Multiagent Systems: Spring 2006

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

AAMAS-2006

- The International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS) is the main conference in the field.
- The fifth edition of the conference took place last week in Japan.
- The plan for today is to introduce a couple of the papers presented at that conference (concentrating on mechanism design).

Facts and Figures

- AAMAS-2006: 5th International Joint Conference on Autonomous Agents and Multiagent Systems, Hakodate, Japan, 8–12 May 2006
- Papers: 550 submissions; 127 accepted as full papers (62 with oral presentation); additionally 135 poster papers
- Attendance: $\sim$500 (?)
- Website: `http://www.fun.ac.jp/aamas2006/`

Main Sessions

- Simulation and modelling
- Logics for agent systems
- Argumentation and negotiation
- Agent planning and search
- Robotics
- Computational complexity in agent systems
- Learning and evolution
- Cooperation and coordination
- Task and resource allocation
- Ontologies and web services
- Architectures: BDI and MDPs
- Believable agents
- Auctions and electronic markets
- Trust and reputation

Papers

- R. Cavallo. *Optimal Decision-making with Minimal Waste: Strategyproof Redistribution of VCG Payments*.
- T. Matsuo, T. Ito, R.W. Day, and T. Shintani. *A Robust Combinatorial Auction Mechanism against Shill Bidders*.
- Briefly: my own papers

Paper on Redistribution of VCG Payments

- Nice properties of the Vickrey-Clarke-Groves (VCG) mechanism:
  - *Strategy-proofness:* truth-telling is dominant strategy
  - *Efficiency:* outcomes maximise utilitarian social welfare
  - *Weak budget-balance:* sum of payments to centre not negative
- Profit for the centre implementing the mechanism is appropriate in the context of auctions, but not necessarily in other cases.
- Example: *full budget-balance* would be desirable in a mechanism used to decide what TV channel to watch together
- This paper explores to what extent it is possible to give back some of the payments required by the VCG mechanism to the agents, without sacrificing strategy-proofness or efficiency.

R. Cavallo. *Optimal Decision-making with Minimal Waste: Strategyproof Redistribution of VCG Payments*. Proc. AAMAS-2006.

Recap: VCG Mechanism

- Notation: set of agents $\mathcal{A} = \{1, \ldots, n\}$; set of agreements $\mathcal{X}$; true valuations $v_i : \mathcal{X} \to \mathbb{R}$; reported valuations $\hat{v}_i : \mathcal{X} \to \mathbb{R}$
- *Allocation rule:* choose an agreement x^* that is efficient according to the reported valuations

$$x^* \in \text{argmax}_{x \in \mathcal{X}} \sum_{j=1}^{n} \hat{v}_j(x)$$

- *Pricing rule:* each agent pays the difference between what the others *could* have achieved without him and what they *did* achieve

$$p_i \quad = \quad \sum_{j \neq i} \hat{v}_j(x^*_{-i}) - \sum_{j \neq i} \hat{v}_j(x^*)$$

  Here x^*_{-i} is the top agreement for the society $\mathcal{A} \setminus \{i\}$.

- Properties: efficiency; strategy-proofness; weak budget-balance (for resource allocation with monotonic valuations)

Groves Mechanisms

Recall that the VCG mechanism is an instance of the family of Groves mechanisms, where the payment rule takes this form:

$$p_i \quad = \quad h_i - \sum_{j \neq i} \hat{v}_j(x^*)$$

Here, h_i is any function not depending on $\hat{v}_i$.

Strategy-proofness is guaranteed for any such mechanism.

Idea: If we can reduce the payments of agents but still make sure that the resulting mechanism remains an instance of the class of Groves mechanisms, then we can get closer to full budget-balance whilst keeping strategy-proofness.

Surplus Guarantee

Fixing the reported valuations $\hat{v}_j$ for all agents $j \in \mathcal{A} \setminus \{i\}$, we may consider the VCG payments p_k for all agents $k \in \mathcal{A}$ as a function of the valuation $\hat{v}_i$ reported by agent i.

Then S_i is the minimum *surplus* (sum of payments to centre) guaranteed whatever i chooses to report:

$$S_i \quad = \quad \min_{\hat{v}_i} \sum_{k \in \mathcal{A}} p_k(\hat{v}_i)$$

- Note that S_i cannot be controlled by agent i ($\leadsto$ requirement of h_i).
- Idea will be to give a kick-back of $1/n$ of S_i to agent i.

Redistribution Mechanism

The mechanism proposed by Cavallo is like the VCG mechanism, put with the following modified pricing rule:

$$p_i \quad = \quad \sum_{j \neq i} \hat{v}_j(x^*_{-i}) - \sum_{j \neq i} \hat{v}_j(x^*) - \frac{S_i}{n}$$

Clearly, this is still *strategy-proof*, as it is an instance of the class of Groves mechanisms. Intuitively, it is also still *weakly budget-balanced*.

Furthermore, it can be shown that this the closest we can get to *full budget-balance* (under some technical conditions).

In some cases, this will be a major improvement over VCG, in other cases S_i may be 0 and the two mechanisms are the same.

The paper also discusses a special case ("all-or-nothing domains") in detail, where the redistribution mechanism is particularly simple and the wasted surplus goes to 0 as n goes to ∞.

Paper on Detecting Shill Bidders

- Despite being strategy-proof, the VCG mechanism is not immune to *shill bidding* (bidding under several false identities).
- M. Yokoo has studied this problem in detail (see references given during mechanism design class). Matsuo *et al*. propose a relatively simple method for detecting shill bidders ...
- Caveat: their method is not perfect (some innocent bidders may be falsely classified as shill bidders). In fact, there can be no such method that would also retain all the other desirable properties of the VCG mechanism.

T. Matsuo, T. Ito, R.W. Day, and T. Shintani. *A Robust Combinatorial Auction Mechanism against Shill Bidders*. Proc. AAMAS-2006.

Recap: Shill Bidding

Despite being strategy-proof, the VCG mechanism is prone to manipulation by *shill* bidding (aka. *false-name* bidding). Example:

Agent 1: $(\{a\}, 0)$, $(\{b\}, 0)$, $(\{a, b\}, 4)$

Agent 2: $(\{a\}, 1)$, $(\{b\}, 1)$, $(\{a, b\}, 2)$

Agent 1 wins. But agent 2 could instead submit bids *under two names:*

Agent 1: $(\{a\}, 0)$, $(\{b\}, 0)$, $(\{a, b\}, 4)$

Agent 2: $(\{a\}, 4)$, $(\{b\}, 0)$, $(\{a, b\}, 0)$

Agent 2': $(\{a\}, 0)$, $(\{b\}, 4)$, $(\{a, b\}, 0)$

Agent(s) 2 (and 2') will win and not pay anything! This form of manipulation is particularly critical for *electronic* auctions, as it is easier to create multiple identities online than in real life.

Approach

- Run a normal VCG to determine default allocation and payments.
- Re-run the auction with each bidder excluded in turn. If bidder i is doing worse when bidder j is excluded, then the two are suspected of shill bidding and put together in a "shill bidding group".
- Re-run the auction one more time, this time treating shill bidding groups as individuals ("merging" valuations).
- Allocate items and charge payments to innocent bidders.
- For each shill bidding group, rune a new auction amongst its members for the items won by the group, using the group payment as reserve price.

Paper on Multilateral Concessions

- Objective: study how the *monotonic concession protocol* may be extended to negotiation between *more than two agents*

- Main problem: what does it mean to make a *concession* to a *group* of opponents? Possible definitions include:
  - *Strong concession:* Make a proposal that is strictly better for each of the other agents.
  - *Utilitarian concession:* Make a proposal such that the sum of utilities of the other agents increases.

- The paper discusses the properties of the resulting protocols: *termination*, *compositionality*, *deadlock-freedom*, *verifiability*.

U. Endriss. *Monotonic Concession Protocols for Multilateral Negotiation*. Proc. AAMAS-2006.

Paper on Tractable Negotiation

- Objective: exploit structure of utilities to design protocols that make it feasible for agents to negotiate an optimal allocation
- Applies to *tree-structured domains:* terms used to represent utilities in the k-additive form can be arranged as a tree
- Approach: (i) in round l, simplify all utilities by dropping terms of size $> l$; (ii) only negotiate deals over redistribution of items of one term of size $\leq l$; (iii) use a "bank" to compensate for losses
- The proposed negotiation protocols rely on two main ideas:
  - *Guidance:* reduce number of options to be considered
  - *Compensation:* relax notion of rationality to make protocol work and compensate agents (theorem: no loss for bank)

Y. Chevaleyre, U. Endriss, and N. Maudet. *Tractable Negotiation in Tree-structured Domains*. Proc. AAMAS-2006.

Paper on Rationality and Fairness in Negotiation

- Objective: study (experimentally) how resource allocations evolve in view of *egalitarian social welfare* when *rational* agents negotiate
- Study different *payment functions:* rational deals result in social surplus; how should it be divided amongst the participants?
- Experiments: comparing outcomes for different payment functions, numbers of agents, and numbers of resources
- Results: difficult to make general statements, but one interesting observation is that the seemingly most "equitable" payment function does not give the best results (explanation: loss in efficiency means less to distribute to get high fairness rating)

S. Estivie, Y. Chevaleyre, U. Endriss, and N. Maudet. *How Equitable is Rational Negotiation?* Proc. AAMAS-2006.

Conclusion

- This final lecture has been an attempt to offer a glimpse at some of the most recent research in MAS.
- We have discussed the following AAMAS-2006 papers:
  - Paper on redistribution of VCG payments by R. Cavallo
  - Paper on detecting shill bidders by T. Matsuo *et al.*
  - Brief summary of my own papers
- The papers (and the entire proceedings) are available from me for those who are interested.