LECTURE 10: REACHING AGREEMENT

An Introduction to Multiagent Systems

CIS 716.5, Spring 2005

Reaching Agreements

- How do agents *reaching agreements* when they are self interested?

- In an extreme case (zero sum encounter) no agreement is possible — but in most scenarios, there is potential for *mutually beneficial agreement* on matters of common interest.

- The capabilities of *negotiation* and *argumentation* are central to the ability of an agent to reach such agreements.

Mechanisms, Protocols, and Strategies

- Negotiation is governed by a particular *mechanism*, or *protocol*.

- The mechanism defines the "rules of encounter" between agents.

- *Mechanism design* is designing mechanisms so that they have certain desirable properties.

- Given a particular protocol, how can a particular *strategy* be designed that individual agents can use?

Mechanism Design

Desirable properties of mechanisms:

- *Convergence/guaranteed success*.

- *Maximising social welfare*.

- *Pareto efficiency*.

- *Individual rationality*.

- *Stability*.

- *Simplicity*.

- *Distribution*.

Auctions

- An auction takes place between an agent known as the *auctioneer* and a collection of agents known as the *bidders*.
- The goal of the auction is for the auctioneer to allocate the *good* to one of the bidders.
- In most settings the auctioneer desires to maximise the price; bidders desire to minimise price.

Auction Parameters

- Goods can have

  private value   *public/common value*;   *correlated value*

- Winner may pay according to

  first price;   *second price*.

- Bids may be

  open cry   *sealed bid*.

- Bidding may be:

  one shot;   *ascending*   *descending*.

English Auctions

- Most commonly known type of auction:
  - *first-price*,
  - *open cry*,
  - *ascending*.
- Dominant strategy is for agent to successively bid a small amount more than the current highest bid until it reaches their valuation, then withdraw.
- Susceptible to:
  - *winners curse*;
  - *shills*.

Dutch Auctions

Dutch auctions are examples of *open-cry descending* auctions:

- auctioneer starts by good at artificially high value;
- auctioneer lowers offer price until some agent makes a bid equal to the current offer price;
- the good is then allocated to the agent that made the offer.

First-Price Sealed-Bid Auctions

First-price sealed-bid auctions are *one-shot auctions*:

- there is a single round;
- bidders submit a sealed bid for the good;
- good is allocated to agent that made highest bid.
- winner pays price of highest bid.

Best strategy is to *bid less than true valuation*.

Vickrey Auctions

- Vickrey auctions are:
  - *second-price*;
  - *sealed-bid*.
- Good is awarded to the agent that made the highest bid; at the price of the *second highest* bid.
- *Bidding to your true valuation is dominant strategy in Vickrey auctions*.
- Vickrey auctions susceptible to *antisocial* behavior.

Negotiation

- Auctions are *only* concerned with the allocation of goods: richer techniques for reaching agreementsare required.
- *Negotiation* is the process of reaching agreements on matters of common interest.
- Any negotiation setting will have four components:
  - A negotiation set: possible proposals that agents can make.
  - A protocol.
  - Strategies, one for each agent, which are private.
  - A rule that determines when a deal has been struck and what the agreement deal is.

Negotiation usually proceeds in a series of rounds, with every agent making a proposal at every round.

Negotiation in Task-Oriented Domains

Imagine that you have three children, each of whom needs to be delivered to a different school each morning. Your neighbour has four children, and also needs to take them to school. Delivery of each child can be modelled as an indivisible task. You and your neighbour can discuss the situation, and come to an agreement that it is better for both of you (for example, by carrying the other's child to a shared destination, saving him the trip). There is no concern about being able to achieve your task by yourself. The worst that can happen is that you and your neighbour won't come to an agreement about setting up a car pool, in which case you are no worse off than if you were alone. You can only benefit (or do no worse) from your neighbour's tasks. Assume, though, that one of my children and one of my neigbours's children both go to the same school (that is, the cost of carrying out these two deliveries, or two tasks, is the same as the cost of carrying out one of them). It obviously makes sense for both children to be taken together, and only my neighbour or I will need to make the trip to carry out both tasks.

TODs Defined

- A TOD is a triple

$$\langle T, Ag, c \rangle$$

where:

  - T is the (finite) set of all possible tasks;
  - $Ag = \{1, \ldots, n\}$ is set of participant agents;
  - $c : \wp(T) \to \mathbb{R}^+$ defines *cost* of executing each subset of tasks:

- An *encounter* is a collection of tasks

$$\langle T_1, \ldots, T_n \rangle$$

where $T_i \subseteq T$ for each $i \in Ag$.

Deals in TODs

- Given encounter $\langle T_1, T_2 \rangle$, a *deal* will be an allocation of the tasks $T_1 \cup T_2$ to the agents 1 and 2.

- The *cost* to i of deal $\delta = \langle D_1, D_2 \rangle$ is $c(D_i)$, and will be denoted $cost_i(\delta)$.

- The *utility* of deal δ to agent i is:

$$utility_i(\delta) = c(T_i) - cost_i(\delta).$$

- The *conflict deal*, Θ, is the deal $\langle T_1, T_2 \rangle$ consisting of the tasks originally allocated.
 Note that

$$utility_i(\Theta) = 0 \quad \text{for all } i \in Ag$$

- Deal δ is *individual rational* if it gives positive utility.

The Negotiation Set

- The set of deals over which agents negotiate are those that are:

  - individual rational
  - pareto efficient.

- Individually rational: agents won't be interested in deals that give negative utility since they will prefer the conflict deal.

- Pareto efficient: agents can always transform a non-Pareto efficient deal into a Pareto efficient deal by making one agent happier and none of the others worse off.

The Negotiation Set Illustrated

The Monotonic Concession Protocol

Rules of this protocol are as follows. . .

- Negotiation proceeds in rounds.

- On round 1, agents simultaneously propose a deal from the negotiation set.

- Agreement is reached if one agent finds that the deal proposed by the other is at least as good or better than its proposal.

- If no agreement is reached, then negotiation proceeds to another round of simultaneous proposals.

- In round $u + 1$, no agent is allowed to make a proposal that is less preferred by the other agent than the deal it proposed at time u.

- If neither agent makes a concession in some round $u > 0$, then negotiation terminates, with the conflict deal.

The Zeuthen Strategy

Three problems:

- What should an agent's first proposal be?
  Its most preferred deal

- On any given round, *who should concede*?
  The agent least willing to risk conflict.

- If an agent concedes, then *how much* should it concede?
  Just enough to change the balance of risk.

Willingness to Risk Conflict

- Suppose you have conceded a *lot*. Then:

  – Your proposal is now near to conflict deal.
  – In case conflict occurs, you are not much worse off.
  – You are *more willing* to risk confict.

- An agent will be *more willing* to risk conflict if the difference in utility between its current proposal and the conflict deal is *low*.

Nash Equilibrium Again. . .

The Zeuthen strategy is in Nash equilibrium: under the assumption that one agent is using the strategy the other can do no better than use it himself. . .

This is of particular interest to the designer of automated agents. It does away with any need for secrecy on the part of the programmer. An agent's strategy can be publicly known, and no other agent designer can exploit the information by choosing a different strategy. In fact, it is desirable that the strategy be known, to avoid inadvertent conflicts.

Deception in TODs

Deception can benefit agents in two ways:

- *Phantom and Decoy tasks*.
  Pretending that you have been allocated tasks you have not.

- *Hidden tasks*.
  Pretending *not* to have been allocated tasks that you have been.

Argumentation

- Argumentation is the process of attempting to convince others of something.

- Gilbert (1994) identified 4 modes of argument:

  1. *Logical mode*.
     "If you accept that A and that A implies B, then you must accept that B".
  2. *Emotional mode*.
     "How would you feel if it happened to you?"
  3. *Visceral mode*.
     "Cretin!"
  4. *Kisceral mode*.
     "This is against Christian teaching!"

Logic-based Argumentation

Basic form of logical arguments is as follows:

$$Database \vdash (Sentence, Grounds)$$

where:

- *Database* is a (possibly inconsistent) set of logical formulae;

- *Sentence* is a logical formula known as the *conclusion*; and

- *Grounds* is a set of logical formulae such that:

  1. *Grounds* $\subseteq$ *Database*; and
 2. *Sentence* can be proved from *Grounds*.

Attack and Defeat

- Argumentation takes into account the relationship between arguments.

- Let (ϕ_1, Γ_1) and (ϕ_2, Γ_2) be arguments from some database $\Delta \ldots$
  Then (ϕ_2, Γ_2) can be defeated (attacked) in one of two ways:

  1. (ϕ_1, Γ_1) *rebuts* (ϕ_2, Γ_2) if $\phi_1 \equiv \neg\phi_2$.
  2. (ϕ_1, Γ_1) *undercuts* (ϕ_2, Γ_2) if $\phi_1 \equiv \neg\psi$ for some $\psi \in \Gamma_2$.

- A rebuttal or undercut is known as an *attack*.

Abstract Argumentation

- Concerned with the overall structure of the argument (rather than internals of arguments).

- Write $x \rightarrow y$

  – "argument x attacks argument y";
  – "x is a counterexample of y; or
  – "x is an attacker of y".

  where we are not actually concerned as to what x, y are.

- An *abstract argument system* is a collection or arguments together with a relation "$\rightarrow$" saying what attacks what.

- An argument is *out* if it has an undefeated attacker, and *in* if all its attackers are defeated.

An Example Abstract Argument System

Argumentation and Communication

- We have two agents, P and C, each with some knowledge base, Σ_P and Σ_C.

- Each time one makes an assertion, it is considered to be an addition to its *commitment store*, $CS(P)$ or $CS(C)$.

- Thus P can build arguments from $\Sigma_P \cup CS(C)$, and C can use $\Sigma_C \cup CS(P)$.

- We assume that dialogues start with P making the first move.

- The outcomes, then, are:

  – P generates an argument both classify as IN, or
  – C makes P's argument OUT.

Argumentation Protocol

- A typical persuasion dialogue would proceed as follows:

  1. P has an acceptable argument (S, p), built from Σ_P, and wants C to accept p.
  2. P asserts p.
  3. C has an argument $(S', \neg p)$.
  4. C asserts $\neg p$.
  5. P cannot accept $\neg p$ and challenges it.
  6. C responds by asserting S'.
  7. P has an argument $(S'', \neg q)$ where $q \in S'$, and challenges q.
  8. . . .

Argumentation Protocol II

- This process eventually terminates when

$$\Sigma_P \cup CS(P) \cup CS(C)$$

and

$$\Sigma_C \cup CS(C) \cup CS(P)$$

eventually provide the same set of IN arguments and the agents agree.

Summary

- This lecture has looked at different mechanisms for reaching agreement between agents.
- We started by briefly describing different auction mechansims.
- We then passed on to looking at negotiation, where agents make concessions and explore tradeoffs.
- Finally, we looked at argumentation, which allows for much more complex interactions.