MAXIMAL CLASSES OF UTILITY FUNCTIONS FOR EFFICIENT ONE-TO-ONE NEGOTIATION: EXTENDED ABSTRACT

Yann Chevaleyre [a] Ulle Endriss [b] Nicolas Maudet [a]

[a] *LAMSADE, University of Paris-Dauphine, France*
[b] *ILLC, University of Amsterdam, The Netherlands*

Abstract

We investigate the properties of an abstract negotiation framework where agents autonomously negotiate over allocations of indivisible resources. In this framework, reaching a socially optimal allocation may require very complex multilateral deals. Therefore, we are interested in identifying classes of utility functions such that any negotiation conducted by means of deals involving only a single resource at at time is bound to converge to an optimal allocation whenever all agents model their preferences using these functions. We show that the class of modular utility functions is not only sufficient but also maximal in this sense.

The full version of this paper has appeared in the Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI-2005).

1 Multiagent Resource Allocation

The problem of *multiagent resource allocation* has recently received a lot of attention in the Artificial Intelligence community [2]. Much work has focussed on combinatorial auctions, where the allocation procedure is centralised [3]. A different perspective is taken when one assumes that the allocation process is truly *distributed*, in the sense that agents autonomously negotiate over the bundles of resources they hold. This assumption is justified in many applications where no central authority can be relied upon to decide on the final allocation. In this case, the system designer will typically seek to set up the system in such way that it guarantees certain desirable properties, without directly interfering in the negotiation process itself.

One such desirable property would be to be able to guarantee that the system will converge towards an allocation of resources that is *socially optimal*. There is a whole range of notions of "social optimality" that would be applicable in such a scenario [1]; here we concentrate on maximising the *utilitarian social welfare* of an allocation, *i.e.* on maximising the sum of individual utilities. We assume that all agents are *rational* and *myopic* in the sense of never accepting a deal that would result in a negative payoff. It is possible to show that any deal that increases utilitarian social welfare does result in a positive payoff for all the agents involved (provided agents can use monetary side payments) and *vice versa* [4]. However, to be able to ensure convergence to an optimal allocation, it can be necessary to implement complex *multilateral* deals (between more than two agents) over any number of resources at a time [5]. This is typically not practical. It is therefore important to investigate under what circumstances convergence to an optimal allocation can be guaranteed by means of sequences of very simple deals.

It turns out that the structural complexity of our negotiation framework largely stems from the fact that agents may use any kind of utility function to model their preferences over alternative bundles of resources. If we introduce suitable restrictions and only allow for utility functions belonging to certain classes of functions to be used, then a simpler negotiation regime may suffice. Here we are interested in classes of utility functions that *permit 1-deal negotiation*, *i.e.* where negotiation by means of mutually beneficial deals over a single resource at a time (and thereby only involving two agents each) is guaranteed to converge to an allocation with maximal utilitarian social welfare.

2 Modular Functions are Sufficient

Our first result establishes that the class of *modular* utility functions is *sufficient* to permit 1-deal negotiation. Recall that a utility function u is called modular iff we have $u(R_1 \cup R_2) = u(R_1) + u(R_2) - u(R_1 \cap R_2)$ for all bundles R_1, R_2. This is equivalent to saying that u can be represented as the sum of utilities assigned to single items (possibly with a non-zero base utility assigned to the empty bundle). That is, in modular domains we cannot model any synergies between different resources. It is therefore not very surprising that negotiation over one item at a time is sufficient to reach an optimal allocation.

3 Modular Functions are not Necessary

Our next result shows that modularity of agents' utility functions, while being sufficient, is not a *necessary* condition for permitting 1-deal negotiation. More surprisingly, we have shown that there can be *no* class of utility functions that would be both sufficient and necessary in this sense. In other words, there is no unique largest class of utility functions such that 1-deal negotiation can guarantee outcomes with maximal social welfare whenever all agents use utility function belonging to that class. The proof of this result is in fact very simple; it proceeds by identifying two classes of functions, each of which is sufficient but the union of which is not.

4 Modular Functions are Maximal

Our main result establishes the surprising fact that the class of modular utility functions is *maximal* in the sense that no class of utility functions strictly including the modular functions would still be sufficient for 1-deal negotiation. The proof is constructive. Given a non-modular utility function for one agent, it shows how to compose modular utility function for all the other agents such that there is an initial allocation from which no optimal allocation can be reached by means of (mutually beneficial) 1-deals alone.

We consider this not only a surprising result, but also a useful characterisation of negotiation domains that can be handled reliably using simple negotiation protocols, catering only for *Contract Net*-like deals over single items between pairs of agents rather than the full range of multilateral deals foreseen in our abstract multiagent resource allocation framework. Such theoretical results affect both the design of agents and of negotiation mechanisms. For instance, if a given mechanism can only handle 1-deals, then it may be inappropriate to design myopic agents with very rich preference structures to use such a mechanism.

References

[1] K. J. Arrow, A. K. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare*. North-Holland, 2002.

[2] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. A. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 2006. To appear.

[3] P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2006. To appear.

[4] U. Endriss, N. Maudet, F. Sadri, and F. Toni. On optimal outcomes of negotiations over resources. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2003)*. ACM Press, 2003.

[5] T. W. Sandholm. Contract types for satisficing task allocation: I Theoretical results. In *Proceedings of the AAAI Spring Symposium: Satisficing Models*, 1998.