Multi-Agent Systems

Sept. 2000, Bahia Blanca

University Nacional del Sur

- **Last two weeks in September**.
- **Tentative Dates:** Tuesday, Sept. 19th, Thursday, Sept. 21st, Friday, Sept. 22nd, Tuesday, Sept. 26th, Thursday, Sept. 28th, Friday, Sept. 29th.
- **Time:** From 4–6 pm, unless otherwise indicated.
- Lecture Course is on theoretical issues, emphasis on mathematical-logical foundations.

Overview

Chapter 4. Distributed Decision Making

4.1 Evaluation Criteria

4.2 Voting

4.3 Auctions

4.4 Bargaining

4.5 General Market Criteria

4 Distributed Decision Making

Two and a half lectures: first lecture up to 4.3, second lecture 4.3 – 4.5, half lecture from 4.5 to the end.

Classical DAI: System Designer fixes an Interaction-Protocol which is uniform for all agents. The designer also fixes a strategy for each agent.

What is a the outcome, assuming that the protocol is followed and the agents follow the strategies?

MAI: Interaction-Protocol is given. Each agent determines its own strategy (maximizing its own good, via a utility function, without looking at the global task).

What is the outcome, given a protocol that guarantees that each agent's desired local strategy is the best one (and is therefore chosen by the agent)?

4.1 General Evaluation Criteria

We need to compare negotiation protocols. Each such protocol leads to a solution. So we determine how good these solutions are.

Social Welfare: Sum of all utilities

Pareto Efficiency: A solution $\boldsymbol{x}$ is Pareto-optimal (also called efficient), if

> there is no solution $\boldsymbol{x'}$ with: (1) $\exists$agent ag : $\mathbf{ut}_{ag}(\boldsymbol{x'}) > \mathbf{ut}_{ag}(\boldsymbol{x})$
> (2) $\forall$agents ag' : $\mathbf{ut}_{ag'}(\boldsymbol{x'}) \geq \mathbf{ut}_{ag'}(\boldsymbol{x})$.

Individual rational: if the payoff is higher than not participating at all.

Stability:

Case 1: Strategy of an agent depends on the others.

The profile $S^*_A = \langle S^*_1, S^*_2, \ldots, S^*_{|A|} \rangle$ is called a Nash-equilibrium, iff

$\forall i : S^*_i$ is the best strategy for agent i if all the others choose $\langle S^*_1, S^*_2, \ldots, S^*_{i-1}, S^*_{i+1}, \ldots, S^*_{|A|} \rangle$.

Case 2: Strategy of an agent does not depend on the others.
Such strategies are called dominant.

Prisoner's Dilemma

		Prisoner 2	
		cooperate	defect
Prisoner 1	cooperate	(3,3)	(0,5)
	defect	(5,0)	(1,1)

- **Social Welfare:** Both cooperate,
- **Pareto-Efficiency:** All are Pareto optimal, except when both defect.
- **Dominant Strategy:** Both defect.
- **Nash Equilibrium:** Both defect.

4.2 Voting

Agents give input to a mechanism and the outcome of it is taken as a solution for the agents.

Motivation: 3 candidates, 3 voters

	1	2	3
w_1	A	B	C
w_2	B	C	A
w_3	C	A	B

Figure 4.1: Nonexistence of desired preference ordering.

Comparing A and B: majority for A. Comparing A and C: majority for C. Comparing B and C: majority for B. **Desired Preference ordering:** $A > B > C > A$????

- Let $\mathcal{A}$ the set of agents, O the set of possible outcomes.
  (O could be equal to $\mathcal{A}$, or a set of laws).
- The **voting** of agent i is described by a binary relation

$$\prec_i \subseteq O \times O,$$

  which we assume to be asymmetric, strict and transitive. We denote by ***Ord*** the set of all such binary relations.
- Often, not all subsets of O are *votable*, only a subset $V \subseteq 2^O \setminus \{\emptyset\}$.
  Each $v \in V$ represents a possible "set of candidates". The voting model then has to select some of the elements of v.
- Each agent votes independently of the others. But we also allow that only a subset is considered. Let therefore be

$$U \subseteq \prod_{i=1}^{|\mathcal{A}|} \mathit{Ord}.$$

- A social choice rule wrt. U is a function

$$\mathbf{f}^* : U \to \mathit{Ord}; (\prec_1, \ldots, \prec_{|\mathbb{A}|}) \mapsto \prec^*$$

For each $V \subseteq 2^O \setminus \{\emptyset\}$ the function $\mathbf{f}^*$ w.r.t. U induces a choice function $\mathbf{C}_{\langle \prec_1, \ldots, \prec_{|\mathbb{A}|} \rangle}$ as follows:

$$\mathbf{C}_{\langle \prec_1, \ldots, \prec_{|\mathbb{A}|} \rangle} =_{def} \begin{cases} V & \longrightarrow & V \\ v & \mapsto & \mathbf{C}_{\langle \prec_1, \ldots, \prec_{|\mathbb{A}|} \rangle}(v) = \max_{\prec^*|_V} v \end{cases}$$

$\max_{\prec^*|_V} v$ is the set of all maximal elements in v according to $\prec^*|_V$.

Each tupel $u = (\prec_1, \ldots, \prec_{|\mathbb{A}|})$ determines the election for all possible $v \in V$.

What are desirable properties for $\mathbf{f}^*$?

Pareto-Efficiency: for all $o, o' \in O$: $(\forall \mathbf{i} \in \mathbf{A} : o \prec_{\mathbf{i}} o')$ implies $o \prec^* o'$.

Independance of Irrelevant Alternatives: for all $o, o' \in O$:

$$(\forall \mathbf{i} \in \mathbf{A} : \ o \prec_{\mathbf{i}} o' \ \text{ iff } \ o \prec'_{\mathbf{i}} o') \Rightarrow (o \prec^* o' \ \text{ iff } \ o \prec'^* o') \, .$$

Note that this implies in particular

$$(\forall \mathbf{i} \in \mathbf{A} : \prec_{\mathbf{i}}|_v = \prec'_{\mathbf{i}}|_v) \ \Rightarrow \ \forall o, o' \in v, \forall v' \in V \text{ s.t. } v \subseteq v' : (o \prec^*|_{v'} o' \ \text{ iff } \ o \prec'^*|_{v'} o')$$

The simple **majority vote** protocol does not satisfy the Independance of irrelevant alternatives.

We consider 7 voters ($\mathcal{A} = \{w_1, w_2, \ldots, w_7\}$) and $O = \{a,b,c,d\}$, $V = \{\{a,b,c,d\},\{a,b,c\}\}$. The columns in the following table represent two different preference orderings of the voters: one is given in black, the second in red.

	$\prec_1$ ($\prec_1$)	$\prec_2$ ($\prec_2$)	$\prec_3$ ($\prec_3$)	$\prec_4$ ($\prec_4$)	$\prec_5$ ($\prec_5$)	$\prec_6$ ($\prec_6$)	$\prec_7$ ($\prec_7$)
a	1 (2)	1 (2)	1 (1)	1 (1)	2 (2)	2 (2)	2 (2)
b	2 (3)	2 (3)	2 (2)	2 (2)	1 (1)	1 (1)	1 (1)
c	3 (4)	3 (4)	3 (3)	3 (3)	3 (3)	3 (3)	3 (3)
d	4 (1)	4 (1)	4 (4)	4 (4)	4 (4)	4 (4)	4 (4)

Let $\prec^*$ be the solution generated by the $\prec_i$ and $\prec^*$ the solution generated by the $\prec_i$. Then we have for $i = 1, \ldots, 7$: $b \prec_i a$ iff $a \prec_i b$, but $b \prec^* a$ and $a \prec^* b$. The latter holds because on the whole set O, for $\prec^*$ a gets selected 4 times and b only 3 times, while for $\prec^*$ a gets selected only 2 times but b gets still selected 3 times. The former holds because we even have $\prec_i|_{\{a,b,c\}} = \prec_i|_{\{a,b,c\}}$.

The introduction of the irrelevant (concerning the relative ordering of a and b) alternative d changes everything: the original majority of a is split and drops below one of the less preferred alternatives (b).

Theorem 4.1 (Arrows Theorem)
If the choice function $\mathbf{f}^*$ *is (1) pareto efficient and (2) independent from irrelevant alternatives, then* ***there always exists a dictator****: for all* $U \subseteq \prod_{\mathfrak{i}=1}^{|\mathfrak{A}|} \mathit{Ord}$

$$\exists \mathfrak{i} \in \mathfrak{A}: \ \forall o, o' \in O: \ o \prec_{\mathfrak{i}} o' \ \text{ iff } \ o \prec^* o'.$$

To be more precise: for all $U \subseteq \prod_{\mathfrak{i}=1}^{|\mathfrak{A}|} \mathit{Ord}$

$$\exists \mathfrak{i} \in \mathfrak{A}: \forall \langle \prec_1, \ldots, \prec_{|\mathfrak{A}|} \rangle \in U: \ \forall o, o' \in O, o \prec_{\mathfrak{i}} o' \ \text{ iff } \ o\, \mathbf{f}^*(\langle \prec_1, \ldots, \prec_{|\mathfrak{A}|} \rangle)\, o'.$$

Ways out:

1. Choice function is not always satisfied.
2. Independence of alternatives is dropped.

The Theorem of Arrow can be even more generalized by weakening the assumption that $\prec^*$ needs to be transitive. In fact, it also holds when using the following definition.

- A social choice rule wrt. U is a function

$$\mathbf{f}^* : U \longrightarrow \mathcal{C}(V); (\prec_1, \ldots, \prec_{|A|}) \mapsto \mathbf{C}_{\langle \prec_1, \ldots, \prec_{|A|} \rangle},$$

  where $\mathbf{C}_{\langle \prec_1, \ldots, \prec_{|A|} \rangle}$ **is any function** from V into 2^O satisfying (1) $\mathbf{C}_{\langle \prec_1, \ldots, \prec_{|A|} \rangle}(v) \neq \emptyset$ and (2) $\mathbf{C}_{\langle \prec_1, \ldots, \prec_{|A|} \rangle}(v) \subseteq v$.
  Such a function simply selects a subset of v: the elected members of the list v.
  No other assumptions about this function are made.

Binary protocol

Pairwise comparison. Not only introduction of irrelevant alternatives but also the ordering may drastically change the outcome.

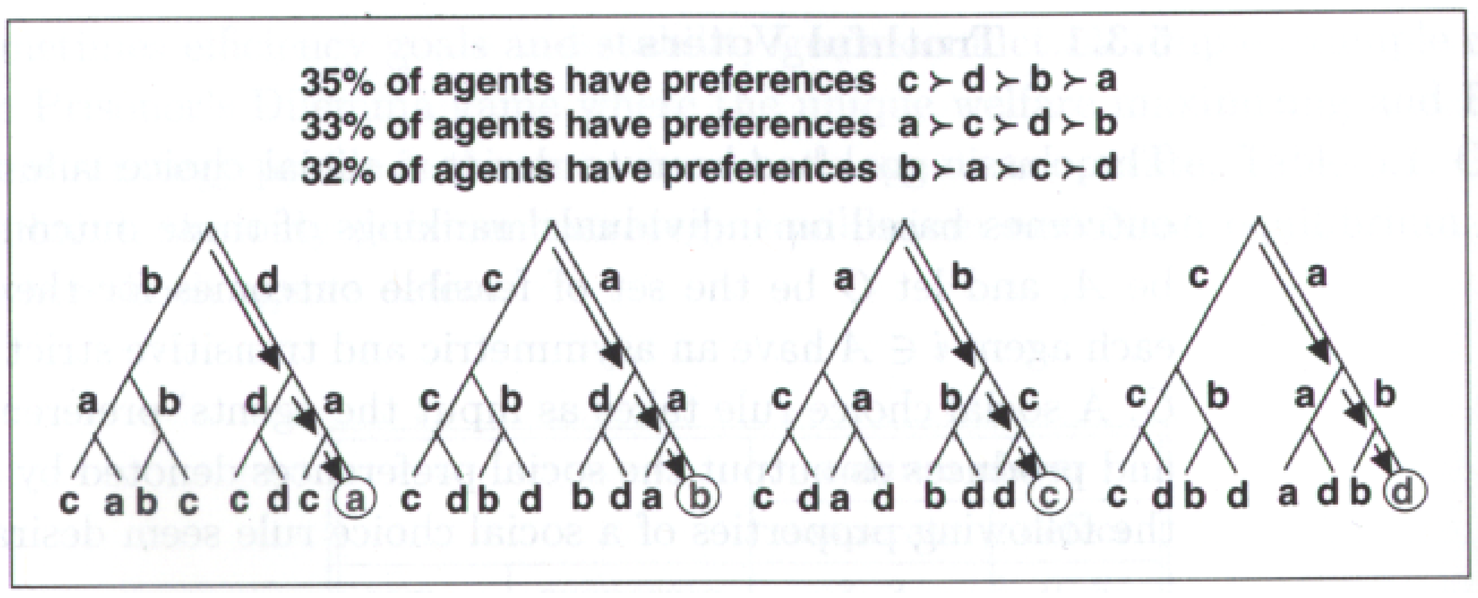

Figure 4.2: Four different orderings and four alternatives in a binary protocol.

Last ordering: d wins, but all agents prefer c over d.

Borda protocol

First gets $|O|$ points, second $|O|-1, \ldots$ Then it is summed up, across voters. The alternative with the highest count wins.

Agent	**Preferences**
1	$a \succ b \succ c \succ d$
2	$b \succ c \succ d \succ a$
3	$c \succ d \succ a \succ b$
4	$a \succ b \succ c \succ d$
5	$b \succ c \succ d \succ a$
6	$c \succ d \succ a \succ b$
7	$a \succ b \succ c \succ d$
Borda count	c wins with 20, b has 19, a has 18, d loses with 13
Borda count with d removed	a wins with 15, b has 14, c loses with 13

Figure 4.3: Winner turns loser and loser turns winner

4.3 Auctions

While voting binds all agents, Auctions are always deals between 2.

Types of auctions:

first-price open cry: (English auction), as usual.

first-price sealed bid: one bids without knowing the other bids.

dutch auction: (descending auction) the seller lowers the price until it is taken.

second-price sealed bid: (Vickrey auction) Highest bidder wins, but the price is the second highest bid!

Three different auction settings:

private value: Value depends only on the bidder (cake).

common value: Value depends only on other bidders (treasury bills).

correlated value: Partly on own's values, partly on others.

What is the best strategy in Vickrey auctions?

Theorem 4.2 (Private-value Vickrey auctions)
The ***dominant strategy*** *of a bidder in a Private-value Vickrey auction* ***is to bid the true valuation****.*

Therefore it is equivalent to english auctions.

Vickrey auctions are used to

- allocate computation resources in operating systems,
- allocate bandwith in computer networks,
- control building heating.

Are first-price auctions better for the auctioneer than second-prize auctions?

Theorem 4.3 (Expected Revenue)
All 4 types of protocols produce the same expected revenue to the auctioneer (assuming (1) private value auctions, (2) values are independently distributed and (3) bidders are risk-neutral).

Why are second price auctions not so popular among humans?

1. Lying auctioneer.
2. When the results are published, subcontractors know the true valuations and what they saved. So they might want to share the profit.

Inefficient Allocation and Lying at Vickrey

Auctioning heterogenous, **interdependent** items.

Example 4.1 (Task Allocation)
Two delivery tasks t_1, t_2. Two agents. $\leadsto$ ***blackboard.***

The global optimal solution is not reached by auctioning independently and truthful bidding.

t_1 goes to agent **2** (for a price of **2**) and t_2 goes to agent **1** (for a price of 1.5).

Even if agent **2** considers (when bidding for t_2) that he already got t_1 (so he bids **cost**$(\{t_1,t_2\}) -$ **cost**$(\{t_1\}) = 2.5 - 1.5 = 1)$ he will get it only with a probability of 0.5.

What about full lookahead ? $\rightsquigarrow$ **blackboard.**

Therefore:

- It pays off for agent **1** to bid more for t_1 (up to 1.5 more than truthful bidding).
- It does not pay off for agent **2**, because agent **2** does not make a profit at t_2 anyway.
- Agent **1** bids 0.5 for t_1 (instead of 2), agent **2** bids 1.5. Therefore agent **1** gets it for 1.5. Agent **1** also gets t_2 for 1.5.

Does it make sense to countersperculate at private value Vickrey auctions?

Vickrey auctions were invented to avoid counterspeculation. But what if the private value for a bidder is uncertain? The bidder might be able to determine it, but he needs to invest c.

Example 4.2 (Incentive to counterspeculate)

*Suppose bidder 1 does not know the (private-) value v_1 of the item to be auctioned. To determine it, he needs to invest **cost**. We also assume that v_1 is uniformly distributed: satisfies $v_1 \in [0,1]$.*

For bidder 2, the private value v_2 of the item is fixed: $0 \leq v_2 < \frac{1}{2}$. So his dominant strategy is to bid v_2.

*Should bidder 1 try to invest **cost** to determine his private value? How does this depend on knowing v_2?*

⤳ **blackboard.**

Answer: Bidder 1 should invest **cost** if and only if

$$v_2 \geq (2\mathbf{cost})^{\frac{1}{2}}.$$

4.4 Bargaining

Axiomatic Bargaining

We assume two agents 1, 2, each with a utility function $\mu_i : E \longrightarrow \mathbb{R}$. If the agents do not agree on a result e the fallback e_{fallback} is taken.

Example 4.3 (Sharing 1 Peso)
How to share 1 Peso?

Agent 1 *offers* ρ $(0 < \rho < 1)$. *Agent* 2 *agrees!*

Such deals are individually rational and each one is in Nash-equilibrium!

Therefore we need axioms!

Axioms on the global solution $\mu^* = \langle \mu_1(e^*), \mu_2(e^*) \rangle$.

Invariance: Absolute values of the utility functions do not matter, only relative values.

Symmetry: Changing the agents does not influence the solution e^*.

Irrelevant Alternatives: If E is made smaller but e^* still remains, then e^* remains the solution.

Pareto: The players can not get a higher utility than $\mu^* = \langle \mu_1(e^*), \mu_2(e^*) \rangle$.

Theorem 4.4 (Unique Solution)
The four axioms above uniquely determine a solution. This solution is given by

$$e^* = \arg\max_e \{(\mu_1(e) - \mu_1(e_{fallback})) \times (\mu_2(e) - \mu_2(e_{fallback}))\}.$$

Strategic Bargaining

No axioms: view it as a game!

Example revisited: Sharing 1 Peso.

Protocol with finitely many steps: The last offerer just offers ε. This should be accepted, so the last offerer gets $1 - \varepsilon$.

This is unsatisfiable. Ways out:

1. Add a discountfactor δ: in round n, only the δ^{n-1}th part of the original value is available.
2. Bargaining costs: bargaining is not for free—fees have to be paid.

Finite Games: Suppose $\delta = 0.9$. Then the outcome depends on # rounds.

Round	1's share	2's share	Total value	Offerer
⋮	⋮	⋮	⋮	⋮
$n-3$	0.819	0.181	0.9^{n-4}	2
$n-2$	0.91	0.09	0.9^{n-3}	1
$n-1$	0.9	0.1	0.9^{n-2}	2
n	1	0	0.9^{n-1}	1

Infinite Games: δ_1 factor for agent **1**, δ_2 factor for agent **2**.

Theorem 4.5 (Unique solution for infinite games)

In a discounted infinite round setting, theres exists a unique Nash equilibrium : Agent 1 gets $\frac{1-\delta_2}{1-\delta_1\delta_2}$. Agent 2 gets the rest. Agent is reached in the first round.

Proof:

Round	1's share	2's share	Offerer
⋮	⋮	⋮	⋮
$t-2$	$1-\delta_2(1-\delta_1\bar{\pi}_1)$		1
$t-1$		$1-\delta_1\bar{\pi}_1$	2
t	$\bar{\pi}_1$		1
⋮	⋮	⋮	⋮

Bargaining Costs

Agent 1 pays c_1, agent 2 pays c_2.

$c_1 = c_2$: Any split is in Nash-equilibrium.

$c_1 < c_2$: Agent 1 gets all.

$c_1 > c_2$: Agent 1 gets c_2, agent 2 gets $1 - c_2$.

4.5 General Equilibrium Mechanisms

A theory for efficiently allocating goods and resources amond agents, based on market prices.

Goods: Given $n > 0$ goods g (coffee, mirror sites, parameters of an airplane design). We assume $g \neq g'$ but within g everything is indistinguishable.

Prices: The market has prices $\mathbf{p} = [p_1, p_2, \ldots, p_n] \in \mathbb{R}^n$: p_i is the price of the good i.

Consumers: Consumer i has $\boldsymbol{\mu_i}(\mathbf{x})$ encoding its preferences over consumption bundles $\mathbf{x_i} = [x_{i1}, \ldots, x_{in}]^t$, where $x_{ig} \in \mathbb{R}^+$ is consumer i's allocation of good g. Each consumer also has an initial endowment $\mathbf{e_i} = [e_{i1}, \ldots, e_{in}]^t \in \mathbb{R}$.

Producers: Use some commodities to produce others: $\mathbf{y_j} = [y_{j1}, \ldots, y_{jn}]^t$, where $y_{jg} \in \mathbb{R}$ is the amount of good g that producer j produces.
$\boldsymbol{Y_j}$ is a set of such vectors $\mathbf{y}$.
Profit of producer j: $\mathbf{p} \times \mathbf{y_j}$, where $\mathbf{y_j} \in \boldsymbol{Y_j}$.

Profits: The profits are divided among the consumers (given predetermined proportions Δ_{ij}): Δ_{ij} is the fraction of producer j that consumer i owns (stocks). Profits are divided according to Δ_{ij}.

Definition 4.1 (General Equilibrium)
$(\mathbf{p}^*, \mathbf{x}^*, \mathbf{y}^*)$ *is in general equilibrium, if the following holds:*

I. *The markets are in equilibrium:*

$$\sum_i \mathbf{x}_i^* = \sum_i \mathbf{e}_i + \sum_j \mathbf{y}_j^*$$

II. *Consumer i maximizes preferences according the prices*

$$\mathbf{x}_i^* = \mathit{arg\ max}_{\{\mathbf{x}_i \in \mathbb{R}_+^n \mid cond_i\}} \mu_i(\mathbf{x}_i)$$

where $cond_i$ *stands for* $\mathbf{p}^* \times \mathbf{x}_i \leq \mathbf{p}^* \times \mathbf{e}_i + \sum_j \Delta_{ij} \mathbf{p}^* \times \mathbf{y}_i$.

III. *Producer j maximizes profit wrt. the market*

$$\mathbf{y}_i^* = \mathit{arg\ max}_{\{\mathbf{y}_j \in Y_j\}} \mathbf{p}^* \times \mathbf{y}_j$$

Theorem 4.6 (Pareto Efficiency)
Each general equilibrium is pareto efficient.

Theorem 4.7 (Coalition Stability)
Each general equilibrium ***with no producers*** *is coalition-stable: no subgroup can incease their utilities by deviating from the equilibrium and building their own market.*

Theorem 4.8 (Existence of an Equilibrium)
Let the sets Y_j *be closed, convex and bounded above. Let* μ_i *be continous, strictly convex and strongly monotone. Assume further that at least one bundle* $\mathbf{x_i}$ *is producible with only positive entries* x_{il}.

Under these assumptions a general equilibrium exists.

4.6 Meaning of the assumptions

Formal definitions: $\leadsto$ **blackboard.**

Convexity of Y_j: Economies of scale in production do not satisfy it.

Continuity of the μ_i: Not satisfied in bandwith allocation for video conferences.

Strictly convex: Not satisfied if preference increases when he gets more of this good (drugs, alcohol, dulce de leche).

In general, there exist more than one equilibrium.

Theorem 4.9 (Uniqueness)
If the society-wide demand for each good is non-decreasing in the prices of the other goods, then a unique equilibrium exists.

Positive example: increasing price of meat forces people to eat potatoes (pasta).

Negative example: increasing price of bread implies that the butter consumption decreases.

References

Arisha, K., F. Ozcan, R. Ross, V. S. Subrahmanian, T. Eiter, and S. Kraus (1999, March/April). IMPACT: A Platform for Collaborating Agents. *IEEE Intelligent Systems 14*, 64–72.

Bratman, M., D. Israel, and M. Pollack (1988). Plans and Resource-Bounded Practical Reasoning. *Computational Intelligence 4*(4), 349–355.

Dix, J., S. Kraus, and V. Subrahmanian (2001). Temporal agent reasoning. *Artificial Intelligence to appear*.

Dix, J., M. Nanni, and V. S. Subrahmanian (2000). Probabilistic agent reasoning. *Transactions of Computational Logic 1*(2).

Dix, J., V. S. Subrahmanian, and G. Pick (2000). Meta Agent Programs. *Journal of Logic Programming 46*(1-2), 1–60.

Eiter, T., V. Subrahmanian, and G. Pick (1999). Heterogeneous Active Agents, I: Semantics. *Artificial Intelligence 108*(1-2), 179–255.

Eiter, T., V. Subrahmanian, and T. J. Rogers (2000). Heterogeneous Active Agents, III: Polynomially Implementable Agents. *Artificial Intelligence 117*(1), 107–167.

Eiter, T. and V. S. Subrahmanian (1999). Heterogeneous Active Agents, II: Algorithms and Complexity. *Artificial Intelligence 108*(1-2), 257–307.

Georgeff, M. and A. Lansky (1987). Reactive Reasoning and Planning. In *Proceedings of the Conference of the American Association of Artificial Intelligence*, Seattle, WA, pp. 677–682.

Rao, A. S. (1995). Decision Procedures for Propositional Linear-Time Belief-Desire-Intention Logics. In M. Wooldridge, J. Müller, and M. Tambe (Eds.), *Intelligent Agents II – Proceedings of the 1995 Workshop on Agent Theories, Architectures and Languages (ATAL-95)*, Volume 890 of *LNAI*, pp. 1–39. Berlin, Germany: Springer-Verlag.

Rao, A. S. and M. Georgeff (1991). Modeling Rational Agents within a BDI-Architecture. In J. F. Allen, R. Fikes, and E. Sandewall (Eds.), *Proceedings of the International Conference on Knowledge Representation and Reasoning*, Cambridge, MA, pp. 473–484. Morgan Kaufmann.

Rao, A. S. and M. Georgeff (1995, June). Formal models and decision procedures for multi-agent systems. Technical Report 61, Australian Artificial Intelligence Institute, Melbourne.

Subrahmanian, V., P. Bonatti, J. Dix, T. Eiter, S. Kraus, F. Özcan, and R. Ross (2000). *Heterogenous Active Agents*. MIT-Press.

Weiss, G. (Ed.) (1999). *Multiagent Systems*. MIT-Press.