Protocols for negotiating complex contracts

Citation: Klein, M. et al. "Protocols for Negotiating Complex Contracts." IEEE Intelligent Systems 18, 6 (November 2003): 32–38 © 2003 IEEE

As Published: http://dx.doi.org/10.1109/MIS.2003.1249167

Publisher: Institute of Electrical and Electronics Engineers (IEEE

Persistent URL: http://hdl.handle.net/1721.1/113705

Version: Original manuscript: author's manuscript prior to formal peer review

Protocols for Negotiating Complex Contracts[1]

Mark Klein
Massachusetts Institute of Technology
m_klein@mit.edu

Peyman Faratin
Massachusetts Institute of Technology
peyman@mit.edu

Hiroki Sayama
University of Electro-Communications
sayama@hc.uec.ac.jp

Yaneer Bar-Yam
New England Complex Systems Institute
yaneer@necsi.org

Abstract

Work to date on negotiation protocols has focused on defining contracts consisting of one or a few independent issues. Many real-world contracts, by contrast, are much more complex, consisting of multiple inter-dependent issues and intractably large contract spaces. This paper describes a simulated annealing based approach appropriate for negotiating such complex contracts that achieves near-optimal outcomes for negotiations with binary issue dependencies.

Keywords: non-linear mediated single text unmediated proposal exchange negotiation, multiple interdependent issues, prisoner's dilemma

1. 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. An agent starts with 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 (Figure 1):

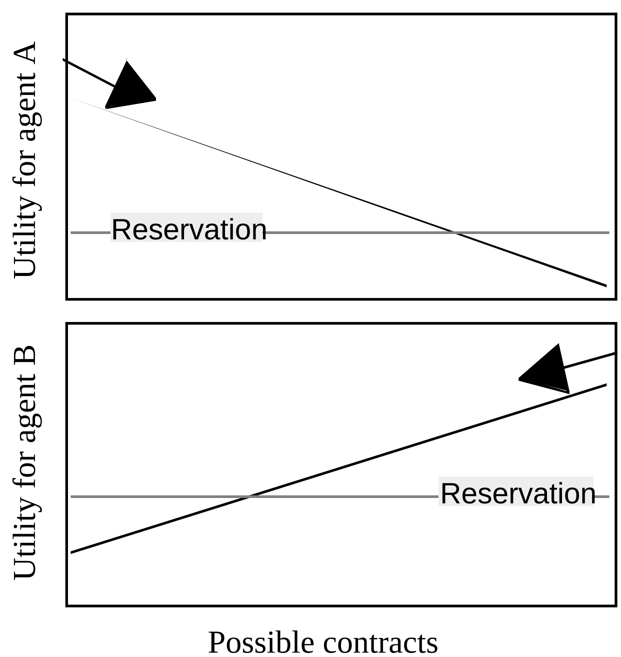

Figure 1: The proposal exchange model of negotiation, applied to a simple contract. The Y axis represents the utility of a contract to each agent. Each point on the X axis represents a possible contract,

[1] To appear in: IEEE Intelligent Systems Journal, special issue on agents and markets.

ordered in terms of its utility to agent B. Since there is no need to negotiate over issues that both parties agree upon, we only consider issues where improvement for one party represents a decrement for the other. The arrows represent how agents begin with locally optimal proposals, and concede towards each other, with their subsequent proposals, as slowly as possible. Note that we have, for presentation purposes, ‘flattened’ the contract space onto a single dimension, but there should actually be one dimension for every issue in the contract.

This is a perfectly reasonable approach for simple contracts. 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. Simple contract negotiations thus typically progress as follows:

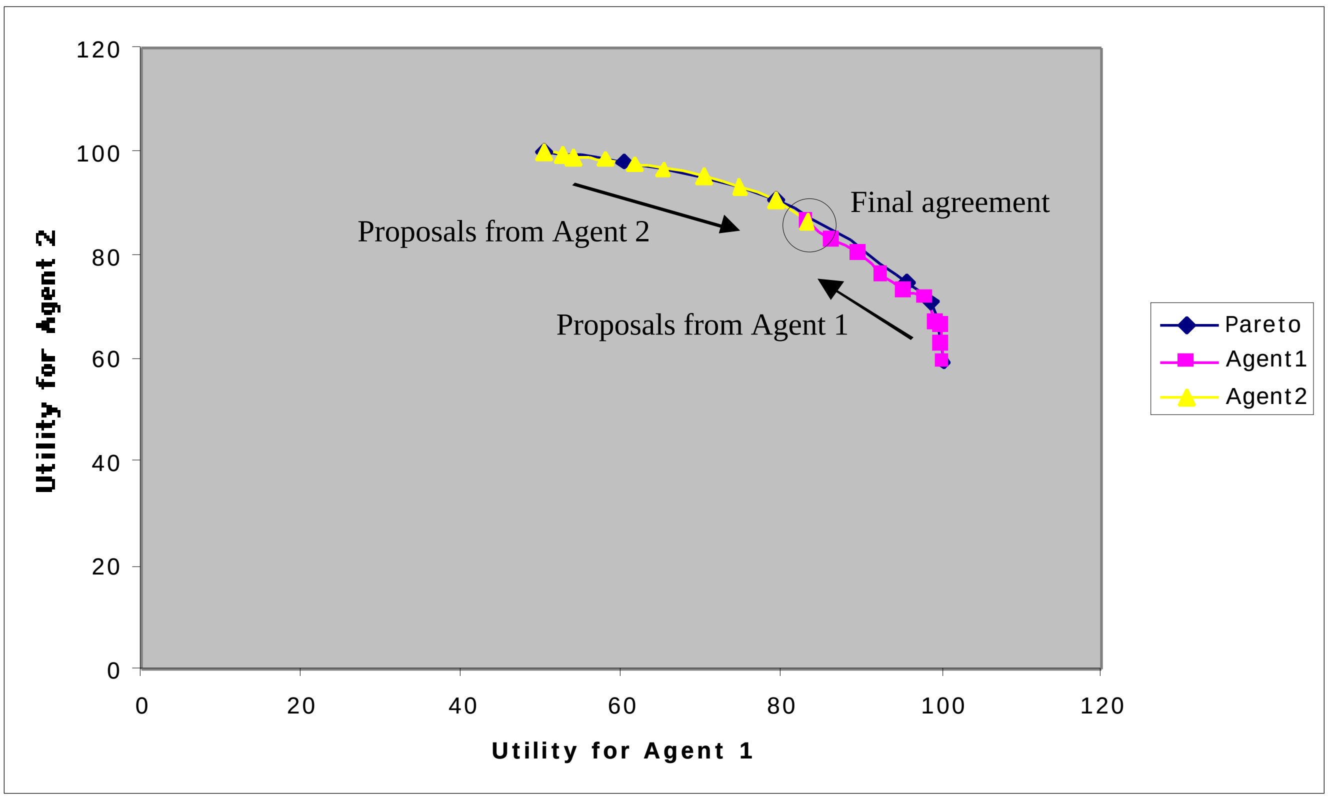

Figure 2. The utilities for the proposals made in a typical simple contract negotiation. The contract consisted in this case of 40 binary issues. Each agent starts with a locally optimal proposed contract (at the extremes of the Pareto frontier) and is required to reduce the Hamming distance (number of issues with different values) between the two agents’ proposals, until an agreement is reached. With simple contracts, this results in optimal outcomes. The Pareto frontier, representing the set of optimal contracts, was estimated by applying an annealing optimizer to differently weighted sums of the two agents’ utility functions.

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 distinct 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. The value to me of a given couch, for example, depends on whether it is a good match with the chair I plan to purchase with it. Such issue interdependencies lead to nonlinear utility functions with *multiple* local optima [1]:

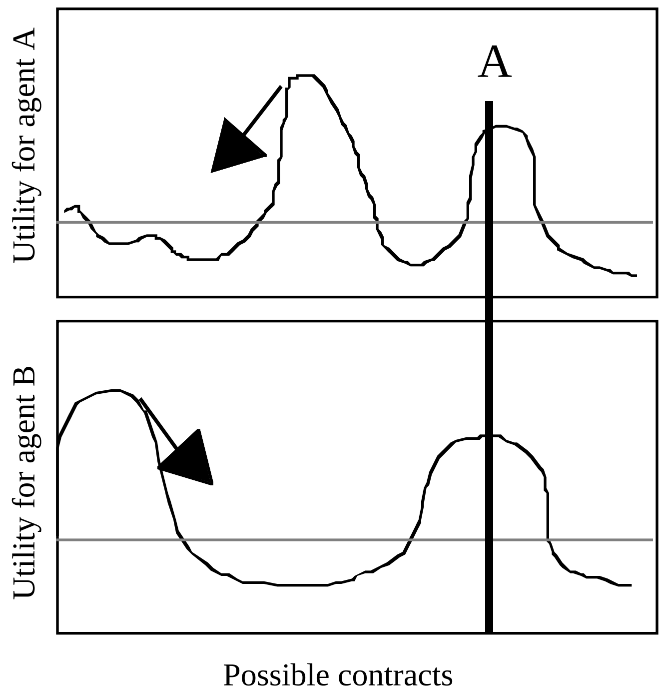

Figure 3: An example of proposal exchange applied to a complex contract. Because of issue inter-dependencies, the utility functions have multiple optima. The arrows show what happens when each agent begins at a local optimum and concedes towards the other: win-win solutions (such as that represented by contract A) found elsewhere in the contract space can be missed.

In such contexts, an agent finding its own ideal contract becomes a nonlinear optimization problem, difficult in its own right. Simply conceding toward the other agents' proposals can result in the agents missing contracts that would be superior from both their perspectives (e.g. the contract labeled "A" in figure 3 above). Standard negotiation techniques thus typically produce the following behavior when applied to complex contract negotiation:

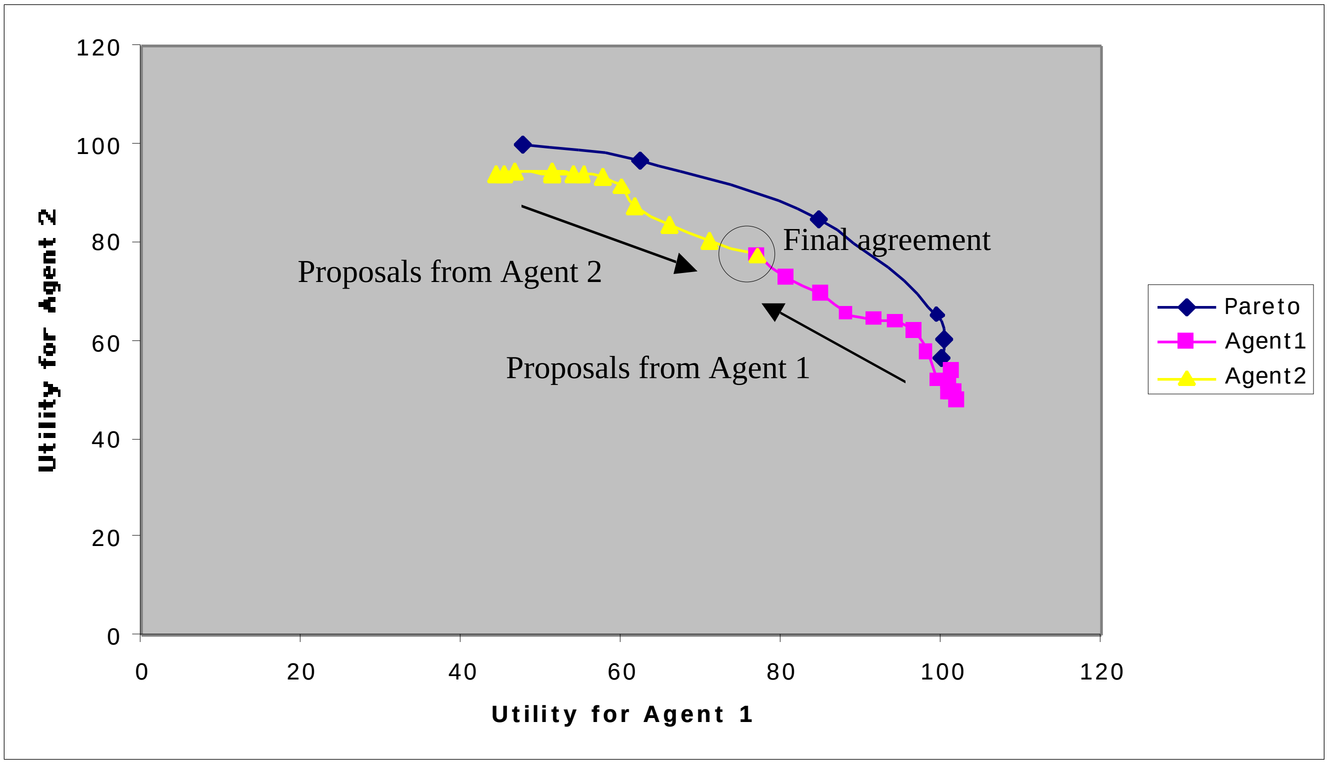

Figure 4. The utilities for the proposals made in a typical complex contract negotiation. This example differs from figure 2 only in that a nonlinear utility function was used by each agent (details below). As we can see, the minimal concession protocol that works optimally for simple contracts produces outcomes, for complex contracts, that are substantially sub-optimal.

The agents start with an approximation to their ideal contract and diverge increasingly from the Pareto frontier as they converge upon an agreement. The degree of sub-optimality depends on the details of the utility function. In our experiments, for example, the final contracts' averaged 94% 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. It is striking that such relatively forgiving multi-optima utility functions lead to substantially sub-optimal negotiation outcomes.

These sub-optimal outcomes represent a fundamental weakness with current negotiation techniques. The only way to ensure that subsequent proposals track the Pareto frontier, and thus conclude with a Pareto optimal result, is to be able to identify the proposal that represents the minimal concession from the current one. But in a utility function with multiple optima, that proposal may be quite distant from the current one, and the only way to find it is to exhaustively enumerate all possible contracts. This is computationally infeasible, however, due to the sheer size of the contract space. Since the utility functions are quite complex, it is in addition no longer practical for one agent to infer the other's utility function. Complex contracts therefore require different negotiation techniques which allow agents to find 'win-win' contracts in intractable multi-optima search spaces in a reasonable amount of time. In the following sections we describe a family of negotiation protocols that make substantial progress towards achieving these goals. The paper is structured as follows. We begin by describing how a well-known non-linear optimization technique (simulated annealing) can be integrated the mediated single text negotiation protocol to produce an approach that offers near-optimal outcomes for complex contract negotiations. We reveal the prisoner's dilemma that results from this approach, and propose a refined protocol, based on parity-maintaining annealing mediator, that resolves that problem. We conclude with describing an unmediated version of the negotiation protocol that is also effective at producing near-optimal outcomes with complex contracts.

2. Mediated Single Text Negotiation

A standard approach for dealing with complex negotiations in human settings is the mediated single text negotiation [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:

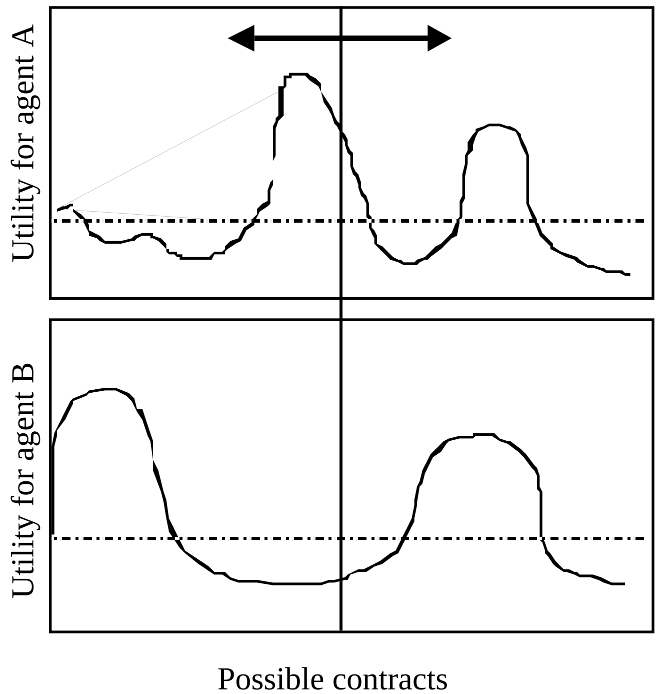

Figure 5: Single text negotiation. The vertical line represents the current proposed contract, and subsequent proposals move that line in the contract space.

Here, the vertical line represents the contract currently proposed by the mediator. Each new contract moves the line to a different point on the X axis. The goal is to find a contract that is sufficiently good for both parties.

We defined a simple simulation experiment to help us explore how well this approach actually works. In this experiment, there were two agents negotiating to find a mutually acceptable contract consisting of a vector S of 100 boolean-valued issues, each issue assigned the value 0 or 1, corresponding to the presence or absence of a given contract clause. This defined a space of 2^100, or roughly 10^30, possible contracts. Each agent had a utility function calculated using its own 100x100 influences matrix H, wherein each cell represents the utility increment or decrement caused by the presence of a given pair of issues, and the total utility of a contract is the sum of the cell values for every issue pair present in the contract:

$$U = \sum_{i=1}^{100} \sum_{j=1}^{100} H_{ij} S_j S_j$$

The influence matrix therefore captures the bilateral dependencies between issues, in addition to the value of any individual contract clause. For our experiments, the utility matrix was initialized to have random values between –1 and +1 in each cell. A different influences matrix was used for each simulation run, in order to ensure our results were not idiosyncratic to a particular configuration of issue inter-dependencies.

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 its 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.

3. 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
Agent1 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 while paired annealers do very well. If both agents are hill-climbers they both get a poor payoff, since it is difficult to find many contracts that represent an improvement for both parties. A typical negotiation with two hill-climbers looks like the following:

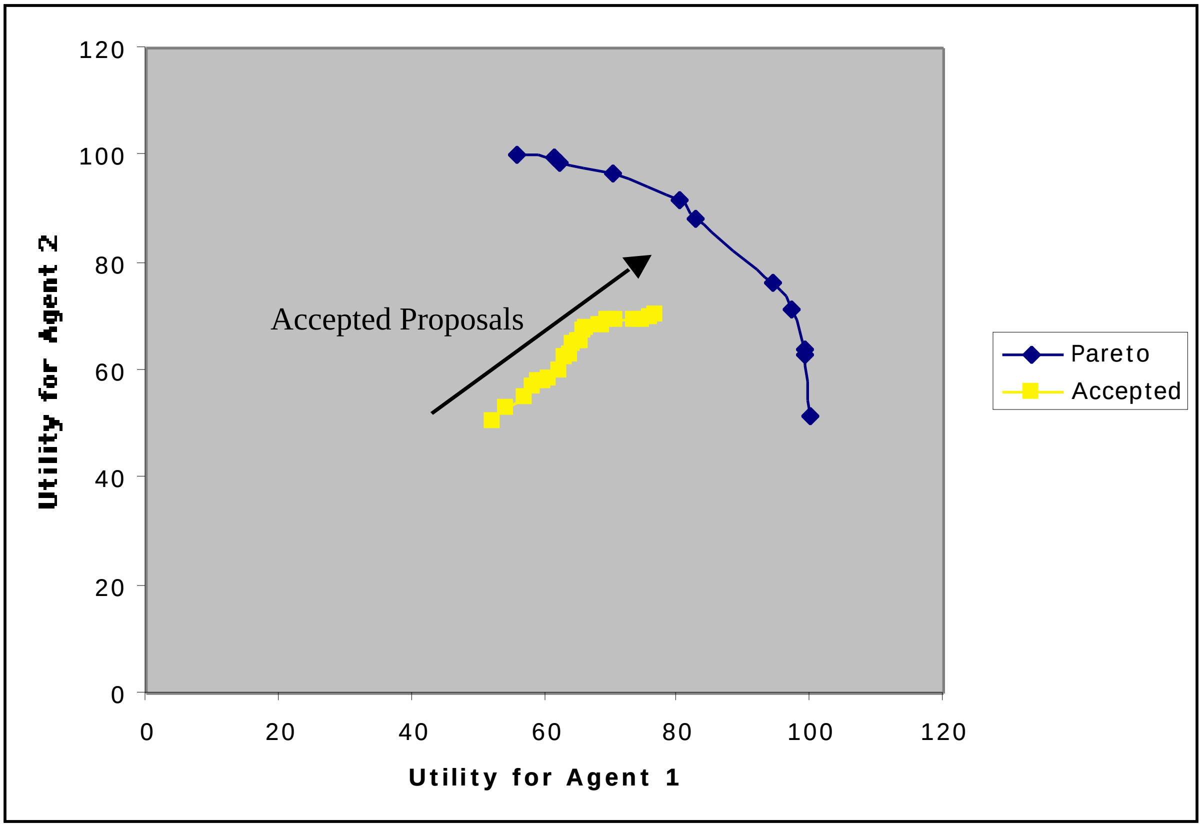

Figure 6: The utilities for the accepted proposals in a typical single text complex contract negotiation with two hill-climbers. The mediator's initial proposal is at the lower left, and the subsequent accepted proposals move towards higher utilities for both agents.

As we can see, in this case the mediator was able to find only a handful of contracts that increased the utility for both hill-climbers, and ended up with a poor final social welfare.

Near-optimal social welfare can be achieved, by contrast, when both agents are annealers, willing to initially accept individually worse contracts so they can find win-win contracts later on:

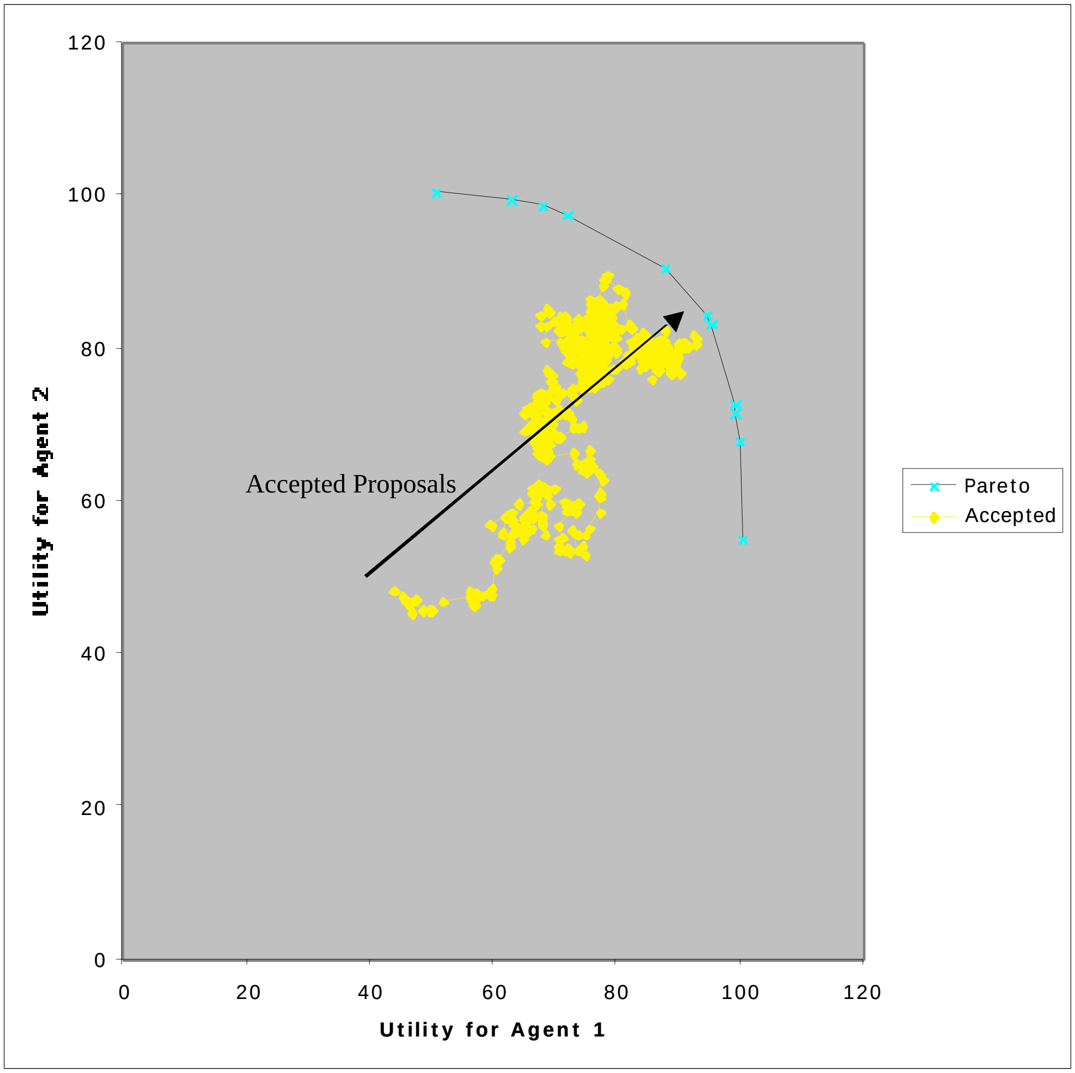

Figure 7: The utilities for the accepted proposals for a typical single text complex contract negotiation with two annealers. Some of the accepted proposals actually cause utility decrements for one or both agents, but the final result is a near-optimal contract.

The agents entertain a much wider range of contracts, eventually ending very near the Pareto frontier.

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 (Figure 8). 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:

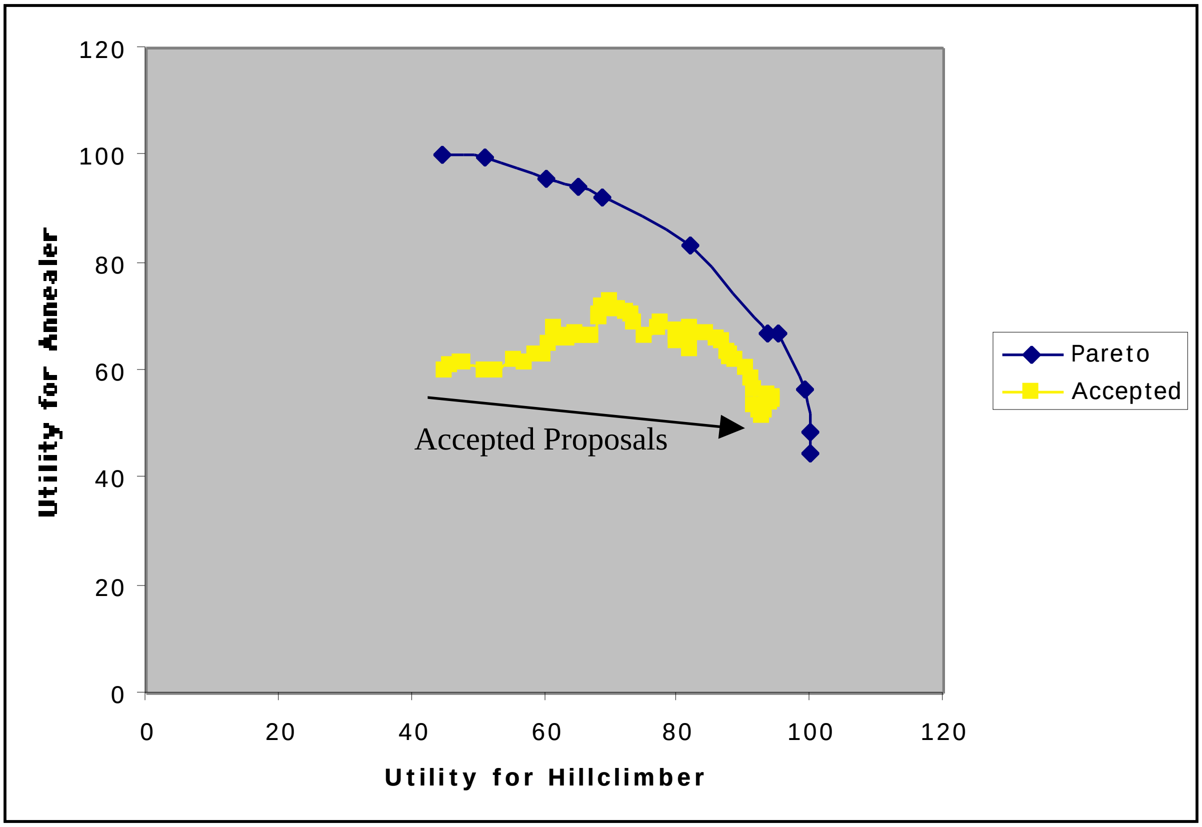

Figure 8: The utilities for the accepted proposals for a typical single text complex contract negotiation with an annealer and a hill climber. Note that the hill climber achieves a near-optimal contract at the expense of the annealer.

This reveals a dilemma. In negotiation contexts we typically can not assume that agents will be altruistic, and we must as a result design protocols such that the individually most beneficial negotiation strategies also produce the greatest social welfare [3]. In our case, however, 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 I). If all agents do this, however, then they forego the higher individual utilities they would get if they both annealed. Individual rationality thus drive the agents towards the strategy pairing with the *lowest individual and social welfare.* This is thus an instance of the prisoner's dilemma. It has been shown that this dilemma can be avoided if we assume repeated interactions between agents [4], but we would prefer to have a negotiation protocol that incents socially beneficial behavior without that difficult-to-enforce constraint. Several straightforward approaches to this problem, however, prove unsuccessful. One possibility is to simply reduce the annealer's willingness to make concessions. This can indeed eliminate the conceder's penalty, but at the cost of achieving social welfare values only slightly better than that achieved by two hill climbers. Another option is to have agents switch from being an annealer to a hill-climber if they determine, by observing the proposal acceptance rates of their opponents, that the other agent is being a hill-climber. We found, however, that it takes too long to determine the type of the other agent: by the time it has become clear, much of the contract utility has been committed, and it is too late to recover from the consequences of having started out as an annealer. See [5] for details.

4. The Annealing Mediator

We were able to define a negotiation protocol that avoids the prisoner's dilemma entirely in mediated single-text negotiation of complex contracts. The trick is simple: 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. In our original protocol, the mediator would simply propose modifications of the last contract both negotiating agents accepted. In our refined protocol, 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 are free to remain hill-climbers and thus avoid the potential of 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. We describe the details of our protocol, and our evaluations thereof, below.

In our initial implementations each agent gave a simple accept/reject vote for each proposal from the mediator, but we found that this resulted in final social welfare values significantly lower than what we earlier achieved using annealing agents. In our next round of experiments we accordingly modified the agents so that they provide additional information to the mediator in the form of vote strengths: each agent annotates an accept or reject vote as being *strong* or *weak*. The agents were designed so that there are roughly an equal number of weak and strong votes of each type. This maximizes the informational content of the vote strength annotations. When the mediator receives these votes, it maps them into numeric values (strong accept = 1, weak accept = 0, weak reject = -1, strong reject = -2) and adds them together to produce an aggregate score. A proposal is accepted by the mediator if the score is non-negative, i.e. if both agents voted to accept it, or if a weak reject by one agent is overridden by a strong accept from the other. The mediator can also accept rejected contracts (i.e. those with a negative agrregate score) using the annealing scheme described above. This approach works surprisingly well, achieving final social welfare values that average roughly 99% of optimal despite the fact that the agents each supply the mediator with only two bits of information. We found, in fact, that increasing the number of possible vote weights did not increase final social welfare. This is because the strong/weak vote annotations are sufficient to allow the system to pursue social welfare-increasing contracts that cause a utility decrement for one agent.

5. Incentives for Truthful Voting

Any voting scheme introduces the potential for strategic non-truthful voting by the agents, and our scheme is no exception. 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.

As we can see, 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? We found that simply placing a limit on the number of strong votes each agent can use does not work. If the limit is too low, we effectively lose the benefit of vote weight information and get the lower social welfare values that result. If the strong vote limit is high enough to avoid this, then all an exaggerator has to do is save all of it's strong votes till the end of the negotiation, at which point it can drag the mediator towards making a series of proposals that are inequitably favorable to it.

Another possibility is to enforce overall parity in the number of "overrides" each agent gets. A override occurs when a contract supported by one agent (the "winner") is accepted by the mediator over the objections of the other agent. Overrides are what drags a negotiation towards contracts favorable to the winner, so it makes sense to make the total number of overrides equal for each agent. But this is not enough, because exaggerators always win disproportionately more than the truth-teller.

The solution, we found, came from enforcing parity between the number of overrides given to each agent *throughout* the negotiation, 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, moreover, 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 this work? Why, in particular, 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. Recall that in this case a truthful agent, offering a mix of strong and weak votes, is paired with an exaggerator for whom at least some weak accepts and rejects are presented as strong ones. The truthful agent can therefore only get an override via annealing (see Table 3), and this is much more likely when its vote was a strong accept rather than a weak one. In other words, the truthful agent spends its tokens almost exclusively on contracts that truly offer it a strong utility increase. The exaggerator, on the other hand, will spend tokens to elicit a override 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.

6. The Unmediated Single Text Protocol

The protocol we have just considered worked well in the contexts studied but 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 proposing and voting 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 aggregate score (ΔU) is calculated as for the annealing mediator, and 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 its' 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, averaging 99% of optimal, while requiring fewer proposal exchanges (averaging about 200 exchanges per negotiation).

7. Contributions

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. While some previous work has studied multi-issue negotiation (e.g. [6] [7] [8]) the issue utilities in these efforts are treated as independent, so the utility functions for each agent are linear, with single optima. As we have seen, however, the introduction of multiple optima changes the game drastically. Multi-attribute auctions [9] represent another scheme for dealing with multiple issues, wherein one party (the buyer) publishes its' utility function, and the other parties (the sellers) make bids that attempt to maximize the utility received by the buyer. If none of the bids are satisfactory, the buyer modifies its' published utility function and tries again. This introduces a search process, and the problem with this approach is that it does not provide any guidance for how the parties involved should control their search through the vast space of possibilities. 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.

8. Next Steps

There are many other promising avenues for future work in this area. The high social welfare achieved by our approach partially reflect the fact that the utility functions for each agent, based as they are solely on binary dependencies, are relatively easy to optimize. Higher-order dependencies, common in many real-world contexts, are known to generate more challenging utility landscapes [10]. We hypothesize that it may be necessary to adapt non-linear optimization techniques such as genetic algorithms into the negotiation context in order to address this challenge. Another possibility involves agents providing limited information about their utility functions to the mediator or to each other in order to facilitate more intelligent search through very large contract spaces. Agents can, for example, tell the mediator which issues are heavily dependent upon each other, allowing the mediator to focus its attention within tightly-coupled issue 'clumps', leaving other less influential issues till later. We hypothesize that agents may be incented to tell the truth in order to ensure that negotiations can complete in an acceptable amount of time. Finally, we would like to derive formal incentive compatibility proofs (i.e. concerning when agents are incented to vote truthfully) for our protocols. New proof techniques will probably be necessary because previous results in this area have made strong assumptions concerning the shape of the agent utility functions that do not hold with complex contracts.

9. Acknowledgements

This work was supported by funding from the DARPA Control of Agent-Based Systems (CoABS) program, and the NSF Computation and Social Systems program.

10. 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.
3. Rosenschein, J.S. and G. Zlotkin, *Rules of encounter : designing conventions for automated negotiation among computers*. Artificial intelligence. 1994, Cambridge, Mass.: MIT Press. xxi, 229.
4. Axelrod, R., *The Evolution Of Cooperation*. 1984: Basic Books.
5. Klein, M., P. Faratin, and Y. Bar-Yam, *Using an Annealing Mediator to Solve the Prisoners' Dilemma in the Negotiation of Complex Contracts*, in *Proeedings of the Agent-Mediated Electronic Commerce (AMEC-IV) Workshop*. 2002, Springer-Verlag.
6. Faratin, P., C. Sierra, and N.R. Jennings, *Using similarity criteria to make negotiation trade-offs*. Proceedings Fourth International Conference on MultiAgent Systems. IEEE Comput. Soc., 2000.
7. Jonker, C.M. and J. Treur. *An Agent Architecture for Multi-Attribute Negotiation*. In the proceedings of *Proceedings of IJCAI-01*. 2001.
8. Fatima, S.S., M. Wooldridge, and N. R.Jennings. *MultiIssue Negotiation Under Time Constraints*. In the proceedings of *Proceedings of the 2002 International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-02)*. 2002. Bologna, Italy: ACM.
9. Bichler, M. and J. Kalagnanam, *Bidding Languages and Winner Determination in Multi-Attribute Auctions*. European Journal of Operational Research, In Press.
10. Kauffman, S.A., *The origins of order: self-organization and selection in evolution*. 1993: Oxford University Press.