Chapter XLVIII
An Annealing Protocol for Negotiating Complex Contracts

Mark Klein
Massachusetts Institute of Technolocy, USA

Peyman Faratin
Massachusetts Institute of Technolocy, USA

Hiroki Sayama
University of Electro-Communications, Japan

Yaneer Bar-Yam
New England Complex Systems Institute, USA

ABSTRACT

Work to date on negotiation protocols has focused almost exclusively on defining contracts consisting of one or a few independent issues and a relatively small number of possible contracts. Many real-world contracts, by contrast, are much more complex, consisting of multiple interdependent issues and intractably large contract spaces. This chapter describes a simulated annealing-based approach appropriate for negotiating such complex contracts that achieves near-optimal social welfare for negotiations with binary issue dependencies.

INTRODUCTION

Work to date on negotiation protocols has focused almost exclusively on negotiating what we can call *simple* contracts—that is, contracts consisting of one or a few independent issues (Faratin, Sierra, & Jennings, 2000; Ehtamo, Ketteunen, & Hamalainen, 2001; Fisher, Ury, & Patton, 1991; Raiffa, 1982). These protocols work, in general, via the iterative exchange of proposals and counterproposals. An agent starts with a contract that is optimal for that agent and makes concessions, in each subsequent proposal, until either an agreement is reached or the negotiation is abandoned because the utility of the latest proposal has fallen below the agents' reservation value (see Figure 1).

This is a perfectly reasonable approach for simple contracts. Since issues are independent,

Figure 1. The proposal exchange model of negotiation, applied to a simple contract. Each point on the X axis represents a possible contract. The Y axis represents the utility of a contract to each agent.

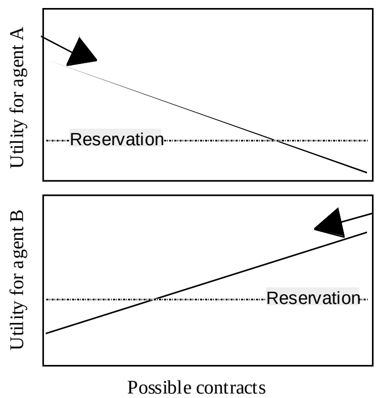

the utility of a contract for each agent can be calculated as the weighted sum of the utility for each issue. The utility function for each agent is thus a simple one, with a single optimum and a monotonic drop-off in utility as the contract diverges from that ideal. Simple contract negotiations thus typically progress as in Figure 2.

As we can see, the proposals from each agent start at their own ideal, and then track the Pareto frontier until they meet in the middle with an optimal agreement. This happens because, with linear utility functions, it is easy for an agent to identify the proposal that represents the minimal concession: the contract that is minimally worse than the current one is "next" to the current one in the contract space and can be found by moving in the direction with the smallest aggregate utility slope. The simplicity of the utility functions, moreover, makes it feasible for agents to infer enough about their opponents that they can identify concessions that are attractive to each other, resulting in relatively quick negotiations.

Real-world contracts, by contrast, are generally much more complex, consisting of a large number of inter-dependent issues. A typical contract may have tens or even hundreds of

Figure 2. A typical negotiation for a simple contract. The contract consisted in this case of 40 binary issues. Each agent was required to reduce the Hamming distance (number of issues with different values) between successive proposals until an agreement was reached. The Pareto frontier was estimated by applying an annealing optimizer to differently weighted sums of the two agents' utility functions.