

Appendix C

Properties of General English Auctions

CERTAIN classes of negotiation problems lend themselves to strategies ensuring that no agent can gain by lying. Truth incentive protocols, among which Generalized Vickrey Auction (GVA) is one of the most famous, can then be used to centrally compute fair and efficient solutions. However, for problems where truth incentive protocols lead to poor management of the social welfare (e.g. problems with false name bids) and for problems that allow no truth incentive protocols, English Auctions are preferred to GVA (Silaghi *et al.* 2001a).

C.1 Introduction

Having agents represent the interests of their owners is desirable in many practical applications. Automated negotiation is a process whereby a distributed network of agents agree on decisions on behalf of their owners. Agents negotiate on *resources* and their decisions are conditioned by *constraints* (e.g. costs, existence,...). When the available information is sub-optimally used, local decisions can lead to losses for some parties involved in negotiations. Bad decisions can also result in a decrease of the social welfare by inefficient resource allocation. There is consequently a demand for automated negotiation techniques that are fair and acceptable to each of the involved parties.

In the automated multi-agent setting, the work described in (Zlotkin & Rosenschein 1992) has brought a new and revolutionary idea, based on concepts from Game Theory. It proves that certain problems from the class called Task Oriented Domains can be solved by *truth incentive* protocols¹. A protocol is truth incentive if any participant cannot gain more than by telling the whole truth about its problem. Additional problems were shown to allow truth incentive protocols and the best known examples are the one item auctions. They can be solved with the Vickrey protocol (Vickrey 1961). An extension of this protocol, Generalized Vickrey Auctions (GVA) (Varian 1995), has also been proposed for multiple-items auctions, namely auctions where individual pricing for items is different from grouped pricing. Truth incentive protocols naturally allow automatic centralized resolution and this is a big success of AI in general. Unfortunately, even if the GVA protocol (Varian 1995) guarantees a certain degree of equity for many multiple-items auctions, it is not always truth incentive (Yokoo *et al.* 2000a). The outcome for this complication, illustrated in (Yokoo *et al.* 2000b), is that with public constraints, the social welfare is sub-optimally managed. General auctions as well as other types of negotiations may be truth incentive even if resources and parties are involved in other known negotiations. However, if there exist *unknown connections* with future negotiations, revealing the truth presents a risk for involved parties.

Definition C.1 *The unknown connections of a given problem P consist of all future negotiations for which not all details are known and that share resources with P.*

¹Also referred as *incentive compatible mechanisms*.

In particular, truth incentive-ness is penalized by the following property, related to the theorem 7.1 presented in (Sandholm 1996):

Property C.1 *If a particular constraint on a resource x of an agent can ever be involved in an unknown future problem that allows no truth incentive protocol, then no truth incentive protocol can be safely used for any problem requiring to reveal x .*

This property does not mean that no truth incentive mechanism exists for the known part of the problem. Rather it states that involved parties might prefer not to reveal their constraints due to external unknown conditions. We therefore introduce the next definition.

Definition C.2 (Globally truth incentive) *Let P be a problem allowing a truth incentive mechanism M . M is globally truth incentive if P does not have any unknown connection. The corresponding truth incentive mechanism is then globally truth incentive.*

Example 3.40 *Let us imagine that a multi-provider bandwidth reservation problem $P1$ cannot allow truth incentive protocols with respect to the structure of the internal networks (e.g. by considering that due to the mechanisms currently used on the corresponding market, revealing this structure may completely destroy the competitiveness of their owner).*

Let us imagine that the negotiation mechanism $M2$ for the problem $P2$ of buying cables for the providers needs the structures of the networks of the participants. Let us imagine that $M2$ is truth incentive with respect to these structures of the internal networks.² Therefore, if the auctioneer of the cable negotiation cannot be a trusted party for the problem $P1$, then the cable negotiation problem $P2$ cannot allow for globally truth incentive mechanisms, and $M2$ is not globally truth incentive.

This property has implications in many problems. Making abstraction of it, even if optimal in the present, may be less good in the future. For simplicity, in the remaining part of this annex we refer problems that do not allow global truth incentive mechanisms due to unknown connections as being non-truth incentive problems.

For this kind of problems and for problems with false name bids, English Auctions are preferred to GVA since they do not require the agents to reveal everything.

C.2 Problem Statement

The English Auction negotiation mechanism is a good candidate for solving non-globally truth incentive problems since it offers a certain degree of privacy. For example, an agent may win the auction without revealing the highest price it can pay. Contrary to GVA, English Auctions are inherently distributed. The one item English Auctions mechanism is well understood and widely used in practice. Due to the complexity of the English Auctions for multiple-items auctions, GVA has been the most used solving mechanism even when it leads to less suitable solutions.

C.2.1 Fairness

The quality of a negotiation protocol mainly depends on its ability to compute fair solutions. In the following we give a set of definitions for characterizing solutions in problems with hidden costs. These definitions mainly adapt the commonly used ones to our situation.

Definition C.3 (Initiator) *The initiators in GEA are agents. They perform functions typically attributed to auctioneers in English Auctions: they offer deals to a set of waiting agents and collect bids, choosing the ones that increase their revenue.*

The cost of a solution is given by the sum of the costs of the agents. Note that since a price requested by an agent is treated as a negative cost, the sum of all the costs paid by all agents in

²Actually it may not be truth incentive when the reservation price of the auctioneer for future auctions can be changed.

a set of agents, A , is equal to the sum paid by agents in A to factors outside A . Any negotiation is started by a subset of A called *initiators*. We assume that the initiators are self-interested. The sum of the costs paid by initiators to some agent A_i is the price of the solution towards A_i .

Definition C.4 (Solution Cost) *The cost of a solution is given by the sum of the prices asked by non-initiator agents to initiators for agreeing on the deals composing the solution.*

Definition C.5 (Utility) *The utility of an agent is defined as the difference between the price it asks and the cost it pays for the chosen alternative.*

A rational agent prefers to offer alternatives that increase its utility³. Therefore, even if the utilities are hidden, it is beneficial for agents to reveal the order of their preferences, whenever Pareto-optimality is a concern.

Definition C.6 (Pareto-optimal solution) *A solution is Pareto-optimal if any other solution is either equally preferred for all agents, or worse for at least one agent, given the order defined by the utilities of each agent on solutions.*

In our case, the utilities are considered secret, but agents can reveal preferences for alternatives, and moreover, the order induced on alternatives by preferences can be considered to be the same as the order induced by utilities. We call Declared-Pareto-optimal solution, a Pareto optimal solution computed for the DisCSP (prices and preferences) declared by the agents.

Definition C.7 (Declared-Pareto-optimal) *A solution is Declared-Pareto-optimal if any other solution is either equally preferred for all agents, or worse for at least one agent, given the order defined by prices and declared preferences.*

When preferences would not be considered, any solution minimizing the solution cost is Declared-Pareto-optimal. Taking into account preferences, some of the solutions minimizing the solution cost may no longer be Declared-Pareto-optimal.

Definition C.8 (Estimated Social Welfare) *An estimated social welfare solution (ESW) is a Declared-Pareto-optimal solution (this implies that it has minimal Solution Cost).*

Guaranteeing that a solution is ESW is possible with complete search techniques. We also want the ESW solution to be chosen impartially (fairness).

Definition C.9 (Fairness) *When several ESWs are candidate, fairness consists in giving them equal probability to be chosen.*

C.2.2 Real Social Welfare

Mechanisms can try to reach social welfare solutions.

Definition C.10 (Social Welfare Solution) *A social welfare solution (SW) for a problem involving a set of agents A is defined as a solution maximizing the sum of all the utilities of the agents in A .*

GEA cannot guarantee SW solutions since the utilities of the agents are considered secret. However, for some problems, GEA has chances to approach a SW solution.

Definition C.11 (Equivalent Solutions) *A problem with equivalent solutions is a problem where the difference between the quality (worth (Zlotkin & Rosenschein 1992)) of its solutions is equal to the difference between the cost of achieving the respective solutions (the solutions are equally good).*

³Alternatively, the notion of *worth* (Zlotkin & Rosenschein 1992) can be similarly used.

It is worth mentioning that for problems with Equivalent Solutions, an ESW gives the best possible estimation of the real Social Welfare (SW). This is the case of a bandwidth allocation problem where any two paths in the network are equally good as long as it has the required bandwidth and quality of service.

Proposition C.2 *For problems with equivalent solutions, a social welfare solution (SW) for a set of agents A can be computed as a solution minimizing the sum of all the costs paid by all the agents in A for agreeing on the alternatives composing the solution, whenever the utilities of the agents are proportional to the prices they ask.*

Proof. If the utilities are proportional with prices, then getting a SW solution means maximizing the sum of utilities which is proportional to the sum of received prices (modeled as minus of paid costs). This is equivalent to minimizing paid costs. \square

Example 3.41 *Consider an auction for buying a car and two bidders that offer cars with equal worth. Consider the price requested for each car to be proportional to the cost of the corresponding bidder for producing and offering that car. Then, the choice of the cheapest car is the social welfare solution since the maximal worth was obtained with the lowest overall cost, and therefore the sum of utilities is maximized.*

While we can never guarantee that *the utilities of the agents are proportional to the prices they ask*, this may often be an acceptable assumption. For problems that do not allow for truth incentive mechanisms, GEA can get closer to SW solutions than other approximative truth incentive mechanisms do (Yokoo *et al.* 2000a).

C.3 Related Work

Researchers have already related *negotiation* and *Distributed CSPs* from both sides. On one side, the *negotiation* is seen as a technique for solving distributed CSPs. The authors of (Lander & Lesser 1993) propose a min-conflict heuristic technique called *negotiation search* as a means of converging towards a solution in a distributed problem with heterogeneous components. On the other side agents have also been proposed for solving by negotiation over-constrained resource allocation problems in (Conry *et al.* 1991; Khedro & Genesereth 1994). Frameworks for over-constrained distributed problems with public constraints are presented in (Hirayama & Yokoo 1997). Our approach shares common concepts with the framework proposed in (Sathi & Fox 1989) for resource allocation. We present in (Silaghi *et al.* 2001k) an algorithm that can be used for the framework that is used here to model GEAs. Simultaneously, (Vauvert & Fallah-Seghrouchni 2001) discusses an approach for coalition formation, which cannot find all solutions and requires exponential space, but has an interesting and complementary approach to fairness. Other algorithms for DisCSPs are described in (Collin *et al.* 1991b; Solotorevsky *et al.* 1996a; Armstrong & Durfee 1997; Havens 1997; Yokoo *et al.* 1998; Silaghi *et al.* 2000a; Bessi ere *et al.* 2001). An overview of known types of auctions was given in (Sandholm & Suri 2000).

C.4 The Negotiation Protocol

A negotiation is viewed as a multi-criteria optimization problem where the agents have to find a solution maximizing their utilities while respecting their constraint on resources. Auctions are special case of negotiations where the negotiation ends when a subset of the agents (auctioneers) cannot improve any longer their utility. In GEA such problems are solved by iterative improvement of ESW solutions according to the following protocol.

- a1 Compute the best solutions (ESW) satisfying the constraints so far imposed by the agents and retain one of them.
- a2 If any solution was found at a1, publish the ESW as an any-time solution.
- a3 If any agent wants to relax the constraints it imposes, go to a1.

a4 If any solution was found at *a1*, return the estimated ESW and stop.

a5 Return failure and stop.

When the algorithm used at step *a1* is complete, the solution of a GEA is a global optima (i.e. no better solution can be constructed by the agents).

If the prices are modified with a minimal increment and the set of alternatives is finite, the previous protocol is safe to converge in finite time as long as the agents are stable in the order on their preferences and commit to their agreements (monotonicity in finite domains).

A solution S of a distributed problem may not need the agreement of some particular agent A_i . In that case we say that A_i is *inactive* for S . Conversely, we say that A_i is *active* for S if its agreement is necessary for choosing S . This provides a means to model a facet of *competition* useful for ameliorating the ESW. An agent, *inactive* for the current solution, may indeed want to make concessions to become *active*.

C.5 Summary

We present an approach to negotiation for problems where no globally truth incentive mechanism is available. The importance of the GAP framework is acknowledged by a recent work which stresses the importance of constraints in negotiations, feature that it offers naturally.

In conjunction with appropriate reordering heuristics in search algorithms, Dynamic DisCSPs can be used in general negotiation problems (e.g. GEA) to provide fair environments.

The presented framework inherits from Constraint Reasoning generality and flexibility in modeling. Among the problems that remain to be solved we mention the legal aspects of the deals, the development of rational strategies for agents enforcing local privacy, and computing fair probabilities in choosing ESWs.

