Center for eBusiness

RESEARCH BRIEF

Volume IV Number 2 May 2004

PROTOCOLS FOR NEGOTIATING COMPLEX CONTRACTS[1]

Mark Klein, *Principal Research Associate, MIT Sloan School of Management*

Peyman Faratin, *Senior Postdoctoral Associate,*

MIT Computer Science & Artificial Intelligence Lab

Hiroki Sayama, *University of Electro-Communications*

Yaneer Bar-Yam, *New England Complex Systems Institute*

Introduction

Work to date on negotiation protocols has focused on negotiating what we can call ‘simple’ contracts, i.e. contracts consisting of one or a few independent issues. These protocols work via the iterative exchange of proposals and counter-proposals. Since issues are independent, 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. The simplicity of the utility functions 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. Even with only 50 issues and two alternatives per issue, we encounter a search space of roughly 10^15 possible contracts, too large to be explored exhaustively. The value of one issue selection to an agent, moreover, will often depend on the selection made for another issue. Such issue interdependencies lead to nonlinear utility functions with multiple local optima.

In such contexts, an agent finding its own ideal contract becomes a nonlinear optimization problem. Simply conceding toward the other agents’ proposals can result in the agents missing contracts that would be superior from both their perspectives. Standard negotiation techniques thus typically produce the following behavior when applied to complex contract negotiation:

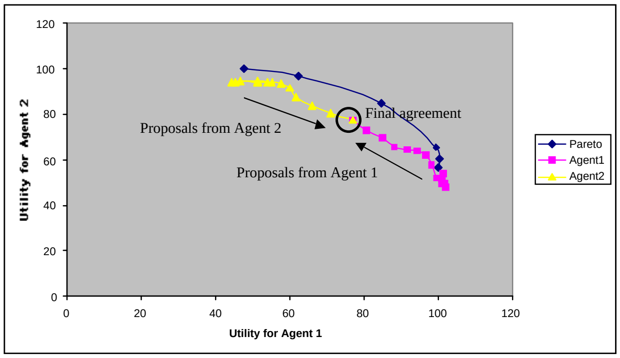

Figure 1. The utilities for the proposals made in a typical complex contract negotiation.

The minimal concession protocol that works optimally for simple contracts produces outcomes, for complex contracts, that are substantially sub-optimal.

In our experiments, for example, the final contracts’ averaged 88% of optimal. This is a substantial decrement when you consider that the utility functions we used for each agent were, individually, quite easy to optimize: a simple steepest ascent search averaged final utility values roughly 97% of those reached by a nonlinear optimization algorithm.

Complex contracts require different negotiation techniques which allow agents to find ‘win-win’ contracts in intractable multi-optima search spaces in a reasonable amount of time.

A Negotiation Algorithm for Complex Contracts

Our starting point was mediated single text negotiation, a standard approach for dealing with complex negotiations in human settings [2]. In this process, a mediator proposes a contract that is then critiqued by the parties in the negotiation. A new, hopefully better proposal is then generated by the mediator based on these responses. This process continues, generating successively better contracts, until some agreed-upon stopping point (e.g. the reservation utility value is met or exceeded for both parties). We can visualize this process as follows:

---

[1] Appears in: Klein, M., et al., Protocols for Negotiating Complex Contracts. IEEE Intelligent Systems, 2003. 18(6): p. 32 - 38.

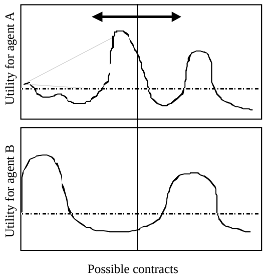

Figure 2: Single text negotiation. The vertical line represents the current proposed contract, and subsequent proposals move that line in the contract space.

In our approach, the mediator proposes a contract that is initially generated randomly. Each agent then votes to accept or reject the contract. If both vote to accept, the mediator mutates the contract (by randomly flipping one of the issue values) and the process is repeated. If one or both agents vote to reject, a mutation of the most recent mutually accepted contract is proposed instead. The process is continued for a fixed number of proposals. Note that this approach can straightforwardly be extended to a N-party (i.e. multi-lateral) negotiation, since we can have any number of parties voting on the contracts.

We defined two kinds of agents: 'hill-climbers' and 'annealers'. The hill-climbers use a very simple decision function: they accept a mutated contract only if it's utility to them is greater than that of the last contract both agents accepted. Annealers are more complicated. Each annealer has a virtual 'temperature' T, such that it will accept contracts worse than last accepted one with the probability:

$$P(\text{accept}) = \min(1, e^{-\Delta U/T})$$

where ΔU is the utility change between the contracts. In other words, the higher the virtual temperature, and the smaller the utility decrement, the greater the probability that the inferior contract will be accepted. The virtual temperature of an annealer gradually declines over time so eventually it becomes indistinguishable from a hill-climber. Annealing has proven effective in single-agent optimization, because it can travel through utility valleys on the way to higher optima [1]. This suggests that annealers can be more successful than hill-climbers in finding good negotiation outcomes.

The Prisoner's Dilemma

Negotiations with annealing agents did indeed result in substantially superior final contract utilities, but as the payoff table below shows, there is a catch:

	Agent 2 hill-climbs	Agent 2 anneals
Agent 1 hill-climbs	.86 .73/.74	.86 .99/.51
Agent 1 anneals	.86 .51/.99	.98 .84/.84

Table 1: The optimality of the negotiation outcomes for different pairings of annealing and hill-climbing agents. The top value in each cell represents how close the social welfare value of the final contract is to optimal. The pair of values below it represent how close the final contract is to the optimum for the Agent 1 and Agent 2, respectively.

As expected, paired hill-climbers do relatively poorly (since it is difficult to find many contracts that represent an improvement for both parties) while paired annealers do very well (because they are willing to accept individually worse contracts earlier on so they can find win-win contracts later).

If one agent is a hill-climber and the other is an annealer, however, the hill-climber does extremely well but the annealer fares correspondingly poorly. This pattern can be understood as follows. When an annealer is at a high virtual temperature, it becomes a chronic conceder, accepting almost anything beneficial or not, and thereby pays a "conceder's penalty". The hill-climber 'drags' the annealer towards its own local optimum, which is not very likely to also be optimal for the annealer.

This reveals a dilemma. In our case, even though annealing is a *socially* dominant strategy (i.e. annealing increases social welfare), annealing is not an *individually* dominant strategy. Hill-climbing is dominant, because no matter what strategy the other agent uses, it is better to be a hill-climber (Table 1). If all agents do this, however, then they forego the higher individual utilities they would get if they both annealed. Individual rationality thus drives the agents towards the strategy pairing with the *lowest individual and social welfare.* This is thus an instance of the prisoner's dilemma.

The Annealing Mediator

We were able to modify this protocol so that it avoids the prisoner's dilemma entirely. Rather than requiring that the negotiating agents anneal, and thereby expose themselves to the risk of being dragged into bad contracts, we moved the annealing into the mediator itself. The mediator is endowed with a time-decreasing willingness to follow up on contracts that one or both agents rejected (following the same inverse exponential regime as the annealing agents). Agents, who mark their accept and reject votes as being 'strong' or 'weak' to help the mediator search the contract space more intelligently, are free to remain hill-climbers and thus avoid making harmful concessions. The mediator, by virtue of being willing to provisionally pursue utility-decreasing contracts, can traverse valleys in the agents' utility functions and thereby lead the agents to win-win solutions.

Incentives for Truthful Voting

Any voting scheme introduces the potential for strategic non-truthful voting by the agents. Imagine that one of the agents always votes truthfully, while the other exaggerates so that its votes are always 'strong'. One might expect that this would bias negotiation outcomes to favor the exaggerator and this is in fact the case:

	Agent 2 exaggerates	Agent 2 tells truth
Agent 1 exaggerates	.92 .81/.81	.93 .93/.66
Agent 1 tells truth	.93 .66/.93	.99 .84/.84

Table 2: The optimality of the negotiation outcomes for truth-telling vs. exaggerating agents with a simple annealing mediator. An exaggeration strategy is individually incented, even though it results in outcomes with lower social welfare.

Even though exaggerating has substantial negative impact on social welfare, agents are individually incented to exaggerate, thus re-creating the prisoner's dilemma we encountered earlier. The underlying problem is simple: exaggerating agents are able to induce the mediator to accept all the proposals that are advantageous to them (if they are weakly rejected by the other agent), while preventing the other agent from doing the same. What we need, therefore, is an enhancement to the negotiation protocol that incents truthful voting, preserving equity and maximizing social welfare.

How can this be done? Let us define an 'override' as what happens when one agent influences the mediator to accept a proposal despite the objections of the other negotiator. The solution, we found, came from enforcing a rough running parity between the number of overrides given to each agent, so neither agent can get more than a given advantage. This way at least rough equity is maintained no matter when (or whether) either agent chooses to exaggerate. The results of this approach were as follows when the override disparity was limited to 3:

	Agent 2 exaggerates	Agent 2 tells truth
Agent 1 exaggerates	.91 .79/.79	.92 .78/.81
Agent 1 tells truth	.92 .81/.78	.98 .84/.84

Table 3: The optimality of the negotiation outcomes for truth-telling vs. exaggerating agents with parity-enforcing mediator. The parity-enforcing mediator makes truth-telling the rational strategy.

When we have truthful agents, we find that this approach achieves social welfare just slightly below that achieved by a simple annealing mediator, while offering a significantly ($p < 0.01$) higher payoff for truth-tellers than exaggerators. We found that the same pattern of results holds for a range of exaggeration strategies, including exaggerating all the time, exaggerating at random, or exaggerating just near the end of the negotiation. Truth-telling is thus both the individually dominant *and* socially most beneficial strategy.

Why does a truth-teller fare better than an exaggerator with this kind of mediator? One can think of this procedure as giving agents 'tokens' that they can use to 'purchase' advantageous overrides, with the constraint that both agents spend tokens at a roughly equal rate. The truthful agent spends its tokens almost exclusively on contracts that truly offer it a strong utility increase. The exaggerator, on the other hand, spends override tokens even when the utility increment it derives is relatively small. At the end of the day, the truthful agent has spend its tokens more wisely and to better effect.

The Unmediated Single Text Protocol

The protocol we have just considered suffers from the disadvantage of requiring a mediator. One issue concerns trust. Since the annealing mediator is empowered to selectively ignore agent votes, there is the risk that it may do so in a way that favors one agent over another (though the use of the parity-enforcing token mechanism does somewhat reduce the potential impact of this problem). Another issue concerns how quickly negotiations converge on a result. The annealing mediator generates new proposals by making random mutations to the last provisionally accepted contract, without taking into account any information about what contracts are preferable or even sensible. As a result, the mediator generates a very high proportion of rejected contracts, which is part of the reason why our experimental runs each involved so many (2500) proposals. The negotiating agents could imaginably provide the mediator with information about their utility functions so that the mediator is able to propose contracts more ‘intelligently’, but this is problematic for a number of reasons including the typical reluctance of self-interested agents to reveal their utility functions to a party that may or may not be worthy of their trust.

An effective unmediated version of the annealing protocol can, fortunately, be defined. It works as follows. Agents each start with a given number of tokens (2 each, in our experiments) and a mutually agreed-upon starting temperature T. A random contract is generated, and one of the negotiating agents is selected at random to propose a small (e.g. single-issue) variant thereof, presumably the variant that most increases the utility of the contract for that agent. The other agent then votes on the proposed variant. The proposals and votes both indicate the strength of their preference for the proposed contract using the scheme described above (i.e. strong reject, weak reject, weak accept, strong accept). The contract is provisionally accepted with probability

$P(\text{accept}) = \min(1, e^{-\Delta U/T})$

where the outcome is determined using the roll of a fair, mutually observable dice. If the decision to accept a proposal represents the over-ride of one agents’ reject vote, the winning agent needs to give one of it’s’ tokens to the over-ridden agent. An over-ride is not permitted if the agent has run out of tokens. The proposer and voter alternate roles thereafter until neither agent can identify any improvements to make to the last accepted contract. Agents in the proposer role may pass but may not repeat proposals. The temperature T declines at a mutually agreed-upon rate during this process. This protocol thus reproduces the key elements of the annealing mediator protocol – a time-dependent annealing regime plus tokens - without the need for a mediator. Our experiments show that this protocol produces results just as good as the annealing mediator, while requiring fewer proposal exchanges (averaging about 200 exchanges per negotiation).

Conclusions

We have shown that negotiation involving complex contracts (i.e. those with many multiple inter-dependent issues) has properties that are substantially different from the simple (independent issue) case that has been studied to date in the negotiation literature, and requires as a result different protocols in order to achieve near-optimal outcomes. This paper presents, as far as we are aware, the first negotiation protocol designed specifically for complex contracts. The essence of our own approach can be summarized simply: conceding early and often (as opposed to little and late, as is typical for independent issue negotiations) is a key to negotiating good complex contracts. Conceding is not individually rational in the face of agents that may choose not to concede, but this problem can be resolved either by introducing a mediator that stochastically ignores agent preferences, or by introducing dice into the negotiation protocol. In both cases, the exchange of tokens when one agent overrides another can be used to incent the truthful voting that enables win-win outcomes.

References

1. Bar-Yam, Y., *Dynamics of complex systems*. 1997, Reading, Mass.: Addison-Wesley. xvi, 848.

2. Raiffa, H., *The art and science of negotiation*. 1982, Cambridge, Mass.: Belknap Press of Harvard University Press. x, 373.

For more information on this topic, please refer to the paper: Klein, M., et al., Protocols for Negotiating Complex Contracts. IEEE Intelligent Systems, 2003. 18(6): p. 32 - 38.

CENTER FOR EBUSINESS MISSION

Founded in 1999, the Center for eBusiness is the largest research center in the history of the MIT Sloan School. Our research is supported by the National Science Foundation and corporate sponsors. We fund more than 45 faculty and more than 60 research projects. Our mission is to be the leading academic source of innovation in management theory and practice for eBusiness.

Examples of Current Focused Research Projects:

- Theory T: Trust-Based Marketing
- Implications of e-Commerce for New Services and Structure of Logistics Systems
- How Do Intangible Assets Affect the Productivity of Computerization Efforts?
- Wireless and Mobile Commerce Opportunities for Payments Services
- Two-Tier Support Business Models
- The Impact of the Internet on the Future of the Financial Services Industry
- Pricing Products and Services in the High-Tech Industry

The Center for eBusiness has recently entered into Phase II, adjusting its agenda to focus more explicitly on business value, while at the same time including technologies beyond the Internet in its purview. The early period of exploration and experimentation is coming to an end and there is now the opportunity, and the necessity, to focus more explicitly on using digital technologies to deliver measurable business value. Amidst all this change, the business fundamentals of investment, revenues, expenses, profits, and satisfying customers have only grown more important. At the same time, a broader, inter-related set of technologies is at our disposal. While the Internet has been an important catalyst, related digital technologies are often at least as relevant.

We are co-located with MIT Sloan's Center for Information Systems Research initiative and the Center for Coordination Science to facilitate collaboration. We also collaborate with the Media Lab and the Program on Internet and Telecoms Convergence.

The Center for eBusiness gratefully acknowledges the support and contributions of its Sponsors.

CENTER FOR EBUSINESS SPONSORS

Founding Sponsors

BT
General Motors
Intel
MasterCard International
PricewaterhouseCoopers
UPS

Research Sponsors

CSK
France Telecom
Nortel Networks
Suruga Bank

Member Sponsors

Amazon
Bank of Tokyo-Mitsubishi
Cisco
GEA
Publicis Technology
Qwest Communications
SAS

CONTACT INFORMATION

Center for eBusiness at MIT
MIT Sloan School of Management
3 Cambridge Center, NE20-336
Cambridge, MA 02142
Telephone: 617/253-7054
Facsimile: 617/452-3231
http://ebusiness.mit.edu/

Glen L. Urban, Chairman
Erik Brynjolfsson, Director
David Verrill, Executive Director
Steve Buckley, Communications & IT
Robynne DeCaprio, Program Administration Associate
Meredith Sampson, Financial Assistant