CMSC 671
Fall 2010

Thu 12/02/10

Multiagent Systems
Overview and Research Directions

Based on the lecture slides "Multiagent Systems",
Marie desJardins, MAPLE, UMBC

Prof. Laura Zavala, laura.zavala@umbc.edu, ITE 373, 410-455-8775

Outline

- What's an Agent?
  - A computer system/program that perceives and acts, and is capable of *autonomous action* in this environment.
    - Might have partial representation of the environment, act under uncertainty and dynamism, and learn.
- Multi-Agent Systems
  - Cooperative multi-agent systems
  - Competitive multi-agent systems
- MAS Research Directions
  - Organizational structures
  - Communication limitations
  - Learning in multi-agent systems

Multi-agent systems

- Jennings et al.'s key properties:
  - Situated
  - Autonomous
  - Flexible:
    - Responsive to dynamic environment
    - Pro-active / goal-directed
    - ***Social interactions with other agents and humans***
- Research questions: How do we design agents to ***interact effectively*** to solve a wide range of problems in many different environments?

Aspects of multi-agent systems

- Cooperative vs. competitive
- Homogeneous vs. heterogeneous

- Interaction protocols and languages
- Organizational structure
- Mechanism design / market economics
- Learning

Topics in multi-agent systems

- Cooperative MAS:
  - Distributed problem solving: Less autonomy
  - Distributed planning: Models for cooperation and teamwork
- Competitive or self-interested MAS:
  - Distributed rationality: Voting, auctions
  - Negotiation: Contract nets

Cooperative Multi-Agent Systems

Typical cooperative MAS domains

- Distributed sensor networks
- Distributed vehicle monitoring
- Distributed delivery
- Teams that play against other teams (robocup soccer annual competition)

Distributed sensing

- Track vehicle movements using multiple sensors
- Distributed sensor network establishment:
  - Locate sensors to provide the best coverage
  - Centralized vs. distributed solutions
- Distributed vehicle monitoring:
  - Control sensors and integrate results to track vehicles as they move from one sensor's "region" to another's

Distributed delivery

- Logistics problem: move goods from original locations to destination locations using multiple delivery resources (agents)
- Dynamic, partially accessible, nondeterministic environment (goals, situation, agent status)
- Centralized vs. distributed solution

Distributed problem solving/planning

- **Cooperative agents, working together to solve complex problems with local information**

- Problem solving in the classical AI sense, distributed among multiple agents
- Agents may be heterogeneous or homogeneous
- DPS implies that agents must be cooperative (or, if self-interested, then rewarded for working together)

Distributed problem solving/planning

- **Cooperative agents, working together to solve complex problems with local information**

- Multiagent Planning
- Distributed Constraint Satisfaction

Multiagent Planning

- Planning with multiple agents
  - Each agent makes its plan
  - Joint actions
    - $<a_1, \ldots, a_n>$ where a_i is the action taken by the *ith* actor
    - Transition model and joint planning problem
      - Complexity of the problem grows exponentially
  - Loosely coupled agents
  - Goals and knowledge base might or might not be shared
    - Can each agent just compute the joint solution and execute its own part?
      - There is no right single joint solution
  - Agents need Coordination
    - Communication (implicit or explicit)

Distributed Constraint Satisfaction

- Looks at solving CSP when there is a collection of agents, each of which controls a subset of the constraint variables.

DCSP: Approach

- If we represent the search problem as a graph, we can solve it by accumulating local computations for each node in the graph
  - Local computations can be executed asynchronously and concurrently

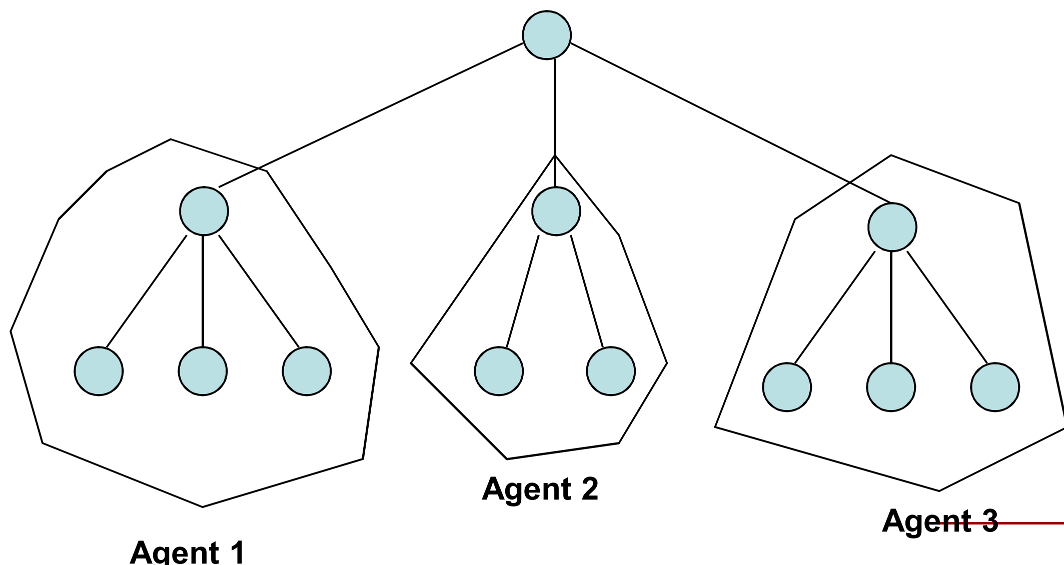

Asynchronous Backtracking

- The processes are priority ordered (by the alphabetical order of the variable identifiers)
- Each process chooses an assignment and communicates it to the neighboring processes (**ok message**)
- Each process maintains the current value of other processes from its viewpoint (**local view**)
  - A value assignment is changed if it is not consistent with the assignments of the higher priority processes
  - If no values are consistent with the higher priority processes, then the process creates a **nogood message** and sends it to the higher priority processes
- All agents wait for and respond to messages

Distributed Constraint Satisfaction

- Looks at solving CSP when there is a collection of agents, each of which controls a subset of the constraint variables.
- Agents need communication too
  - DCS algorithms rely on message exchanging among the agents

Competitive Multi-Agent Systems

Distributed rationality

- **Techniques to encourage/coax/force self-interested agents to play fairly in the sandbox**

- **Voting**: Everybody's opinion counts (but how much?)
- **Auctions**: Everybody gets a chance to earn value (but how to do it fairly?)
- **Contract nets**: Work goes to the highest bidder
- **Issues**:
  - Global utility
  - Fairness
  - Stability
  - Cheating and lying

Pareto optimality

- S is a **Pareto-optimal** solution iff
  - $\forall S' (\exists x\ U_x(S') > U_x(S) \rightarrow \exists y\ U_y(S') < U_y(S))$
  - i.e., no one could be made better off without making someone else worse off
- Social welfare, or global utility, is the sum of all agents' utility
  - If S maximizes social welfare, it is also Pareto-optimal (but not vice versa)

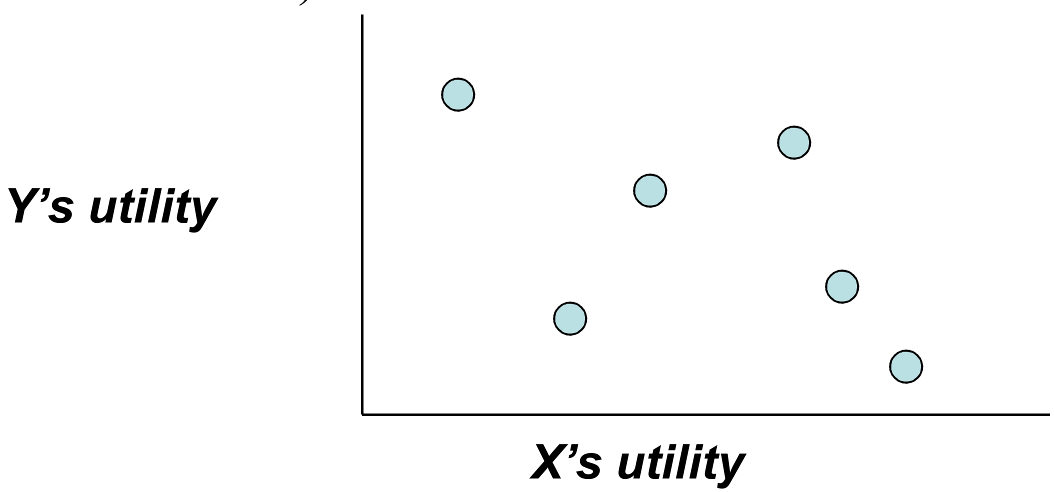

Which solutions are Pareto-optimal?

Which solutions maximize global utility (social welfare)?

Stability

- If an agent can always maximize its utility with a particular strategy (regardless of other agents' behavior) then that strategy is **dominant**
- A set of agent strategies is in **Nash equilibrium** if each agent's strategy S_i is locally optimal, given the other agents' strategies
  - Knowing the other agent's strategies, no agent has an incentive to change strategies

Prisoner's Dilemma

- Demonstrates why two people might not cooperate even if it is in both their best interests to do so
- Two suspects are arrested by the police. The police have insufficient evidence for a conviction, and, having separated the prisoners, visit each of them to offer the same deal.

A \ B	Cooperate	Defect
Cooperate	3, 3	0, 5
Defect	5, 0	1, 1

Prisoner's Dilemma: Exercise

- Strategies are: Defect or Cooperate
- What is the Pareto-optimal and social welfare maximizing solution?
- What is the dominant strategy and Nash equilibrium?

A \ B	Cooperate	Defect
Cooperate	3, 3	0, 5
Defect	5, 0	1, 1

- Pareto optimal: no one could be made better off without making someone else worse off
- Dominant: an agent can always maximize its utility with a particular strategy
- Nash E.: Knowing the other agent's strategies, no agent has an incentive to change strategies

Voting

- How should we rank the possible outcomes, given individual agents' preferences (votes)?
- Six desirable properties (which **can't all simultaneously be satisfied**):
  - Every **combination of votes** should lead to a **ranking**
  - Every **pair of outcomes** should have a **relative ranking**
  - The ranking should be **asymmetric and transitive**
  - The ranking should be **Pareto-optimal**
  - **Irrelevant alternatives** shouldn't influence the outcome
  - **Share the wealth**: No agent should always get their way ☺

Voting protocols

- **Plurality voting**: the outcome with the highest number of votes wins
  - Irrelevant alternatives can change the outcome: The Ross Perot factor
- **Borda voting**: Agents' rankings are used as weights, which are summed across all agents
  - Agents can "spend" high rankings on losing choices, making their remaining votes less influential
- **Binary voting**: Agents rank sequential pairs of choices ("elimination voting")
  - Irrelevant alternatives can still change the outcome
  - Very order-dependent

Auctions

- Many different types and protocols
- All of the common protocols yield Pareto-optimal outcomes
- ***But***… Bidders can agree to artificially lower prices in order to cheat the auctioneer
- What about when the colluders cheat each other?
    - (Now that's *really* not playing nicely in the sandbox!)

Contract nets

- Simple form of negotiation
- Announce tasks, receive bids, award contracts
- Many variations: directed contracts, timeouts, bundling of contracts, sharing of contracts, …
- There are also more sophisticated dialogue-based negotiation models

Agent Communication

Communication models

- Interaction Protocols – FIPA Protocols
  - Specifies flow of messages
  - FIPA Request Protocol, Contract Net Protocol, etc
- Communication Language: Speech act theory, KQML
  - Request, Accept, Reject, etc.
- Content language:
  - Ontologies or any other content language to represent the elements of the domain

FIPA Request Interaction Protocol

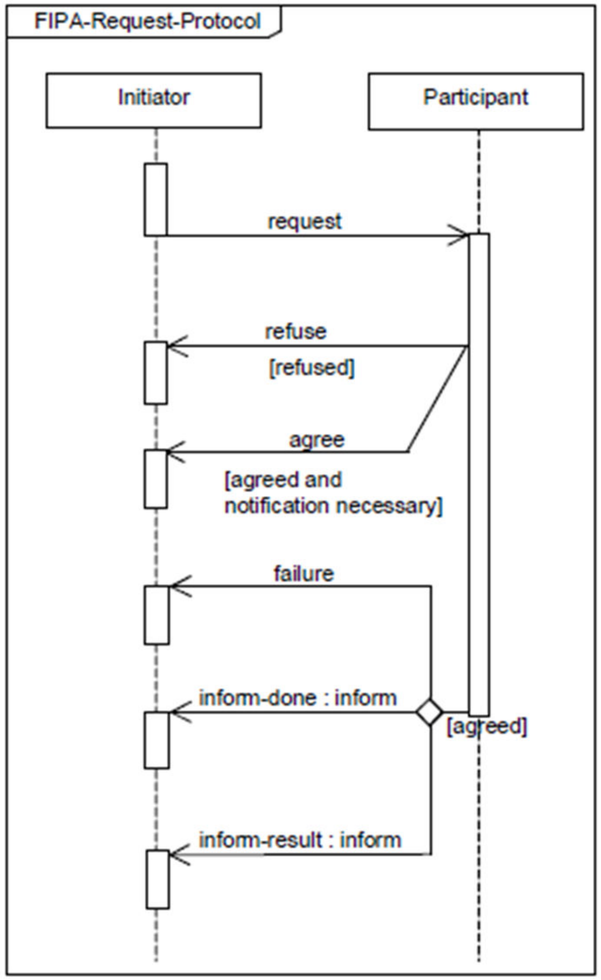

Figure 1: FIPA Request Interaction Protocol

Learning in MAS

- **Emerging field to investigate how teams of agents can learn individually and as groups**
- **Distributed reinforcement learning**: Behave as an individual, receive team feedback, and learn to individually contribute to team performance
  - Iteratively allocate “credit” for group performance to individual decisions
- **Genetic algorithms**: Evolve a society of agents (survival of the fittest)
- **Strategy learning**: In market environments, learn other agents’ strategies

Conclusions and directions

- "Agent" means many different things
- Different types of "multi-agent systems":
  - Cooperative vs. competitive
  - Heterogeneous vs. homogeneous
- Lots of interesting/open research directions:
  - Effective cooperation strategies
  - "Fair" coordination strategies and protocols
  - Learning in MAS
  - Resource-limited MAS (communication, …)