Negotiation Exploiting Reasoning by Projections

Toni Mancini

Dipartimento di Informatica, Università di Roma "La Sapienza", Italy
tmancini@di.uniroma1.it

Abstract. We present a framework that allows two self-motivated distributed agents to perform, in an efficient way, negotiations aiming at achieving mutually satisfactory agreements, when privacy of information is an issue, and no central authority could be used. In particular, each agent has her own constraints to satisfy, as well as her own utility function. Such issues are kept private, and cannot be disclosed to the counterpart. Negotiation is hence carried out by exchanging proposals and by performing sophisticated forms of reasoning on the remote agent's offers, by trying to infer some characteristics of the counterpart, in order to achieve efficient process convergence.

1 Introduction

Automated negotiation among rational agents is an important topic in Distributed Artificial Intelligence, being necessary in several domains, e.g., distributed resource allocation [3], distributed scheduling [11], e-business [9] and, in general, in applications in which: *(i)* no single agent can achieve her own goals without interaction with the others (or she is expected to achieve more utility with interaction), and *(ii)* constraints of various kinds (e.g., security or privacy) forbid the parties to communicate their desiderata to others (the counterpart or a trusted authority), hence traditional centralized approaches, like Mathematical or Constraint Programming cannot be used. In the literature, various protocols and algorithms for negotiation have been proposed (cf., e.g., [13, 9, 12, 8, 10]), all of which can be classified according to several factors, like the negotiation objects, the agents' decision making models, the degree of cooperation and the level of privacy about each agents' constraints and preferences, or the communication and computation costs.

In this paper, we focus on negotiation between two self-motivated, competitive agents, which aim to find a mutually satisfactory agreement without disclosing their private information (constraints and preferences). Differently from much work existing in the literature, *we explicitly deal with the fact that, in many scenarios, agents face with unknown counterparts, for which no strong assumptions (even probabilistic ones) can be made*: understanding their characteristics and behavior is part of the job of the agent. One of the major prices of this setting is that even the termination of the negotiation process is not guaranteed. To this end, we focus on the local agent only (the only one upon which we have control), which aims at *maximizing* her own private utility function (*competitiveness*) while *guaranteeing termination and efficiency* of the negotiation process. In order to do so, she behaves in a sort of *non-obstructionist way*, and acts very carefully in order to avoid the process to go thrashing and last indefinitely, relying on initial expectations about the reasoning capabilities of the counterpart, but

Y. Demazeau et al. (Eds.): 7th International Conference on PAAMS'09, AISC 55, pp. 329–338.
springerlink.com

being ready to revise them, by means of a continuous monitoring of her actions, in case a contradicting behavior is observed.

The starting point of our research is the (cooperative) framework described in [1], in which the task of finding a mutually satisfactory agreement is particularly efficient. In such a framework, in which privacy of information is a major concern, negotiation proceeds by exchanging *proposals*. The other agent can either accept the deal or propose a counter-offer. An agent is also able to *reason* on the other agent's proposals. This form of reasoning is called *reasoning by projections*.

In this paper we extend the framework above in the following directions (Sec. 2): *(i)* we provide the agent with the ability to *dynamically estimate and reason on the quality of reasoning made by the counterpart*, and to consequently adapt her behavior *exploiting different heuristics* suitable for the different cases; *(ii)* we provide the agent with the possibility of *defining her own* (private) *utility function* to maximize. Given that the presence of utility functions could make the agent interested in refusing acceptable deals in order to pursue better ones, guarantee of termination can easily be lost. We equip the agent with capabilities that allow her to *still guarantee convergence*, by reasoning on the impact that her refusals have on the counterpart potential reasoning. A *full functional implementation* is briefly described and experimentally evaluated (Sec. 3, details in Appendix[1] due to space reasons).

2 The Framework

When dealing with negotiation, there are several aspects that must be taken into account. We address the reader to [1] for a discussion on the most important ones. Here, we only describe the assumptions behind our theoretical framework.

Negotiation framework. When two agents start a negotiation, they already agree on the relevant *variables* (or issues). Following [1], we assume that for each negotiation there is a finite list of *real variables* that are involved. Hence, negotiation spaces can be regarded as multi-dimensional real vector spaces. Agents have their own private *feasibility regions* (R_{loc} and R_{rem}). This setting is different from what is often assumed in the literature, since agents don't even know (or probabilistically estimate) possible counterpart's *types*, *bounds* or *most preferred values* for the variables (in other words, it is not a *split-the-pie* game [5] although with incomplete information [10]). Negotiation proceeds with agents alternatively exchanging *proposals* (aka *deals*), as single points in the space of variable assignments. The counterpart can either accept the deal or decline it, by sending a *counter-offer*.

Characteristics of the local agent. We assume that R_{loc} is *convex* (i.e., all points between two acceptable points are acceptable as well), and more particularly *limited* and defined by means of *linear constraints*, hence it is a *bounded polyhedron*. We also assume that R_{loc} is *stable* during time. Convexity and stability, which play a key role in our approach, are very common in many scenarios of practical utility (cf., e.g., [1] and Appendix A.2 for an example). Admissible proposals are in general not equally worth for the agent, that may have her own *utility function*, preferring some solutions to other ones. We assume that the utility function, to be *maximized* by the local agent, is *linear*.

[1] Available online at `http://www.dis.uniroma1.it/~tmancini`