Group Decision and Negotiation **12**: 111–125, 2003

Negotiating Complex Contracts

MARK KLEIN AND PEYMAN FARATIN
Massachusetts Institute of Technology, Building NE20-336, Cambridge, MA 02139, USA
(E-mail: m_klein@mit.edu; peyman@mit.edu)

HIROKI SAYAMA AND YANEER BAR-YAM
New England Complex Systems Institute, 24 Mt. Auburn Street, Cambridge, MA 02138, USA
(E-mail: sayama@necsi.org; yaneer@necsi.org)

Abstract

Work to date on computational models of negotiation has focused almost exclusively on defining contracts consisting of one or a few independent issues and tractable contract spaces. 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 social welfares for negotiations with binary issue dependencies.

Key words: interdependent issues, non-linear negotiation

1. Introduction

Work to date on computational models of negotiation has focused almost exclusively on defining contracts consisting of one or a few independent issues (Ehtamo, Ketteunen, and Hamalainen 2001; Faratin, Sierra, and Jennings 2000). We can frame what these techniques do (Figure 1).

Each point on the X-axis represents a candidate contract.[1] The Y-axis represents the utility of each contract to each agent, where higher is better. Each agent will only accept contracts whose utility is above that agent's reservation value. The utility functions for each issue are typically linear (e.g., as in price), monotonic, or single optimum. Since relative few issues are involved, the space of all possible contracts can be explored exhaustively, and since the issues are independent, the utility functions for each issue are superimposed linearly. The result is that the overall utility function for different possible is *linear*, with a single optimum in the utility function for each agent, and therefore easy to optimize. In such a context, the reasonable strategy is for each agent to start at its own ideal contract, and concede, through iterative proposal exchange, just enough to get the other party to accept the contract. Since the utility functions are simple, it is feasible for one agent to infer enough about the opponent's utility function through observation to make concessions likely to increase the opponent's utility.

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 to hundreds of

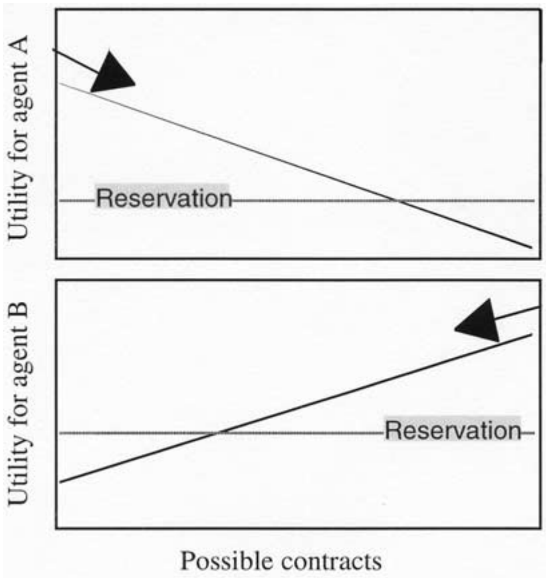

Figure 1. The standard view of negotiation.

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 (Bar-Yam 1997).

In such contexts, an agent finding its own ideal contract becomes a nonlinear optimization problem, difficult in its own right. Simply conceding as slowly as possible from one's ideal can result in the agents missing contracts that would be superior from both agent's perspectives. In Figure 2, for example, if both agents simply concede slowly from their own ideal towards the opponents' ideal, they will miss the better contracts on the right. Exhaustive search for "win-win" contracts, however, is impractical due to the size of the search spaces involved. Finally, since the utility functions are quite complex, it is no longer practical for one agent to learn the other's utility function.

Imagine, for example, that we have a mediated single text negotiation with hill-climbing agents (these terms are defined in the next section). We find that a "single mutation" mediator (one that successively proposes contracts where a single issue value has been changed) does quite well for independent issues (the social welfare averages 98% of optimal) but relatively poorly when the issues have binary dependencies (87% of optimal).

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 negotiation approach that make substan-

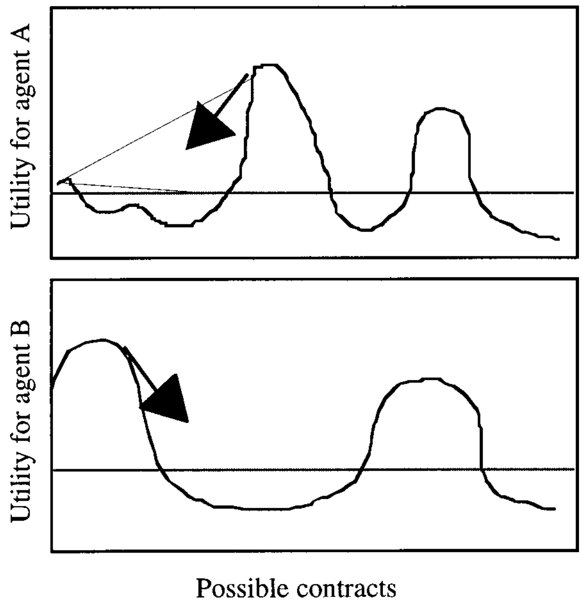

Figure 2. Complex negotiation.

tial progress towards achieving these goals. Our paper is structured as follows. We first describe the negotiation protocol we selected – mediated single text negotiation – and point out how a straightforward application of a well-known nonlinear optimization technique leads to a prisoner's dilemma game wherein the agents are individually incented to use strategies that produce inferior contracts. We then describe several apparently reasonable remedies that don't work, as well as a novel approach – what we call the parity-preserving annealing mediator – which does.

2. Mediated single text negotiation

A standard approach to dealing with complex negotiations in human settings is the mediated single text negotiation (Raiffa 1982). 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 the reservation utility value is met or exceeded for both parties. We can visualize this process (Figure 3).

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 experiment to help us explore how this approach could be instantiated in a computational framework. In this experiment, there were two agents negotiat-

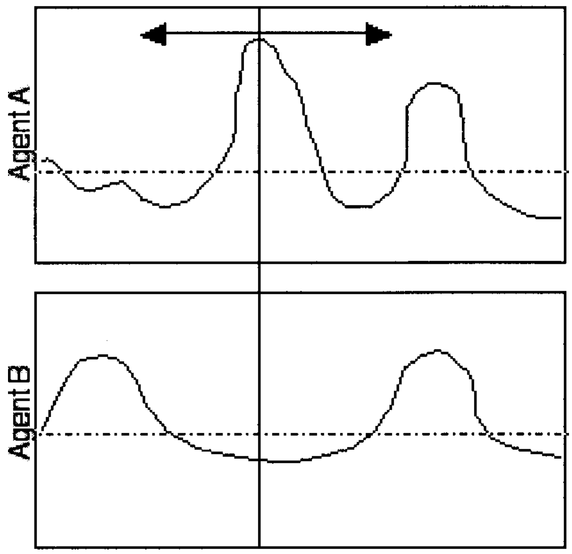

Figure 3. Single text negotiation.

ing 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 100 × 100 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:

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

The influence matrix therefore captures the 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 acceptable 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, implementing a Monte Carlo machine [5]. Each annealer haw a virtual “temperature” T, such that it will accept contracts worse than last accepted one with the probability:

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

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 (Bar-Yam 1997). This suggests that annealers will be more successful than hill-climbers in finding good contracts through the negotiation process.

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 (Table 1).

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 Figure 4.

Figure 4 shows the normalized utilities of the accepted contracts for each agent, plotted next to the pareto-efficient line (estimated by applying an annealing optimizer to different weighted sums of the two agents’ utility functions). As we can see, in this case the mediator was able to find only two contracts that increased the utility for both hill-climbers, and ended up with a poor final social welfare.

Near-optimal social welfares are achieved, by contrast, when both agents are annealers, both willing to initially accept individually worse contracts to help find win-win contracts later on (Figure 5).

Table 1. Annealing vs. hill-climbing agents

	Agent 2 hill-climbs	Agent 2 aneals
Agent 1 hill-climbs	[0.86] 0.73/0.74	[0.86] 0.99/0.51
Agent 1 anneals	[0.86] 0.51/0.99	[0.98] 0.84/0.84

In the table, the cell values are laid out as follows:
[<social welfare optimality.]
<agent 1 optimality >/<agent 2 optimality>

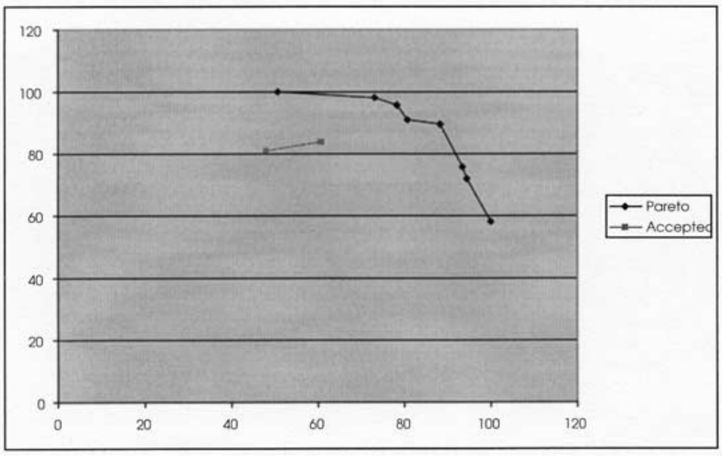

Figure 4. A typical negotiation with two hill-climbers.

As we can see in Figure 5, the agents entertain a much wider range of possible 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 6). This pattern can be understood as follows. When an annealer is at a high virtual temperature, it becomes a

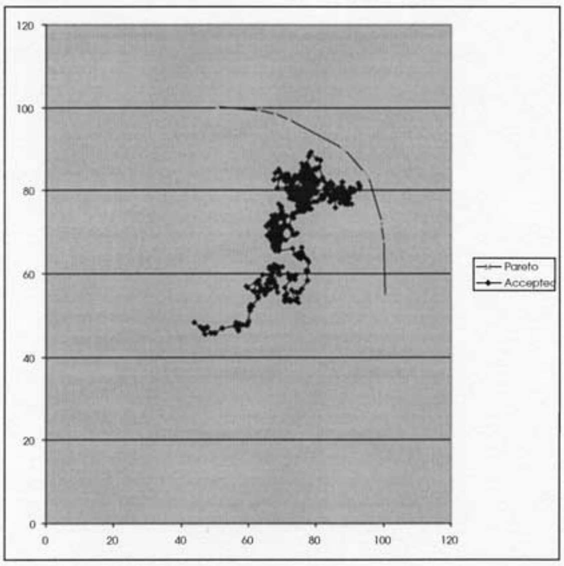

Figure 5. A typical negotiation with two annealers.

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 many negotiation contexts we can not assume agents will be altruistic, and we must as a result design negotiation protocols such that the individually most beneficial negotiation strategies also produce the greatest social welfare (Rosenschein and Zlotkin 1994; Sandholm 1998). 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 1). If all agents do this, however, then they forego the higher individual utilities they would get if they both annealed. The individual strategic considerations thus drive the system towards the strategy pairing with the *lowest 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 (Axelrod 1984), but ideally we would prefer to have a negotiation protocol that incents socially beneficial behavior without that difficult-to-enforce constraint.

4. Partial solutions: adaptive and cold annealers

If both agents could know ahead of time what strategy the other agent is going to use, then all agents would select annealing. In an open system environment we can not rely on self-reports for this, however, since agents are incented to lie, i.e. claim they will use annealing

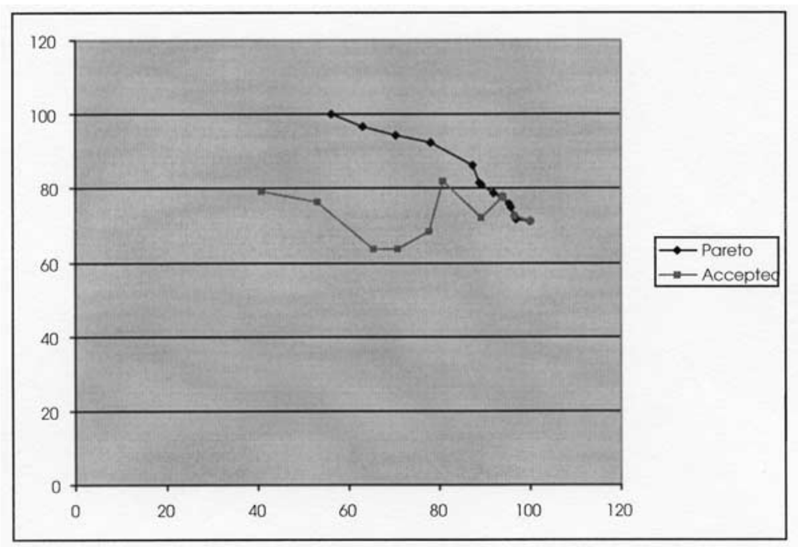

Figure 6. A typical negotiation with an annealer and hill climber.

but actually hill-climb. An agent must thus be able to determine the type of its opponent based purely on observing its behavior. It turns out this is relatively easy to do. An annealer will tend to accept a much higher percentage of proposed contracts that a hill-climber, especially at higher virtual temperatures (Figure 7).

The problem with this "adaptive" approach is that determining the type of an agent in this way takes *time*. Agents must start with a guess concerning the other agent's strategy and then observe its voting behavior to see what it actually uses. As we can see above, the divergence in acceptance rates between annealers and hill-climbers only becomes clear after several hundred proposals have been exchanged. By this time, however, *much of the contract utility has already been committed*, so it is too late to fully recover from the consequences of having guessed wrong (Figure 8).

In our experiments, for example, between 40 and 60% of the final social welfare had already been committed in the first 100 proposal exchanges. The early commitment of utility is a result of the topology of nonlinear utility functions. These functions tend to be fractal (i.e., self-similar at different scales) with the highest optima also tending to be the widest, so the steepest slope tends to occur earlier, and the slope reduces as one gets closer to the summit (Bar-Yam 1997).

Adaptive strategies therefore can not eliminate the prisoner's dilemma, just reduce its magnitude. Let us consider a specific example of an adaptive strategy we can call "tit-for-tat" (T4T). In this strategy, an agent starts as an annealer, and then switches to hill-climbing if the other agent proves to be a hill-climber. One could argue that it is more rational to

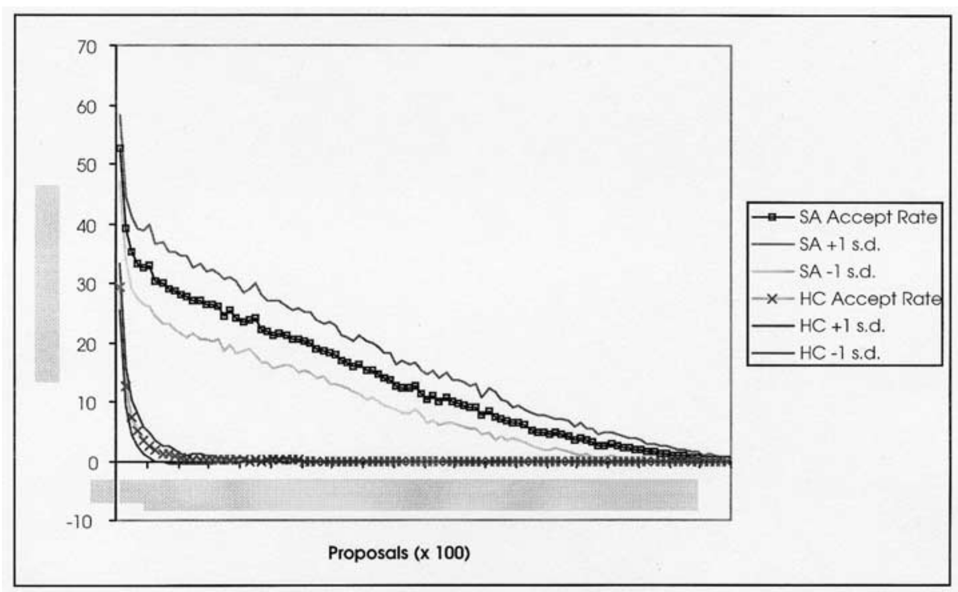

Figure 7. Proposal acceptance percentages for hill-climbers and annealers ± 1 standard deviation.

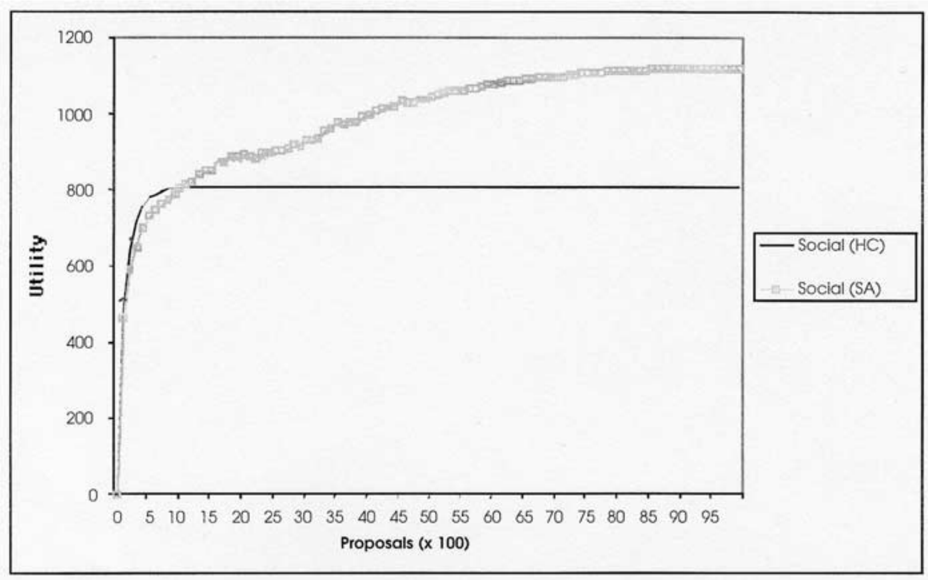

Figure 8. Social welfare values over time, averaged over 100 simulation runs.

start with the individually dominant strategy (hill-climbing), thereby avoiding the conceder's penalty, and then switch to annealing if the other agent is an annealer. But if everyone does this everyone will stay stuck in hill-climbing so we still get poor social welfare values. If we test the annealing-first T4T strategy we get the payoffs in Table 2.

A T4T agent fares just as well as an annealer when paired with an annealer or another T4T agent, and has a reduced conceder's penalty when paired with a hill-climber as compared to an annealer. The strategic picture is thus inconclusive: if you are paired with T4T agent, annealing is your best choice. But if you are paired with an annealer, hill climbing is your best choice. So annealing, the socially most beneficial strategy, is still not individually dominant.

Another strategy for reducing the conceder's penalty is for the annealer to start at a lower temperature, so that it can not be dragged as far from its own optimum (Figure 9).

Table 2. Payoffs with tit-for-tat agents

	Agent 2 hill-climbs	Agent 2 anneals	Agent 2 T4T
Agent 1 hill-climbs	[800]	[880]	[840]
	400/400	700/180	500/340
Agent 1 anneals	[880]	[1100]	[1100]
	180/700	550/550	550/550
Agent 1 T4T	[840]	[1100]	[1100]
	340/500	550/550	550/550

If the annealer agent starts at a low enough temperature (T0 = 3 in this case), the conceder's penalty is in fact eliminated, but at the cost of achieving social welfare values only slightly better than that achieved by two hill climbers.

5. The annealing mediator

We were able to develop 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 and adds them together according to the following simple scheme:

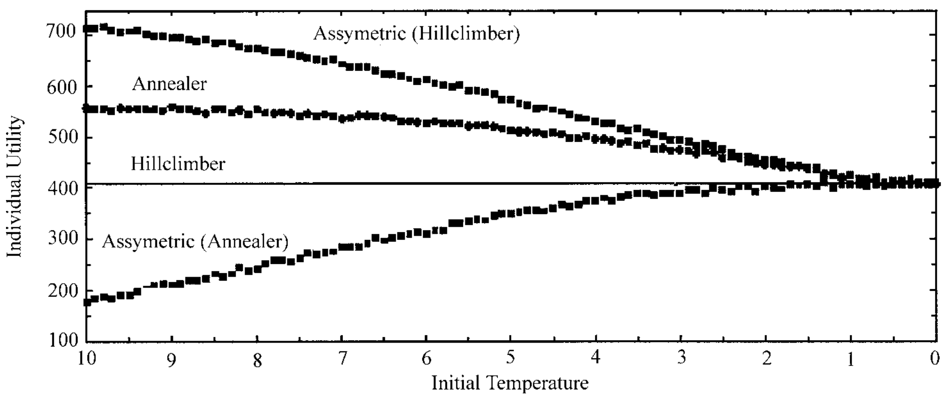

Figure 9. Individual utilities as a function of annealer agent starting temperature.

	Strong accept (1)	Weak accept (0)	Weak reject (–1)	Strong reject (–2)
Strong accept (1)	Accept (2)	Accept (1)	Mixed accept (0)	Weak reject (–1)
Weak accept (0)	Accept (1)	Accept (0)	Weak reject (–1)	Medium reject (–2)
Weak reject (–1)	Mixed accept (0)	Weak reject (–1)	Medium reject (–2)	Strong reject (–3)
Strong reject (–2)	Weak reject (–1)	Medium reject (–2)	Strong reject (–3)	Very strong reject (–4)

A proposal is thus accepted by the mediator 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 in addition occasionally accepts rejected contracts (i.e., with a negative overall 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. This additional bit of information is critical, however, because it allows the system to pursue social welfare-increasing contracts that cause a utility decrement for one agent.

6. 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 (Table 3).

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 game we encountered in our earlier work. 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 ben-

Table 3. Truth-telling vs. exaggerating agents with a simple annealing mediator

	Agent 2 exaggerates	Agent 2 tells truth
Agent 1 exaggerates	[0.92] 0.81/0.81	[0.93] 0.93/0.66
Agent 1 tells truth	[0.93] 0.66/0.93	[0.99] 0.84/0.84

efit 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 "mixed wins" each agent gets. A mixed win occurs when a contract supported by one agent (the "winner") is accepted by the mediator over the objections of the other agent. Mixed wins are what drags a negotiation towards contracts favorable to the winner, so it makes sense to make the total number of mixed wins 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 mixed accepts given to each agent *throughout* the negotiation, so neither agent can get more than a given advantage in the mixed win category. 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 for a mixed win gap limit of 3 (Table 4).

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, stochastically, or lying 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 win in mixed vote situations, 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 some weak accepts and rejects are presented as strong ones. The truthful agent can therefore only win a mixed vote 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, often spends its tokens trying to elicit a mixed win 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.

Table 4. Truth-telling vs. exaggerating agents with parity-enforcing mediator

	Agent 2 exaggerates	Agent 2 tells truth
Agent 1 exaggerates	[0.91] 0.79/0.79	[0.92] 0.78/0.81
Agent 1 tells truth	[0.92] 0.81/0.78	[0.98] 0.84/0.84

7. Contributions

We have shown that negotiation with multiple inter-dependent issues has properties that are substantially different from the independent issue case that has been studied to date in the computational negotiation literature, and requires as a result different negotiation schemes. This paper presents, as far as we are aware, the first computational negotiation approach suited for multiple issues with interdependent utilities. Kowalczyk and Bui (2001) describe a negotiation approach designed for multiple issues, but the issue utilities (as opposed to the viable issue *values*) are independent, so the utility functions for each agent are linear, with single optima. Their work therefore does not address the challenging and important problems that appear when we deal with non-linear utility functions. Multi-attribute auctions represent a related line of work, but while attribute interaction is recognized as important, research to date has generally assumed independence among attributes on both the buyer (bid value calculation) and seller (bid price calculation) sides, placing it into the domain of simple contracts with single-optimum utility functions (Bichler and Kalagnanam 2002; Kalagnanam and Parkes 2003). Auction protocols seem, moreover, poorly suited for the common challenge of very large contract spaces. A typical contract negotiation can easily include 10s to 100s of attributes, allowing trillions of possible contracts. In such contexts it is difficult for sellers to provide buyers with the utility functions that auction protocols require, simply because the contract space is too large for them to have explored exhaustively before hand. An iterative negotiation protocol such as ours, which only requires that parties assess the relative worth of pairs of contracts, appears much more realistic.

The essence of our approach can be summarized simply: conceding early and often (as opposed to little and late, as is typical for independent issue negotiations) is the key to achieving good contracts. We have also demonstrated that negotiation with inter-dependent issues produces a prisoner's dilemma game, and that introducing a mediator that stochastically ignores agent preferences and enforces running parity in agent influence resolves this dilemma. These results, we believe, are potentially relevant to any collaborative decision making task involving interdependent decisions.

8. Next steps

Higher order dependencies

The high social welfare values 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 contexts, are known to generate more challenging utility landscapes (Kauffman 1993), and will be addressed in future work. We speculate that non-linear optimization techniques such as genetic algorithms may represent a good starting point for handling such negotiation challenges.

Faster Negotiations

The simulated annealing approach produces better social welfares than hill-climbing but involves larger numbers of proposal exchanges. Hill-climbers typically reached stability after roughly 100 proposal exchanges, while the annealers approached stable utility values after roughly 800 proposal exchanges (Figure 8). This makes sense because hill-climbers simply climb to the top of the closest utility optimum and then stop, while annealers can, when at a high temperature at least, traverse multiple optima in the utility function. This is a potential problem however because, in competitive negotiation contexts, agents will typically wish to reveal as little information as possible about themselves for fear of presenting other agents with a competitive advantage. The more proposals considered, however, the more information an agent reveals.

What can we do about this? One option is to define better contract alternative generation operators. In our experiments the contract space was explored in random walk fashion, and all the "intelligence" was in the evaluation process. One example of a domain-independent approach is where agents provide the mediator with information concerning which issues are heavily dependent upon each other. Using this information, the mediator can focus its attention within the tightly-coupled issue "clumps", leaving the other less influential issues till later.

Another option is to introduce (limited) cooperative information exchange. It is clear that if agents cooperate they can produce higher contract utilities. Imagine for example that two hill-climbers vote to accept a contract based on whether it increases the social welfare, as opposed to their individual utilities. We have found that if we compare this with two "selfish" hill-climbers, the cooperative hill-climbers *both* benefit individually compared to the selfish case, thereby increasing social welfare as well. Other kinds of cooperation are imaginable. Agents can begin by presenting a list of locally [near-]optimal contracts, and then agree to explore alternatives around the closest matches in their two sets. Note that in the previous work with independent issues, this kind of information exchange has not been necessary because it relatively easy for agents to infer each other's utility functions from observing their negotiation behavior, but with inter-dependent issues and large multiple-optima utility functions this becomes intractable and information exchange probably must be done explicitly. 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.

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.

Notes

1. For simplicity of exposition we show only one dimension in these figures, but there is in actuality one dimension for every issue negotiated over.

References

Axelrod, R. (1984). *The Evolution Of Cooperation*. Basic Books.

Bar-Yam, Y. (1997). *Dynamics of Complex Systems*. Reading, Mass.: Addison-Wesley, xvi, 848.

Bichler, M. and J. Kalagnanam. (2002). *Bidding Languages and Winner Determination in Multi-Attribute Auctions*.

Ehtamo, H., E. Ketteunen, and R. Hamalainen. (1997). "Searching for Joint Gains in Multi-Party Negotiations," *European Journal of Operational Research* 1 (30), 54–69.

Faratin, P., C. Sierra, and N. R. Jennings. (2000). "Using Similarity Criteria to Make Negotiation Trade-Offs. *Proceedings Fourth International Conference on MultiAgent Systems*. IEEE Comput. Soc.

Kalagnanam, J. and D. Parkes. (2003). "Auctions, Bidding and Exchange Design," *in* Simchi-Levi Wu and Shen (eds.), *Supply Chain Analysis in the eBusiness Area*. Kluwer Academic Publishers.

Kalos, M. H. and P. A. Whitlock. (1986). *Monte Carlo Methods*. New York: J. Wiley & Sons.

Kauffman, S. A. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press.

Kowalczyk, R. and V. Bui. (2001). "On Constraint-Based Reasoning in e-Negotiation Agents," *in* F. Dignum and U. Cortes (eds.), *Agent-Mediated Electronic Commerce III, Current Issues in Agent-Based Electronic Commerce*. Springer, 31–46.

Raiffa, H. (1982). *The Art and Science of Negotiation*. Cambridge, Mass.: Belknap Press of Harvard University Press, x, 373.

Rosenschein, J. S. and G. Zlotkin. (1994). *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers*. Artificial intelligence. Cambridge, Mass.: MIT Press, xxi, 229.

Sandholm, T. W. (1998). "Distributed Rational Decision Making," *in* G. Weiss (ed.), *Multi-Agent Systems*.