Distributed Software Development

Self-interested Agents

Chris Brooks

Department of Computer Science

University of San Francisco

23-2: Engineering systems vs Engineering agents

- Recall that at the end of Thursday's class, we were talking about ant algorithms.
  - By specifying a simple set of rules, we can achieve interesting large-scale behavior.
- Ant-type approaches lead us to think about how we can build systems that produce the effects we want.
- "Given that agents will act in a particular way, how can we constrain the environment to achieve a desirable outcome?"
- This method of problem solving is best applied to problems involving self-interested agents.

23-3: Preferences and Utility

- Agents will typically have preferences over outcomes
  - This is declarative knowledge about the relative value of different states of the world.
  - "I prefer ice cream to spinach"
- Often, the value of an outcome can be quantified (perhaps in monetary terms.)
- This allows the agent to compare the utility (or expected utility) of different actions.
- A rational agent is one that maximizes expected utility.
- Self-interested agents each have their own utility function.

23-4: Rationality and protocol design

- By treating participants as rational agents, we can exploit techniques from game theory and economics.
- Assume everyone will act to maximize their own payoff
- How do we structure the rules of the game so that this behavior leads to a desired outcome?
- This approach is called mechanism design.

23-5: Example: Clarke tax

- Assume that we want to find the shortest path through a graph.
- Each edge is associated with an agent.
- Each edge has a privately known transmission cost.
  - Agents might choose to lie about their transmission cost.
- How can we find the shortest path?

23-6: Clarke tax

- Rule:
  - Accept each agent's bid.
  - If they are not on the shortest path, they get 0.
  - If they are on the shortest path, they get:
    - Cost of next shortest path - (cost of shortest path without their contribution).

23-7: Example

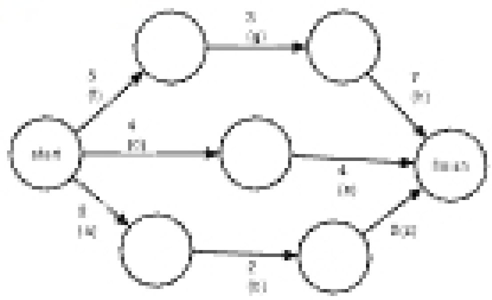

- Assume each agent bids truthfully.
- Agents A, B, and C are each paid 8 - (6 - 2) = 4
  - This is their contribution to the 'best solution'
- Other agents are paid nothing.

23-8: Example

- Why is truth-telling a dominant strategy?
  - What if a underbids?
    - A bids 1: paid 8 - (5 - 1) = 4. No benefit.
  - What if A overbids?
    - A bids 3: paid 8 - (7 - 3) = 4. No benefit.
    - A bids 5. No longer on the shortest path, so A gets 0.
  - What if d underbids?
    - D bids 3: no change.
    - D bids 1: paid 6 - (5 - 1) = 2. But his cost is 4.
  - D overbids: no change.

23-9: Solution concepts

- So how do we evaluate an algorithm or protocol involving self-interested agents?
- Some solutions may be better for some agents and worse for others.
  - Example: cake-cutting problem
- We know that each agent will try to maximize its own welfare
- What about the system as a whole?

23-10: Solution concepts

- There are a number of potential solution concepts we can use:
  - Social welfare - sum of all agent utility.
  - Pareto efficiency
    - Is there a solution that makes one agent better off without making anyone worse off?
  - Individual rationality
    - An agent who participates in the solution should be better off than if it hadn't participated.
  - Stability
    - The mechanism should not be able to be manipulated by one or more agents.
- It's not usually possible to optimize all of these at the same time.

23-11: Stability

- Ideally, we can design mechanisms with **dominant strategies**
  - A dominant strategy is the best thing to do no matter what any other agent does.
  - In the previous example, truth-telling was a dominant strategy.
  - We would say that the mechanism is non-manipulable. (lying can't break it.)
- Unfortunately, many problems don't have a dominant strategy.
- Instead, the best thing for agent 1 to do depends on what agents 2,3,4,... do.

23-12: Nash equilibrium

- This leads to the concept of a ***Nash equilibrium***
- A set of actions is a Nash equilibrium if, for every agent, given that the other agents are playing those actions, it has no incentive to change.
- Example: big monkey and little monkey
  - Monkeys usually eat ground-level fruit
  - Occasionally they climb a tree to get a coconut (1 per tree)
  - A Coconut yields 10 Calories
  - Big Monkey expends 2 Calories climbing the tree.(net 8 calories)
  - Little Monkey expends 0 Calories climbing the tree. (net 10 calories)

23-13: Nash equilibrium

- If BM climbs the tree
  - BM gets 6 C, LM gets 4 C
  - LM eats some before BM gets down
- If LM climbs the tree
  - BM gets 9 C, LM gets 1 C
  - BM eats almost all before LM gets down
- If both climb the tree
  - BM gets 7 C, LM gets 3 C
  - BM hogs coconut
- How should the monkeys each act so as to maximize their own calorie gain?

23-14: Nash equilibrium

- Assume BM decides first
  - Two choices: wait or climb
- LM has four choices:
  - Always wait, always climb, same as BM, opposite of BM.

23-15: Nash equilibrium

- What should Big Monkey do?
- If BM waits, LM will climb (1 is better than 0): BM gets 9
- If BM climbs, LM will wait :BM gets 4
- BM should wait.
- What about LM?
  - LM should do the opposite of BM.
- This is a Nash equilibrium. For each monkey, given the other's choice, it doesn't want to change.
- Each monkey is playing a **best response**.

23-16: Nash equilibrium

- Nash eqilibria are nice in systems with rational agents.
- If I assume other agents are rational, then I can assume they'll play a best response.
- I only need to consider Nash equilibria.
- They are ***efficient*** (in the Pareto sense).
- Problems:
  - There can be many Nash equilibria. (the cake-cutting problem has an infinite number of Nash equilibria)
  - Some games have no Nash equilibrium.
  - There may be ways for groups of agents to cheat.

23-17: Selecting between equilibria

- Given that there are lots of possible Nash equilibria in a problem, how does an agent choose a strategy?
- In some cases, external forces are used to make one equilibrium more attractive.
  - Government regulation, taxes or penalties
- In other cases a natural ***focal point*** exists.
  - There is a solution that is attractive or sensible ***outside the scope of the game.***

23-18: Bilateral and multilateral negotiation

- There are two different ways that we can think about agents negotiating or bargaining with each other.
- Bilateral: negotiation happens one-on-one.
  - Game theory is applicable here.
- Multilateral: Many agents negotiate simultaneously.
  - Markets and auctions are appropriate here.