A Negotiation Protocol for Agents with Nonlinear Utility Functions

Takayuki Ito
Division of Computer Science and Engineering,
Master Course of Techno-Business Administration,
Nagoya Institute of Technology,
Gokiso, Showaku, Nagoya 466-8555, Japan
itota@ics.nitech.ac.jp
http://www-toralab.ics.nitech.ac.jp/~itota/

Mark Klein and **Hiromitsu Hattori**
Center for Coordination Science,
Sloan School of Management,
Massachusetts Institute of Technology,
3 Cambridge Center, NE20-336, Cambridge, MA 02142, USA
{m_klein, hatto}@mit.edu

Introduction

Multi-issue negotiation protocols represent a important field of study since negotiation problems in the real world are often complex ones involving multiple issues. While there has been a lot of previous work in this area (Faratin, Sierra, & Jenning (2002) Soh & Li (2004) Fatima, Wooldridge, & Jennings (2004)) these efforts have, to date, dealt almost exclusively with simple negotiations involving **independent** issues, and therefore linear (single optimum) utility functions. Many real-world negotiation problems, however, involve **interdependent** issues. When designers work together to design a car, for example, the value of a given carburetor is highly dependent on which engine is chosen. The addition of such interdependencies greatly complicates the agent's utility functions, making them nonlinear, with multiple optima. Negotiation mechanisms that are well-suited for linear utility functions, unfortunately, fare poorly when applied to nonlinear problems (Klein *et al*. (2003)).

We propose a multiple-issue negotiation protocol suited for agents with such nonlinear utility functions. Agents generate bids by sampling their own utility functions to find local optima, and then using constraint-based bids to compactly describe regions that have large utility values for that agent. These techniques make bid generation computationally tractable even in large (e.g. 10^{10} contracts) utility spaces. A mediator then finds a combination of bids that maximizes social welfare. Our experimental results show that our method substantially outperforms negotiation methods designed for linear utility functions.

Negotiation with Nonlinear Utilities

Graph (a) in Figure 1 shows an example of a binary constraint between issues 1 and 2. This constraint has a value of 55, and holds if the value for issue 1 is in the range $[3, 7]$ and the value for issue 2 is in the range $[4, 6]$. Every agent has its' own, typically unique, set of constraints.

An agent's utility for a contract $\vec{s}$ is defined as $u_i(\vec{s}) = \sum_{c_k \in C, \vec{s} \in x(c_k)} w_i(c_k, \vec{s})$, where $x(c_k)$ is a set of possible contracts (solutions) of c_k. A constraint c_k has value $w_i(c_k, \vec{s})$ if and only if it is satisfied by contract $\vec{s}$. This expression produces a "bumpy" nonlinear utility space, with

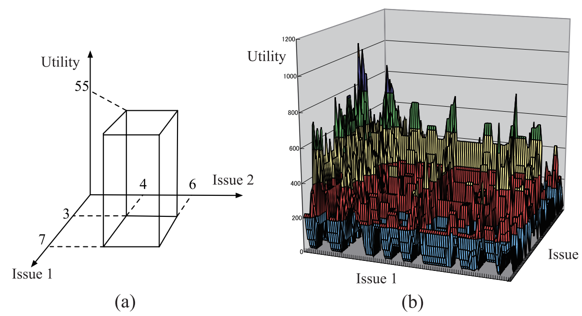

Figure 1: Examples

high points where many constraints are satisfied, and lower regions where few or no constraints are satisfied. This represents a crucial departure from previous efforts on multi-issue negotiation, where contract utility is calculated as the weighted sum of the utilities for individual issues, producing utility functions shaped like flat hyper-planes with a single optimum.

Graph (b) in Figure 1 shows an example of a nonlinear utility space.

We assume, as is common in negotiation contexts, that agents do not share their utility functions with each other, in order to preserve a competitive edge. It will generally be the case, in fact, that agents do not fully know their **own** utility functions, because they are simply too large. If we have 10 issues with 10 possible values per issue, for example, this produces a space of 10^{10} (10 billion) possible contracts, too many to evaluate exhaustively. Agents must thus operate in a highly uncertain environment.

The Proposed Negotiation Protocol

Our auction-based negotiation protocol consists of the following four steps: **(Step 1 : Sampling)** Each agent samples its utility space in order to find high-utility contract regions. A fixed number of samples are taken from a range of random points, drawing from a uniform distribution. **(Step 2 : Adjusting)** There is no guarantee, of course, that a given sample will lie on a locally optimal contract. Each agent, therefore, uses a nonlinear optimizer based on simulated annealing to try to find the local optimum in its neighborhood. **(Step 3 : Bidding)** For each contract $\vec{s}$ found by adjusted sampling, an agent evaluates its utility. If that utility is larger than the reservation value δ, then the agent defines a bid that covers all the contracts in the region which have that utility value. This is easy to do: the agent need merely find the intersection of all the constraints satisfied by that $\vec{s}$. **(Step 4 : Winner Determination)** The mediator identifies the final contract by finding all the combinations of bids, one from each agent, that are mutually consistent, i.e., that specify overlapping contract regions. If there is more than one such overlap, the mediator selects the one with the highest summed bid value (and thus, assuming truthful bidding, the highest social welfare) . The mediator employs breadth-first search with branch cutting to find social-welfare-maximizing overlaps. It is easy to show that, in theory, this approach can be guaranteed to find optimal contracts.

Experiments

Setting

We conducted several experiments to evaluate the effectiveness and scalability of our approach. In each experiment, we ran 100 negotiations between agents with randomly generated utility functions. For each run, we applied an optimizer to the sum of all the agents' utility functions to find the contract with the highest possible social welfare. This value was used to assess the optimality of the negotiation protocols. When possible, we used exhaustive search (EX) to find the optimum contract, but when this became intractable we switched to simulated annealing (SA).

We compared two negotiation protocols: hill-climbing (HC), and our proposed protocol with random sampling (AR). The HC approach implements a mediated single-text negotiation protocol based on hill-climbing. In our implementation, every possible single-issue change was proposed once, so the HC protocol requires only $domainsize \times numberofissues$ iterations for each negotiation (e.g. 100 steps for the 10 issue case with domain $[0, 9]$). We selected this protocol as a comparison case because it represents a typical example of the negotiation protocols that have been applied successfully, in previous research efforts, to linear utility spaces. The details of the parameters are omitted in the interest of space.

Results

Let us first consider the linear utility function (**independent** issue) case that has been the focus of almost all previous work on multi-issue negotiation. As we can see (Table 1) even the simple HC protocol produces essentially optimal results for a wide range of contract space dimensions. Since the issues are independent, the mediator can optimize over each issue independently, first finding the most-favored value for issue 1, then for issue 2, and so on. Once every issue has been optimized, the final contract will generally be very close to optimal.

The story changes dramatically, however, when we move to a nonlinear utility function (**interdependent** issue) case (Figure 2). In this context, HC produces highly suboptimal results, averaging only 40% of optimal, for example, for the 10 issue case. Why does this happen? Since every agent has a "bumpy" (multi-optimum) utility function, the HC mediator's search for better contracts grinds to a halt as soon as any of the agents reach a local optimum, even if a contract which is better for all agents exists somewhere else in the contract space. The AR protocol, by contrast, achieves often near-optimal outcomes for higher-order problems. Since agents using the AR protocol generate bids that cover multiple optima in their utility spaces, our chances of finding contracts that are favored by all agents is greatly increased.

Table 1: Optimality with **linear** utility function, 4 agents

Issues	1	2	3	4	5
HC	0.973	0.991	0.998	0.989	0.986
Issues	6	7	8	9	10
HC	0.987	0.986	0.996	0.988	0.991

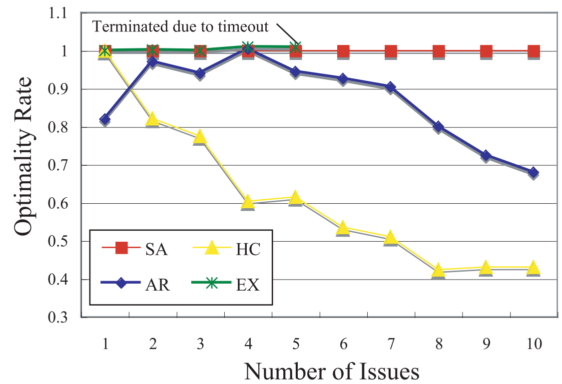

Figure 2: Social welfare with **nonlinear** utility functions

Conclusions

In this paper, we have proposed a novel auction-based protocol designed for the important challenge of negotiation with multiple interdependent issues and thus nonlinear utility functions. Our experimental results show that our method substantially outperforms protocols that have been applied successfully in linear domains.

References

Faratin, P.; Sierra, C.; and Jenning, N. R. 2002. Using similarity criteria to make issue trade-offs in automated negotiations. In *Artificial Intelligence*, 142:205–237.

Fatima, S.; Wooldridge, M.; and Jennings, N. R. 2004. Optimal negotiation of multiple issues in incomplete information settings. In *Proc. of Autonomous Agents and Multi-Agent Systems (AAMAS2004)*.

Klein, M.; Faratin, P.; Sayama, H.; and Bar-Yam, Y. 2003. Negotiating complex contracts. *Group Decision and Negotiation* 12(2):58–73.

Soh, L.-K., and Li, X. 2004. Adaptive, confidence-based multiagent negotiation strategy. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS2004)*.