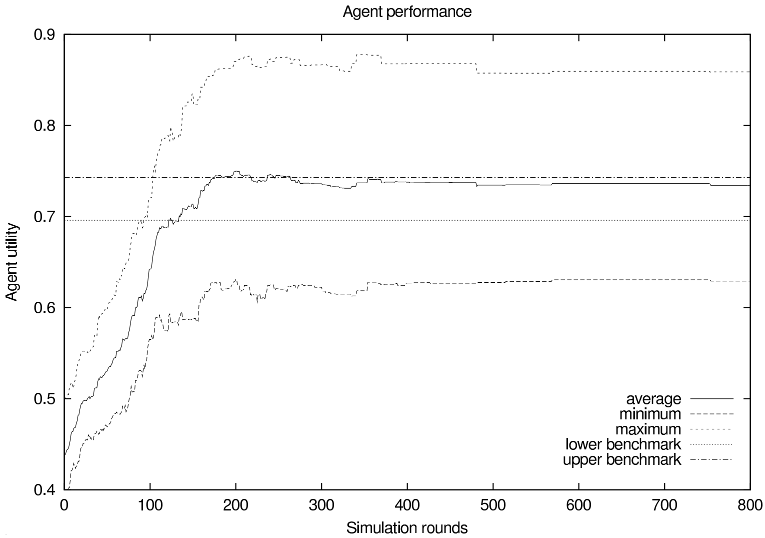

Interest-Based Negotiation

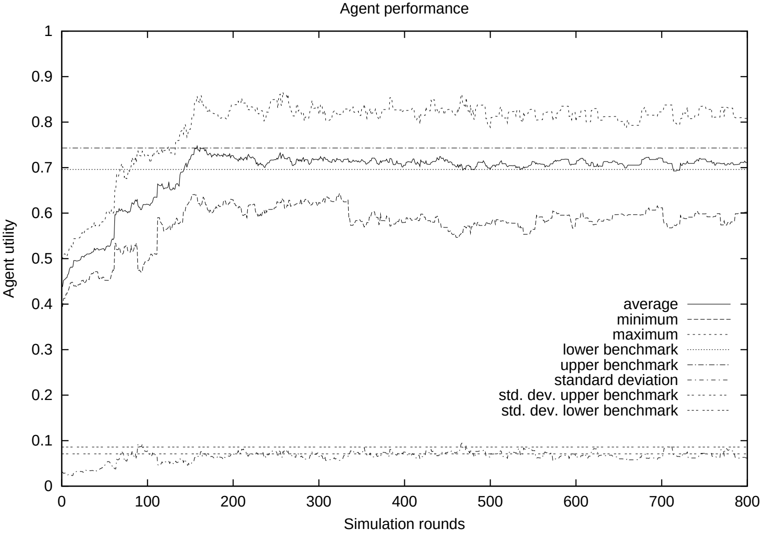

Outline

Outline

Conclusions

- An exercise in the development of adaptive strategies for argument-based negotiation

Conclusions

- An exercise in the development of adaptive strategies for argument-based negotiation
- First attempt to apply learning to complex and expressive argumentation protocols

Conclusions

- An exercise in the development of adaptive strategies for argument-based negotiation
- First attempt to apply learning to complex and expressive argumentation protocols
- Approach computationally tractable (for simple subset of IBN theory), focus on realism

Conclusions

- An exercise in the development of adaptive strategies for argument-based negotiation
- First attempt to apply learning to complex and expressive argumentation protocols
- Approach computationally tractable (for simple subset of IBN theory), focus on realism
- Combination of argumentation frameworks with practical agent architectures crucial for wider acceptance of ABN

Conclusions

- An exercise in the development of adaptive strategies for argument-based negotiation
- First attempt to apply learning to complex and expressive argumentation protocols
- Approach computationally tractable (for simple subset of IBN theory), focus on realism
- Combination of argumentation frameworks with practical agent architectures crucial for wider acceptance of ABN
- Multiagent learning perspective: our approach avoids opponent modelling (which is hardly tractable in large-scale, open multiagent societies)

Outlook

- Optimisation still heuristic, no convergence results

Outlook

- Optimisation still heuristic, no convergence results
- Currently attempting to develop game-theoretic model for argumentation protocols with propositional content to derive provably optimal strategies

Outlook

- Optimisation still heuristic, no convergence results
- Currently attempting to develop game-theoretic model for argumentation protocols with propositional content to derive provably optimal strategies
- Particularly interesting: wise choice of "logical commitments", otherwise future (potentially optimal) statements might be ruled out

Outlook

- Optimisation still heuristic, no convergence results
- Currently attempting to develop game-theoretic model for argumentation protocols with propositional content to derive provably optimal strategies
- Particularly interesting: wise choice of "logical commitments", otherwise future (potentially optimal) statements might be ruled out
- More principled comparison between ABN and PBN necessary, does it pay to introduce added complexity of argumentation?

Outlook

- Optimisation still heuristic, no convergence results
- Currently attempting to develop game-theoretic model for argumentation protocols with propositional content to derive provably optimal strategies
- Particularly interesting: wise choice of "logical commitments", otherwise future (potentially optimal) statements might be ruled out
- More principled comparison between ABN and PBN necessary, does it pay to introduce added complexity of argumentation?
- Long-term goal: mechanism design for argumentation (?)

The End

Thank you for your attention!

Without Frame Learning

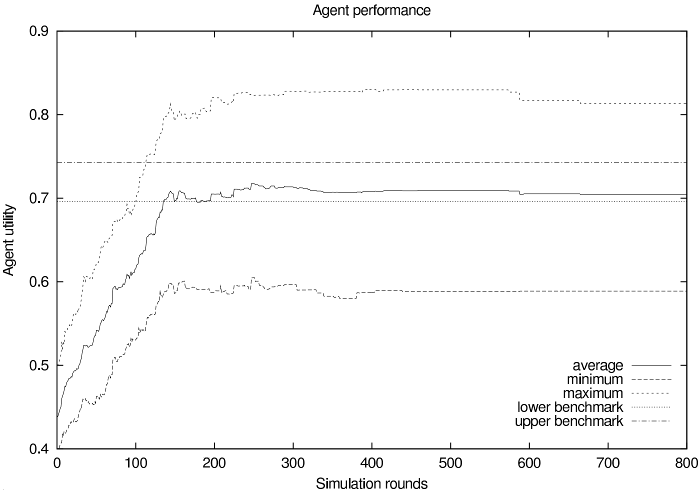

With Frame Learning