Modal Logics of Negotiation and Preference

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

[joint work with Eric Pacuit]

Talk Outline

The motivation behind this line of work is to explore the use of logic as a means of specifying and reasoning about problems in negotiation over the allocation of indivisible resources.

- Some general remarks about *social software*
- Introduction to the problem domain: *negotiation* over resources
- Definition of *our logic* (PDL-style), examples, decidability result
- Modelling *convergence* to socially optimal allocations in our logic
- Discussion of other *points of contact* between dynamic modal logic and the study of negotiation processes

U. Endriss and E. Pacuit. Modal Logics of Negotiation and Preference. Proc. 10th European Conference on Logics in Artificial Intelligence (JELIA-2006).

Social Software

Social software is a research programme that applies tools from *logic* and *computer science* to the study of *social procedures*.

Examples for such social procedures include:

- voting protocols (majority rule, approval voting, . . .)
- fair division algorithms (e.g. "cake-cutting algorithms")

Just as computer programs have properties that can be analysed by means of appropriate logics of programs, social procedures can also be specified and analysed using appropriate logical tools.

Negotiation is yet another example for a social procedure . . .

R. Parikh. Social Software. *Synthese* 132:187–211, 2002.

Negotiation over Indivisible Resources

A finite set of *agents* $\mathcal{A}$ negotiate over the allocation of a finite set of indivisible *resources* $\mathcal{R}$. Some notation:

- An *allocation* is a total function $A : \mathcal{R} \to \mathcal{A}$ specifying *what* is owned by *whom* (set of all allocations: $\mathcal{A}^{\mathcal{R}}$).
- An *atomic deal* $(i \leftarrow r)$ says that resource r is given to agent i ($\leadsto R_{i \leftarrow r} \subseteq \mathcal{A}^{\mathcal{R}} \times \mathcal{A}^{\mathcal{R}}$). A *deal* is a sequence of atomic deals.
- Each agent i has a reflexive and transitive *preference* relation R_i over alternative bundles, inducing also a relation over allocations.

Negotiation is driven by *individual* interests: agents may agree on *any* deals benefitting themselves. As system designers, we are interested in the evolution of allocations and social welfare at the *global* level.

Possible Definitions of Social Welfare

Social choice theorists have come up with a wide range of concepts for assessing the social welfare of an allocation:

- *Utilitarian collective utility:* try to find an allocation maximising the sum of individual utilities (so would need utility functions)
- *Egalitarian collective utility:* find an allocation maximising the utility of the agent that is worst off (*dito*)
- *Envy-freeness:* find an allocation where no agent would rather own the bundle allocated to one of the other agents
- *Pareto optimality:* find an allocation such that no other allocation is better for some agents without being worse for any of the others

For logical modelling, ordinal preferences seem more manageable than cardinal utility functions. For now concentrate on Pareto optimality, but would be nice to tackle the other concepts in the future as well.

Individual Rationality

We assume that agents are *rational* in the sense of accepting only deals that will make them better off (not necessarily strictly).

Example: Agents $\{1, 2, 3\}$ and resources $\{a, b, c\}$. Suppose all three strongly dislike the empty bundle or having more than one item; and:

$$\begin{array}{lcr} \text{(worst)} & \{c\}\, R_1\, \{b\}\, R_1\, \{a\} & \text{(best)} \\ & \{a\}\, R_2\, \{c\}\, R_2\, \{b\} & \\ & \{b\}\, R_3\, \{a\}\, R_3\, \{c\} & \end{array}$$

Initially, they each get their second-best resource (not Pareto optimal).

But *no rational bilateral deal* exists, so imposing that sort of restriction we *cannot reach a Pareto optimal allocation!*

▶ So, when *can* we hope to get to a Pareto optimal allocation?

The Logic $\mathcal{L}_{\langle \mathcal{A},\mathcal{R}\rangle}$: Syntax

Short version: language of PDL with all extras over atomic relations $R_{i \leftarrow r}$ and R_i + special propositions H_{ir} (similar to nominals)

Long version: The set of *relation terms* is the set of all terms we can build from atomic deal relations $R_{i \leftarrow r}$ and preference relations R_i, using these operations:

$$R \; ::= \; R_{i \leftarrow r} \mid R_i \mid R \cup R \mid R \cap R \mid R^{-1} \mid \overline{R} \mid R \circ R \mid R^*$$

Then the set of well-formed *formulas* is constructed as follows:

$$\varphi \; ::= \; p \mid \neg\varphi \mid \varphi \vee \varphi \mid \langle R \rangle \varphi$$

Here p stands for atoms, which include H_{ir} (read: "i holds r") for $i \in \mathcal{A}, r \in \mathcal{R}$. Additional connectives as usual, e.g. $[R]\varphi = \neg\langle R \rangle \neg\varphi$.

Note: For each choice for $\mathcal{A}$ and $\mathcal{R}$ we get a different logic $\mathcal{L}_{\langle \mathcal{A},\mathcal{R}\rangle}$.

The Logic $\mathcal{L}_{\langle \mathcal{A}, \mathcal{R} \rangle}$: Semantics

A *frame* $\mathcal{F} = (\mathcal{A}, \mathcal{R}, \{R_i\}_{i \in \mathcal{A}})$ consists of a set of agents $\mathcal{A}$, a set of resources $\mathcal{R}$, and a preference relation R_i (over allocations) for each agent i. Note that the deal relations $R_{i \leftarrow r}$ are implicit.

A *model* $\mathcal{M} = (\mathcal{F}, V)$ consists of such a frame $\mathcal{F}$ and a *valuation* V mapping atoms to subsets of $\mathcal{A}^\mathcal{R}$ s.t. $V(H_{ir}) = \{A \in \mathcal{A}^\mathcal{R} \mid A(r) = i\}$.

Truth of a formula φ at a world w (an allocation) in a model $\mathcal{M}$:

(1) $\mathcal{M}, w \models p$ iff $w \in V(p)$ for atomic propositions p;

(2) $\mathcal{M}, w \models \neg\varphi$ iff not $\mathcal{M}, w \models \varphi$;

(3) $\mathcal{M}, w \models \varphi \vee \psi$ iff $\mathcal{M}, w \models \varphi$ or $\mathcal{M}, w \models \psi$;

(4) $\mathcal{M}, w \models \langle R \rangle\varphi$ iff there is a $v \in \mathcal{A}^\mathcal{R}$ s.t. wRv and $\mathcal{M}, v \models \varphi$.

Examples: φ holds in some allocation preferred by agent i: $\langle R_i \rangle\varphi$;
ψ holds if we give r_1 or r_2 to agent i: $[R_{i \leftarrow r_1} \cup R_{i \leftarrow r_2}]\psi$

Decidability

PDL with complements is undecidable. Nevertheless:

Proposition 1 (Decidability) *The logic $\mathcal{L}_{\langle \mathcal{A},\mathcal{R} \rangle}$ is decidable.*

<u>Proof:</u> Fixing $\mathcal{A}$ and $\mathcal{R}$ means fixing the set of worlds $\mathcal{A}^{\mathcal{R}}$. Hence, the number of frames is bounded. Considering only the propositional letters occurring in a given formula φ, also the number of relevant models is bounded. Model checking is obviously decidable, so we can "simply" check whether φ is true in all possible models. ✓

Examples

A formula completely specifying the bundle $X \subseteq \mathcal{R}$ held by agent i:

$$\mathrm{BUN}_i^X \quad = \quad \bigwedge_{r \in X} H_{ir} \wedge \bigwedge_{r \in \mathcal{R} \setminus X} \neg H_{ir}$$

Given a partitioning $\langle X_1, \ldots, X_n \rangle$ of $\mathcal{R}$, we can now completely describe an allocation:

$$\mathrm{ALLOC}_{\langle X_1, \ldots, X_n \rangle} \quad = \quad \bigwedge_{i=1}^{n} \mathrm{BUN}_i^{X_i}$$

Each of these formulas works like a nominal.

Pareto Efficiency

An allocation is Pareto optimal *iff* there is no other allocation that is better for some agents without being worse for any of the others.

We can define a relation of *Pareto dominance:*

$$\text{PAR} \quad = \quad \bigcap_{i \in \mathcal{A}} R_i \cap \bigcup_{i \in \mathcal{A}} (R_i \cap \overline{R_i^{-1}})$$

Now we can define a formula characterising *Pareto optimal* allocations:

$$\text{PAR-OPT} \quad = \quad [\text{PAR}]\bot$$

The paper also introduces a *Logic of Pareto Efficiency* for reasoning just about individual and aggregated preferences (but not deals).

Classes of Deals

We can define different classes of deals, conforming to either *structural* or *rationality* constraints. Examples:

- Class of *atomic deals:* $\textsc{atomic} = \bigcup_{i \in \mathcal{A}, r \in \mathcal{R}} R_{i \leftarrow r}$
- Class of *all deals:* $\textsc{all} = \textsc{atomic}^*$
- Class of *rational* deals: nobody suffers, at least one agent gains.

$$\textsc{rat} = \bigcap_{i \in \mathcal{A}} R_i \cap \bigcup_{i \in \mathcal{A}} (R_i \cap \overline{R_i^{-1}})$$

Note that this coincides with Pareto dominance: $\textsc{rat} = \textsc{par}$.

Convergence

A known *convergence result* states that any sequence of rational deals will eventually result in a Pareto optimal allocation, provided deals are not subject to any structural restrictions.

This amounts to saying that the following formula is *valid:*

$$[(\textsc{all} \cap \textsc{rat})^*]\langle(\textsc{all} \cap \textsc{rat})^*\rangle\textsc{par-opt}$$

Also, *no formula* of the following form is valid for $D \subset \textsc{all} \cap \textsc{rat}$:

$$[D^*]\langle D^*\rangle\textsc{par-opt}$$

This means that imposing any structural restrictions whatsoever on the negotiation protocol will make us lose the convergence property.

U. Endriss, N. Maudet, F. Sadri and F. Toni. Negotiating Socially Optimal Allocations of Resources. *Journal of Artificial Intelligence Research* 25:315–348, 2006.

Reachability Properties and Model Checking

Some work on negotiation has investigated the computational complexity of checking whether a given class of deals will guarantee convergence to a top allocation for any given initial allocation.

This corresponds to a *model checking* problem:

- given a model $\mathcal{M}$ encoding a *negotiation scenario*,
- given a relation D encoding a *class of deals*,
- given a formula $\textsc{opt}$ encoding a notion of social optimality,
- check whether $\mathcal{M} \models [D^*]\langle D^* \rangle \textsc{opt}$

What can we get out of this sort of correspondences?

- Use model checking algorithms for negotiation support
- Compare (and obtain) complexity results

P.E. Dunne, M. Wooldridge, and M. Laurence. The Complexity of Contract Negotiation. *Artificial Intelligence* 164(1–2):23–46, 2005

Convergence and Correspondence Theory

Another line of research in negotiation has been to identify conditions on agent preferences that would guarantee convergence for a given class of deals.

This can be mapped to a question in *correspondence theory*:

- given a relation D encoding a *class of deals*,
- given a formula $\textsc{opt}$ encoding a notion of social optimality,
- identify a class of *frames* such that $\mathcal{F} \models [D^*]\langle D^*\rangle\textsc{opt}$

U. Endriss, N. Maudet, F. Sadri and F. Toni. Negotiating Socially Optimal Allocations of Resources. *Journal of Artificial Intelligence Research* 25:315–348, 2006.

Conclusions

- Proposal for a PDL-style logic that allows us to represent
  - individual and aggregated preferences of agents;
  - resource allocations and deals that alter allocations
- Parallels between questions investigated in negotiation on the one hand, and reasoning tasks in modal logic on the other:
  - theorem proving, model checking, correspondence theory, . . .
- Example for a general trend towards applying tools from logic and computer science to the study of problems originating in the socio-economic sciences: social software, algorithmic game theory, computational social choice, . . .