Agent-Based Systems

Michael Rovatsos
mrovatso@inf.ed.ac.uk

Lecture 15 – Summary and Concluding Remarks

Lessons learnt

- Course served as an introduction to the area of agents and multiagent systems
- Today we review the central insights of the past lectures

Lessons learnt: introduction, agency, abstract architectures

- Pervasive trends in computing history raise new requirements
- One possible answer: agents and multiagent systems
- Agent notion fuzzy, criticism & abuse
- Transdisciplinary area (inspiration from philosophy, sociology, psychology, economics, etc.)
- Distinctions to AI, distributed systems, economics, objects, expert systems

Lessons learnt: defining agency, abstract architectures

- Situatedness, autonomy, reactivity, proactiveness, social ability
- Rationality = proactiveness + reactivity
- Formal abstract models: runs, transformer functions, behavioural equivalence, perception and action, internal states
- Telling agents what to do: utilities and the MEU principles, optimal agents, predicate task specifications

Lessons learnt: agent architectures

- Deductive reasoning agents: logic-based agents, concurrent MetateM
- Practical reasoning systems & BDI, planning
- Reactive architectures: subsumption architecture
- Hybrid architectures: vertical & horizontal layering (Touring machines, InteRRaP)

Lessons learnt: communication and interaction

- Agent communication languages
- Speech act theory: communication as action
- Plan-based theory of speech act semantics
- The KQML/KIF and FIPA/ACL languages
- Mentalistic & commitment-based semantics, associated problems
- Interaction protocols, the contract-net protocol

Lessons learnt: multiagent interactions

- Utility- and preference-based model of interaction
- Game-theoretic notions: games, strategies, equilibria
- Prisoners' Dilemma, the evolution of cooperation?
- Critique of game-theoretic models

Lessons learnt: coordination methods

- (Generalized) Partial Global Planning
- Joint intentions: commitments and conventions
- Teamwork-based model of CDPS
- Mutual modelling
- Norms and social laws (off-line design and emergent norms)

Lessons learnt: social choice

- Making group decisions given individuals' preferences
- Simple plurality, sequential voting
- Succinct representations, majority graphs
- Borda count and Slater ranking
- Arrow's impossibility theorem
- Strategic manipulation and it complexity

Lessons learnt: coalition formation

- How to organise collaboration and split gain?
- Cooperative game theory – games with enforceable deals
- The Core concept and the Shapley value
- Induced subgraphs, marginal contribution nets
- Simple games, weighted voting games

Lessons learnt: resource allocation

- How to allocate goods given preferences of agents?
- Auctions: English, Dutch, FPSB, Vickrey
- Incentive compatibility, lying, collusion, shills
- Combinatorial auctions, bidding languages
- The VCG mechanism

Lessons learnt: bargaining

- How to behave in a negotiation to get the best deal?
- Alternating offers protocol, ultimatum games & time
- Negotiation in task-oriented domains
- Monotonic concession protocol & Zeuthen strategy
- Bargaining for resource allocation
- Finding allocations using different contracts

Lessons learnt: argumentation

- Negotiation using the possibility to “give reasons”
- Making decisions in the presence of conflicting knowledge
- Abstract argumentation systems, extensions
- Logic-based argumentation
- Argumentation dialogue systems

Lessons learnt: Logics for multiagent systems

- Logical modelling of MAS
- Modal logic framework, possible worlds semantics
- Axiom systems & accessibility relations (correspondence theory)
- Epistemic logic, common & distribued knowledge

So how does it all come together?

- As said at start of terms: agent-based systems = study of integration of intelligent systems
- Some methods concerned with abstract modelling of systems (abstract architectures, formal logic)
- . . . others with organising joint behaviour of different components (architectures, coordination methods)
- . . . and others with optimisation in the presence of different interests (game-theoretic topics, argumentation)
- All these are pieces in the puzzle
  - but show breadth of techniques used
  - AI legacy vs. maths vs. logic vs. economics vs. distributed systems
  - field still struggles to find a topic that is not also addressed by others
  - that's a good and bad thing!

What we haven't talked about

- Multiagent learning
- Trust and reputation
- Mobile agents
- Matchmaking and brokering
- Multiagent organisations

What we haven't talked about

- Multi-robot systems/distributed sensor networks
- Distributed search & distributed constraint satisfaction
- Agent programming languages and APIs
- Virtual agents, lifelike characters
- Agent-oriented software engineering
- Social computation

The exam

- Two hours, two out of three questions of equal size
- Roughly speaking for k marks the answer should contain k items
- Emphasis on things you can define formally, calculate, or explain
- Occasionally a short discussion question
- All lecture and tutorial material can be examined
- You won't have to do complex or very lengthy calculations

The End

Thanks for your attention and participation, and good luck with the exam!