CM30174: Intelligent Agents

Marina De Vos, Julian Padget

Coalitions / version 0.3

November 9, 2010

Authors/Credits for this lecture

- Onn Shehory: for slides and discussions

Content

Overview

- Fit with preceding material: relationship to Game Theory and Mechanism Design
- Fit with properties of agents: sociability and cooperation – even if self-interested
- Builds on communication, cooperation and negotiation
- Moving towards concept of agent societies

What is an Agent?

- An intelligent agent is a computer system capable of flexible, autonomous action in some environment: the situated agent.

What are *Multi*-Agent Systems?

- An agent can be more useful in the context of others:
  - Can concentrate on tasks within competence
  - Can delegate other tasks
  - Can use ability to communicate, coordinate, negotiate

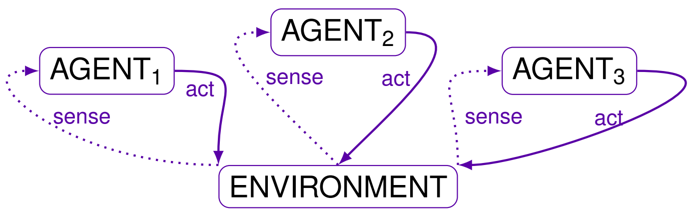

What are *Multi*-Agent Systems?

So, a MAS is a collection of interacting agents? No:

- Needs meaningful ways for agents to interact
- Needs organizational framework
- Needs identification of roles, responsibilities, permissions
- Needs to be verified
- Needs to be validated

What is a Coalition?

- Coalitions are (temporary) collections of individuals working together for the purpose of achieving a task
- Coalition formation is the process whereby an agent decides to cooperate with other agents
- Because
  - Either: task cannot be performed by a single agent
  - Or: task could be performed more efficiently by several
- Agents bring different, complementary capabilities to the coalition
- When the task is completed, the payoff is distributed and agents continue to pursue their own agenda

Games and Cooperation

- Game theory—prisoner's dilemma—concludes defection is best strategy
- Why?
  - no binding agreements
  - utility $\rightarrow$ *individuals* following *individual* action
- Real-world relies on contracts etc.
- Organizations receive revenue then distribute to individuals
- $\rightsquigarrow$ cooperative games

Cooperation via Coalitions

- To perform a task and increase benefits, agents may need to cooperate via coalition formation
- A coalition: a set of agents that agree to cooperate to perform a task
- Given n agents, k tasks, there are $k(2^n - 1)$ different possible coalitions
- The number of configurations is $O(n^{(n/2)})$
- Hence, exhaustive search is infeasible

Issues in Coalition Formation

- Given a task and other agents, which coalitions should an agent attempt to form?
- What mechanism can an agent use for coalition formation?
- What guarantees regarding efficiency and quality can the mechanism provide?
- Once a coalition is created, how should its members handle distribution of work/payoff?
- When, and how, does a coalition dissolve?

Solution Types

- Self-interest vs. benevolence: the mechanisms for benevolent agents are usually much simpler, as such agents do not need means to maintain their own payoff maximization.
- Centralization vs. distribution: central design of coalitions is usually much simpler to execute and enforce than a distributed one.
- Environment super-additivity: in super-additive environments any unification of two coalitions increases overall payoff. Strongly influences the mechanism.

Coalition formation: external

- By imposition: an external agency makes decisions
- Agents advertise skills and prices
- Requestor defines properties of coalition
- External process computes optimal coalition
- Essentially the same as combinatorial auction – same complexity
- See “Generating Coalition Structures with Finite Bound from the Optimal Guarantees”, [Dang and Jennings, 2004]

Coalition formation: internal

- By self-organization: coalitions are established by group (inter)actions
- Process: multi-lateral negotiation
- Identification of tasks (responsibilities)
- Negotiation of outcomes (self-interested agents)
- Examples: Robocup, Robocup rescue
- See “Methods for Task Allocation via Agent Coalition Formation”, [Shehory and Kraus, 1998]
- And “Self-organization through bottom-up coalition formation”, [Sims et al., 2003]

Cooperative Games

- n agents, typically $n > 2$
- $Ag = \{1, ..., n\}$
- coalition C is a subset of Ag that may (or may not) work together
- *grand coalition* is when $C = Ag$
- *singleton* coalition contains just one agent
- A coalition has a utility
- Cooperative game is $\mathcal{G} = \langle Ag, \nu \rangle$, where $\nu : \mathbf{2}^{Ag} \to \mathbb{R}$ is the *characteristic* function
- $\nu(C) = k$ denotes the utility deriving from coalition C
- Does *not* specify distribution of utility
- Does *not* explain how ν is derived

Coalition formation activities

- Coalition structure generation:
  - Partition into exhaustive disjoint coalitions
  - Given $\{a_1, a_2, a_3\}$, $\exists$ seven coalitions:

$$\{a_1\}, \{a_2\}, \{a_3\}, \{a_1, a_2\}, \{a_1, a_3\}, \{a_2, a_3\}, \{a_1, a_2, a_3\}$$

  - And five coalition structures:

$$\{a_1, a_2, a_3\},$$
$$\{\{a_1\}, \{a_2, a_3\}\}, \{\{a_2\}, \{a_1, a_3\}\}, \{\{a_3\}, \{a_1, a_2\}\},$$
$$\{\{a_1\}, \{a_2\}, \{a_3\}\}$$

  Probably not desirable to generate all CSs in advance
- Optimizing coalition value: pooling the tasks and resources of the agents to maximize the coalition value
- Payoff distribution: deciding how to distribute the payoff between coalition members (equally, inputs, outputs, role)

Goals of coalition formation 1/2

- Game theory typically only considers super-additive environments where any two disjoint coalitions are better off merging, resulting in the grand coalition of all agents.
- Inappropriate for real-world problems:
  - Ignores cost of coalition formation
  - Ignores cost of coalition coordination
- For non-super-additive environments aim to maximize social welfare... but known NP-hard problem.

Goals of coalition formation 2/2

- Given a set of agents and a set of tasks, want to identify a mapping between tasks and (sub-)groups of agents because:
  - Either: task *cannot* be performed by a single agent
  - Or: task could be performed *more efficiently* by several
- Overlapping coalitions make problem harder, but in general cost is NP-hard and solutions are approximations to NP-hard algorithms.
- See “Methods for task allocation via agent coalition formation” [Shehory and Kraus, 1998]

Where to start?

- How to decide who to work with? Need criteria
- Goal: to join the *best* coalition
- Criteria may vary, but broadly similar?
- Only some coalitions are attractive
- Coalition can only exist if agents choose to be members
- $\equiv$ agent cannot do better by *defecting*
- Replace "which coalition to join?"
  by "which coalitions are stable?"
- Identify *core* as set of feasible distributions

$$\langle x_1, ..., x_k \rangle \ s.t. \sum_{i \in C} x_i = \nu(C)$$

Computing the core

- Example (from [Wooldridge, 2009], p.273)
- $Ag = \langle 1, 2 \rangle$ and $\nu(\{1\}) = 5, \nu(\{2\}) = 5, \nu(\{1,2\}) = 20$
- Outcomes are $\{\langle 20, 0 \rangle, ..., \langle 0, 20 \rangle\}$ assuming integer valued utility
- The (singelton) coalitions are no worse off in the grand coalition for the outcomes $\{\langle 15, 5 \rangle, ..., \langle 5, 15 \rangle\}$
- Thus core is non-empty and identifies a set of feasible distributions, hence coalition itself is stable
- Issues:
  - core is empty
  - core non-empty, but not "fair"
  - computationally hard for large coalitions

Shapley's axioms

- Define $\mu_i(C)$ as the *marginal contribution* i adds to C and sh_i is i's share of the outcome
- Axioms for fair distribution of coalitional value:
  1. **Symmetry** agents that do the same, get the same

     $$\forall C \in Ag \setminus \{i, j\}, \mu_i(C) = \mu_j(C) \text{ then } sh_i = sh_j$$

  2. **Dummy player** agent that does not enhance coalition, gets individual outcome

     $$\forall C \in Ag \setminus \{i\}, \mu_i(C) = \nu(\{i\}) \text{ then } sh_i = \nu(\{i\})$$

  3. **Additivity** agent that plays several games gets the sum of outcomes from each *when those games are combined*

     $$\mathcal{G}^1 \rightsquigarrow sh_i^1, \mathcal{G}^2 \rightsquigarrow sh_i^2, \mathcal{G}^{1+2} \rightsquigarrow sh_i^{1+2} \text{ then } sh_i^{1+2} = sh_i^1 + sh_i^2$$

Shapley value

- Agent gets *average marginal contribution it makes to a coalition*
- First attempt:

$$sh_i = \frac{1}{2^n - 1} \sum_{C \subseteq Ag \setminus \{i\}} \mu_i(C)$$

- But order matters: join early $\Rightarrow$ big contribution, and vv.
- Need to evaluate wrt *all possible orderings*
- $\Pi(Ag)$ are the permutations of Ag
- $C_i(o) =$ agents preceding i in $o, o \in \Pi(Ag)$
- Second attempt:

$$sh_i = \frac{1}{|Ag|!} \sum_{o \in \Pi(Ag)} \mu_i(C_i(o))$$

- Result is unique

Coalition Formation Exercise

- **Groups:** 3-4 people
- **Objective:** To consider the parameters and process of coalition formation for student group working
- **Plan:**
  - Pair up
  - Core activity [10 mins in all]
    - Identify capabilities of agents
    - Identify capabilities of tasks
    - Identify potential resources
    - Identify protocols and facilities for coalition formation
  - Reflect and discuss [5 mins]

Content

Modelling the Agents

- A set of n agents: $N = \{A_1, A_2, ..., A_n\}$
- Each agent A_i has a vector of capabilities $B_i = \langle b_1^i, ..., b_r^i \rangle$ s.t. $b_j^i \in \mathbb{R}^+$
- Each capability is a property that quantifies the agent's ability to perform a specific action.
- Resources (and thus capabilities) are limited. Each capability may be expendable (fuel) or non-expendable (carrying capacity).
- For each coalition an evaluation function is required to compute the overall capability of a group of agents: that is the element-wise sum of the capabilities $\sum_{i=0}^{k} \text{capability}(A_i)$

Modelling the Tasks

- A set of m independent tasks $T = \{t_1, t_2, ..., t_m\}$
- Each task t_n has a vector of capabilities $B_n = \langle b_1^n, ..., b_r^n \rangle$
- Utility gained from performing a task depends on the capabilities required to perform it. A simple (adequate) measure is a linear function of the resource amount.
- Tasks may be (partially) ordered to capture the need to complete task t_i before starting task t_j.

Modelling the Coalitions

- Thus a coalition C has a vector of capabilities B_C that is the element-wise sum of the capabilities of the member agents
- For overlapping coalitions, an agent may only contribute a fraction of a capability to each coalition of which it is a member – no double counting!
- A coalition C can perform a task t iff B_t satisfies $\bigwedge_{i=0}^{r} b_i^t \leq b_i^C$ and t has no unsatisfied predecessors (if applicable)
- A low cost approximation to the value V of a coalition is the joint utility that the members can reach by cooperating.
- Agents are group-rational: benefit of joining coalition is at least as much as the sum of benefits from not, where benefits derive from completing tasks.

Coalition formation algorithm I: 1/2

- Adapted from [Shehory and Kraus, 1998]
- Self-organization: reflects reality, but can it work computationally?
- Market economy: social, self-interested agents
- Objective: given a set of tasks, the agents partition themselves to maximize system performance
- Static organizations will probably fail in dynamic environments... RE-organization is a necessity
- Problem is equivalent to set partitioning and set covering (NP-complete), but only expensive centralized algorithms exist for optimal solutions.

Coalition formation algorithm I: 2/2

- Practical solution is to use a greedy anytime algorithm
  - greedy: makes locally optimal decision at each step hoping this results in the global optimum for each task
  - anytime: can be stopped at any time returning the best answer so far, more time *may* produce better answers
- Algorithm outline:
  - Stage 1: distributed computation of all possible coalitions and initial coalition values
  - Stage 2: iterative process
    - Recalculate coalition values
    - Agents decide upon preferred coalitions and join them
    - Agents in coalitions do not participate further (no detaching)
    - Repeat Stage 2 until there are no agents left
- An heuristic is to prefer smaller over larger coalitions

Summary of situation so far

- Coalition formation is possible and beneficial in large-scale MAS: simple, low complexity
- With no detachment, coalitions provide gains, a steady state is reached, coalitions are of medium sizes
- With detachment, gains doubled, coalitions are larger but time to reach steady state increases
- Increased detachment rate results in a slow gain then drops to zero
- Prediction of distribution of coalition sizes is possible
- Open issues: heterogeneity, perturbations, complex strategies
- See “Coalition Formation for Large-Scale Electronic Markets” [Lerman and Shehory, 2000]

Content

Request For Proposal (RFP) Coalitions

- Problem properties:
- Tasks can only be performed by groups
- A task is comprised of sub-tasks
- A task has a deadline and a value (discounted over time)
- Agents have private, subjective valuations of sub-tasks
- Agents are self-interested utility maximizers
- Solution approach:
  - Agents negotiate under time pressure to form coalitions
  - Decisions during negotiation are derived via strategies
  - Complete search of the problem space is infeasible
  - Consequently $\Rightarrow$ a simulation-based solution

Coalition negotiation

- Iterative: one proposal at each iteration
- Agents either propose or wait, committed
- More beneficial proposals are preferred
- Time is an issue because of discount
- Agents must follow protocol, can use any strategy for proposal preparation/decision
- Strategy space is intractable
- Can propose some strategies based on heuristics

Strategies

- Goal: decide which coalitions to propose to which agents and accept/reject proposals
- Strategies based on heuristics for ranking coalitions according to desirability
- General guidelines:
  - Inspect RFP tasks and sub-tasks
  - Inspect capabilities and capacities of other agents
  - Compute candidate coalitions, then rank them
- Ranking heuristics:
  - Marginal: prefer coalitions where the estimated marginal profit of the coalition is maximal
  - Expert: prefer coalitions where only a few others have the right capabilities. $\rightsquigarrow$ a better chance of winning
  - Mix of marginal and expert

Coalition Formation Algorithm II 1/3

- A negotiation-based approach
- A sequence of (bilateral) negotiations work incrementally to create a coalition bottom-up — similar to contract-net protocol
- Spectrum of protocols depending on perspective:
  - Local: individual agent utility
  - Global: marginal social utility (self plus partners)

Coalition Formation Algorithm II 2/3

- Assumptions:
  - Agents are sufficiently densely distributed as to be able to communicate with several neighbours.
  - Agents may join and leave the system at any time, but overall population is largely static.
  - Agents are cooperative.
- Problem description:

  Given a composite task T and a set of agents A:
  - *break T into m sub-tasks* $\{t_1, t_2, ...t_m\}$
  - *and A into a coalition structure* $CS = \{A_1, A_2, ..., A_m\}$,
  - *s.t. each sub-task is assigned to one coalition and the utility of each coalition is maximal.*

Coalition Formation Algorithm II 3/3

- Market context:
  - Seller: a coalition manager responsible for the agents that can satisfy the buyer's requirements. The buyer will typically be another coalition manager.
  - Product: the completion of task *T* (overall) or a sub-task (for a given coalition).
  - Value: (of a product) is a function of the marginal utility gains from the transaction.
  - Local marginal utility: is the difference between a coalition's utility before and after a transaction
  - Social marginal utility: is the sum of the local marginal utilities of buyer and seller.
- Adapted from "Self-organization through Bottom-up Coalition Formation" [Sims et al., 2003].

Negotiation: local utility case

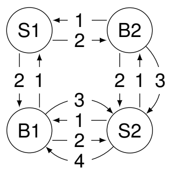

1. Buyer broadcasts message requesting a resource
2. Each "nearby" manager responsible for an agent with that resource responds *if* the local marginal utility of giving up that agent is positive
3. Buyer selects the seller whose product would provide the greatest local marginal utility gain and sends a request
4. Seller selects the buyer that maximizes the *seller's* local utility

Negotiation: social utility case

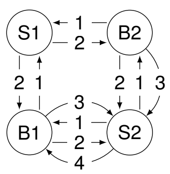

1. Buyer sends a product request [as before]
2. Seller responds regardless of marginal utility, reporting value to buyer
3. Buyer selects seller that maximizes the *sum* of the buyer's and the seller's local marginal utility (if positive) and reports the sum to the seller
4. Seller selects the buyer reporting the highest social marginal utility

Coalition Formation Algorithm III

- Coalition by imposition
  - Anytime algorithm: solutions within a finite bound of optimal
  - Let L be the set of all coalitions
  - Partition L into n subsets where $|L_i| = i$, thus L_1 is the "grand coalition" and L_n is the set of all unitary coalitions.
  - Algorithm searches L_1, L_2 and L_n to establish a bound $b = \lceil n/2 \rceil$
  - Bound is used to initiate the search through *selected* L_k where there is at least one coalition whose cardinality is not less than $\lceil n(q-1)/q \rceil$, where the initial value of q is $\lfloor \frac{n+1}{4} \rfloor$ and decreases to 2.
  - At each step the result is within a bound $b = 2q - 1$ of the optimal for a step-dependent value of q.
- See [Dang and Jennings, 2004] and p.290 of [Wooldridge, 2009]

Content

Summary

- In theory
  - algorithmic approaches
  - complexity analysis
  - metric for fairness
- In practice:
  - heuristic strategies work
  - local decision making
  - scalable

Recommended Reading

- Wooldridge: Ch.13
- [Shehory and Kraus, 1998]: not the first (or last) paper on coalition formation, but very readable and one of the most cited.
- [Sims et al., 2003] describes the bottom-up scenario in more detail.
- [Lerman and Shehory, 2000] provides a more practically oriented paper.
- [Dang and Jennings, 2004] gives the details of the bounded optimal coalition formation technique.

References

Dang, V. D. and Jennings, N. R. (2004).
Generating coalition structures with finite bound from the optimal guarantees.
In *AAMAS '04: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 564–571. IEEE Computer Society.

Lerman, K. and Shehory, O. (2000).
Coalition formation for large-scale electronic markets.
In *Proceedings of Fourth International Conference on Multi-Agent Systems (ICMAS'00)*, pages p167–? IEEE Computer Society, IEEE Press.
DOI Bookmark: http://doi.ieeecomputersociety.org/10.1109/ICMAS.2000.858449.

Shehory, O. and Kraus, S. (1998).
Methods for task allocation via agent coalition formation.
Artificial Intelligence, 101(1–2):165–200.

Sims, M., Goldman, C. V., and Lesser, V. (2003).
Self-organization through bottom-up coalition formation.
In *AAMAS '03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 867–874, New York, NY, USA. ACM Press.

Wooldridge, M. (2009).
An introduction to multiagent systems (second edition).
Wiley.
ISBN: 978-0-470-51946-2.