Protocols for Negotiating Complex Contracts

Mark Klein and Peyman Faratin, *Massachusetts Institute of Technology*

Hiroki Sayama, *University of Electro-Communications*

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

Most real-world contracts are complex, with many possible combinations of interdependent issues. These negotiation protocols, based on simulated annealing, achieve near-optimal outcomes for negotiations with binary issue dependencies.

Research to date on negotiation protocols has focused almost exclusively on defining simple contracts consisting of one or a few independent issues and a relatively small number of possible contracts. Many real-world contracts, in contrast, are much more complex, consisting of multiple interdependent issues and intractably large contract spaces. The family of negotiation protocols we've developed make substantial progress toward achieving near-optimal outcomes for negotiations with binary issue dependencies. We propose a simulated-annealing-based approach, a refined version based on a parity-maintaining annealing mediator, and an unmediated version of the negotiation protocol.

Simple and complex contracts

Negotiation protocols work, in general, via the iterative exchange of proposals and counter-proposals. An agent starts with a contract that's optimal for that agent and makes concessions, in each subsequent proposal, until either an agreement is reached or the negotiation is abandoned because the latest proposal's utility has fallen below the agents' *reservation value*—that is, the minimum level of contract utility that the agent will accept.

Figure 1 shows the proposal exchange model of negotiation, applied to a simple contract. The *y*-axis represents a contract's utility to each agent. Each point on the *x*-axis represents a possible contract, ordered in terms of its utility to agent B. Because there's no need to negotiate over issues that both parties agree on, we consider only issues where improvement for one party represents a decrement for the other. The arrows represent how agents begin with locally optimal proposals and concede toward each other, with their subsequent proposals, as slowly as possible. 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 approach is perfectly reasonable for simple contracts. Because issues are independent, a contract's utility 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 shown in Figure 2. In this example, the contract consists of 40 binary issues. Each agent starts with a locally optimal proposed contract (at the extremes of the Pareto frontier, representing the set of optimal contracts) and is required to reduce the Hamming distance (the number of issues with different values) between the two agents' proposals until the agents reach an agreement. With simple contracts, this results in optimal outcomes. We estimated the Pareto frontier using the standard technique of applying an annealing optimizer to differently weighted sums of the two agents' utility functions.

The proposals from each agent start at the agents' separate ideals and then track the Pareto frontier until they meet in the middle with an optimal agreement. This happens because, with linear utility functions, an agent can easily identify the proposal that represents the minimal concession: the contract that's 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

1094-7167/03/$17.00 © 2003 IEEE
Published by the IEEE Computer Society

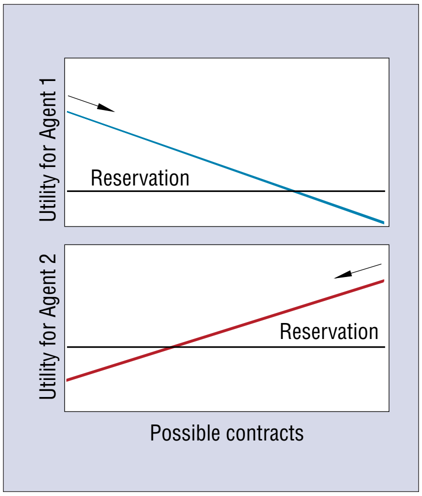

Figure 1. The proposal exchange model of negotiation, applied to a simple contract.

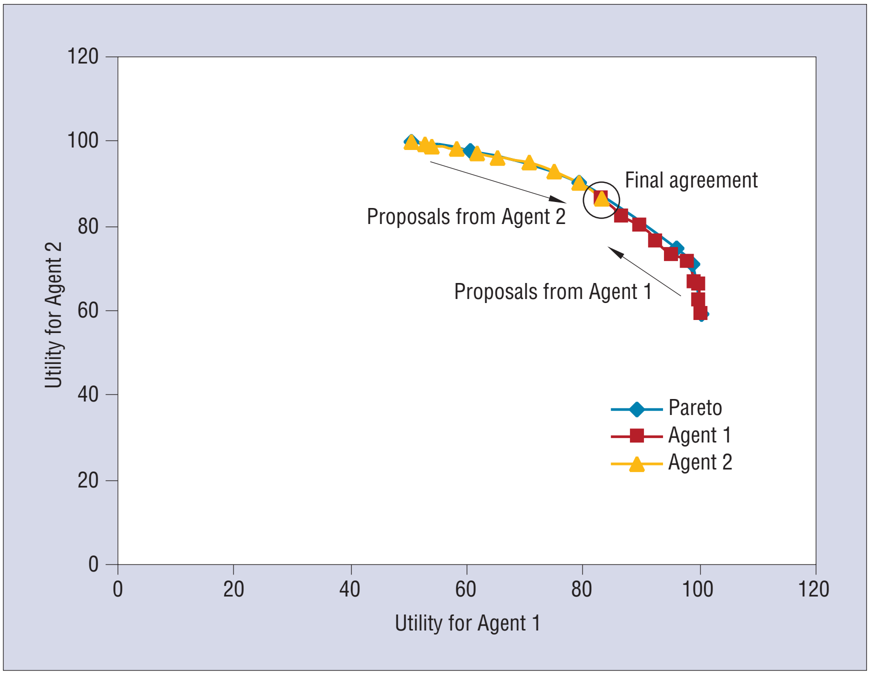

Figure 2. The utilities for the proposals made in a typical simple contract negotiation. The contract consists of 40 binary issues. Each agent starts with a locally optimal proposed contract and must reduce the number of issues with different values between the two proposals until the agents reach an agreement.

utility slope. The utility functions' simplicity, moreover, lets agents 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 interdependent issues. A typical contract might 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. For example, the value to me of a given couch depends on whether it's a good match with the chair I plan to purchase with it.

As Figure 3 shows, such issue interdependencies lead to nonlinear utility functions with multiple local optima.[1] 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 (for example, contract C in Figure 3).

Figure 4 shows how agents behave in complex contract negotiations using standard negotiation techniques. The agents start with an approximation to their ideal contract and diverge increasingly from the Pareto frontier as they converge on an agreement. As you can see, the minimal concession protocol that works optimally for simple contracts produces substantially suboptimal outcomes for complex contracts. The degree of suboptimality depends on the details of the utility function. In our experiments, for example, the final contracts averaged 94 percent of optimal. This is a substantial decrement when you consider that the utility functions we used for each agent were, individually, easy to optimize: a simple steepest-ascent search averaged final utility values roughly 97 percent of those reached by a nonlinear optimization algorithm. It's striking that such relatively forgiving multi-optima utility functions lead to substantially suboptimal negotiation outcomes.

These suboptimal 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 might 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, because of the contract space's sheer size. Also, because the utility functions are quite complex, it's no longer practical for one agent to infer the other's utility function. Complex contracts therefore require different negotiation techniques that let agents find win-win contracts

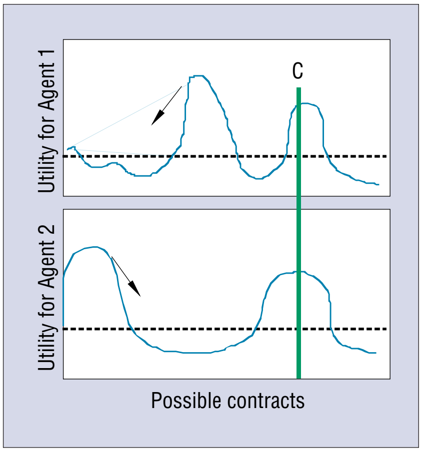

Figure 3. Proposal exchange applied to a complex contract. Because of issue interdependencies, the utility functions have multiple optima. The arrows show what happens when each agent begins at a local optimum and concedes toward the other: they can miss win-win solutions (such as that represented by contract C) found elsewhere in the contract space.

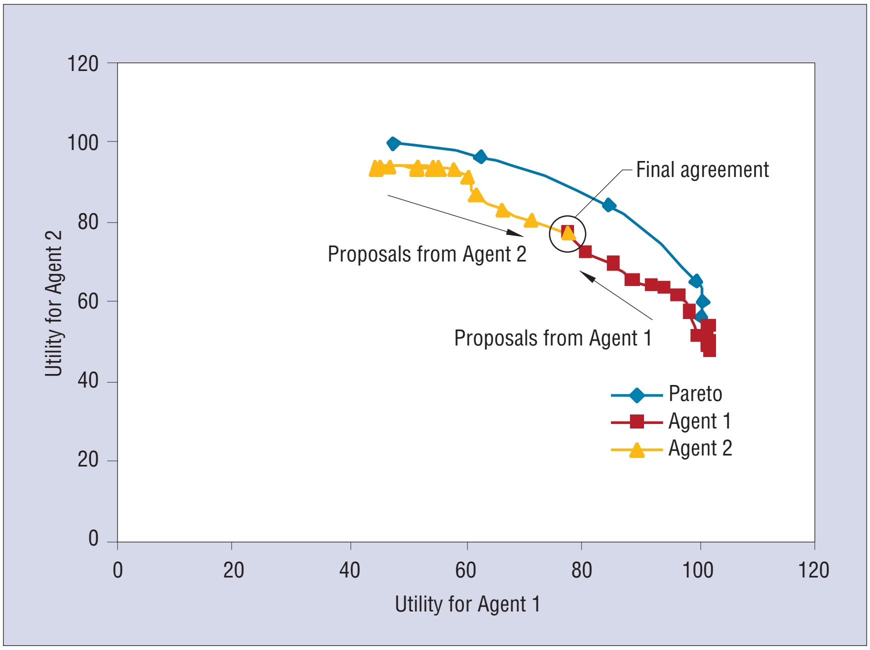

Figure 4. The utilities for the proposals made in a typical complex contract negotiation. This example differs from Figure 2 only in that each agent is using a nonlinear utility function.

in intractable multioptima search spaces in a reasonable amount of time.

Mediated single-text negotiation

A standard approach for dealing with complex negotiations in human settings is *mediated single-text negotiation*.[2] In this process, a mediator proposes a contract that's then critiqued by the parties in the negotiation; the mediator then generates a new, hopefully better proposal based on these responses. This process continues, generating successively better contracts, until some agreed-on stopping point (for example, the reservation utility value is met or exceeded for both parties).

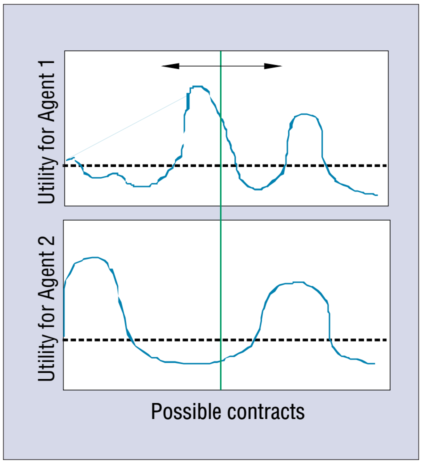

Figure 5. Single-text negotiation. The vertical line represents the current proposed contract; subsequent proposals move that line in the contract space.

Figure 5 illustrates this process. 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's sufficiently good for both parties.

We defined a simple simulation experiment to help us explore how well this approach actually works. In this experiment, two agents negotiated to find a mutually acceptable contract consisting of a vector **S** of 100 Boolean-valued issues. We assigned each issue the value 0 or 1 corresponding to a given contract clause's absence or presence. 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 influence matrix $\boldsymbol{H}$, wherein each cell represents the utility increment or decrement caused by the presence of a given pair of issues, and a contract's total utility is the sum of the cell values for every issue pair in the contract:

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

The influence matrix therefore captures the bilateral dependencies between issues, in addition to any individual contract clause's value. For our experiments, we initialized the utility matrix to have random values between −1 and +1 in each cell. We used a different influence matrix for each simulation run to ensure that our results weren't idiosyncratic to a particular configuration of issue interdependencies.

The mediator proposes a contract that's 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 repeats. If one or both agents vote to reject, the mediator proposes a mutation of the most recent mutually accepted contract instead. The process continues for a fixed number of proposals. We can extend this approach straightforwardly to an *N*-party (multilateral) negotiation, because we can have any number of parties voting on the contracts.

We defined two kinds of agents: *hill-climbers* and *annealers*. Hill-climbers use a 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 the last accepted one with the

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 final contract's social-welfare value is to optimal. The pair of values below it represent how close the final contract is to optimal for Agents 1 and 2, respectively.

	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

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. An annealer's virtual temperature gradually declines over time so eventually its behavior becomes indistinguishable from that of 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 Table 1 shows, there's a catch.

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, because finding many contracts that represent an improvement for both parties is difficult. Figure 6a shows the utilities for the accepted proposals in a typical negotiation with two hill-climbers. In this case the mediator could find only a handful of contracts that increased the utility for both hill-climbers and ended up with a poor final social welfare (sum of the utilities a contract provides for the agents involved) far short of the Pareto frontier.

Near-optimal social welfare can be achieved, in contrast, when both agents are annealers, willing to initially accept individually worse contracts so that they can find win-win contracts later on. Figure 6b shows an example of this, in which the agents enter-

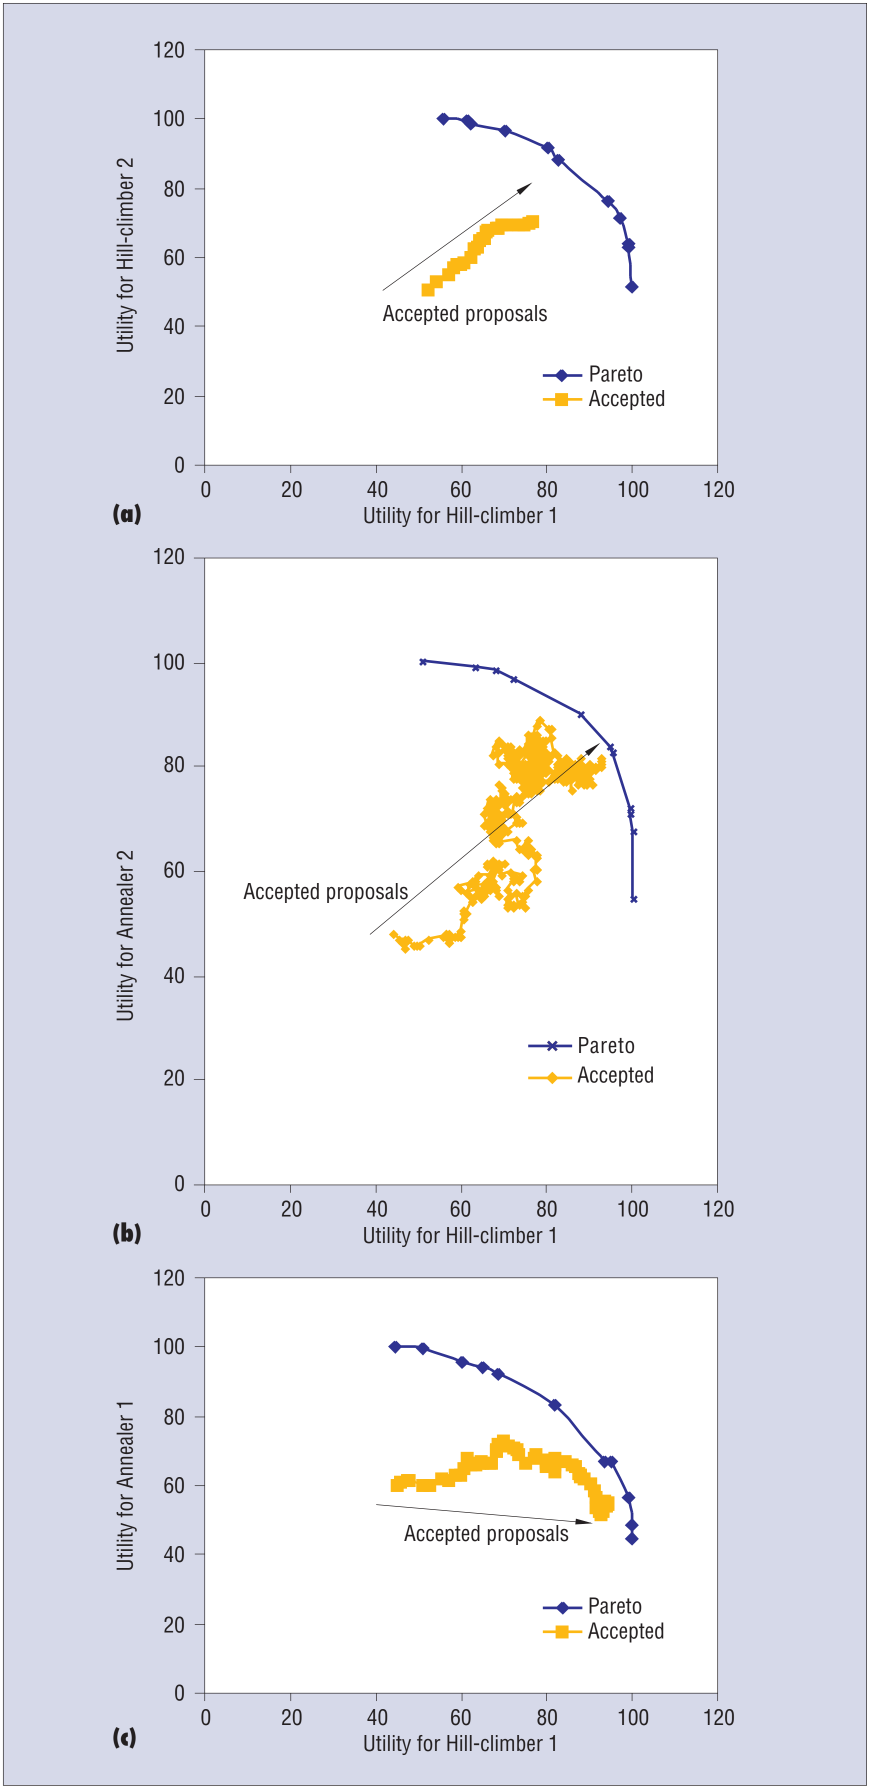

Figure 6. The utilities for the accepted proposals in a typical single-text complex contract negotiation (a) With two hill-climbers: the mediator's initial proposal is at the lower left, and the subsequent accepted proposals move toward higher utilities for both agents. (b) With two annealers: some accepted proposals actually cause utility decrements for one or both agents, but the final result is a near-optimal contract. (c) With an annealer and a hill-climber: the hill-climber achieves a near-optimal contract at the annealer's expense.

Table 2. The optimality of the negotiation outcomes for truthful versus exaggerating agents with a simple annealing mediator. An exaggeration strategy is individually rational, even though it results in outcomes with lower social welfare. The top value in each cell represents how close the final contract's social-welfare value is to optimal; the two values in each cell represent how close the final contract is to optimal for Agents 1 and 2, respectively.

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

tain 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 (see Figure 6c). Why? When an annealer is at a high virtual temperature, it becomes a chronic conceder, accepting almost anything, and thereby pays a "conceder's penalty." The hill-climber "drags" the annealer toward its own local optimum, which is unlikely to also be optimal for the annealer.

This reveals a dilemma. In negotiation contexts, we typically can't assume that agents will be altruistic, so we must 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 (that is, it increases social welfare), it isn't an *individually* dominant strategy. Hill-climbing is dominant because no matter what strategy the other agent uses, it's better to be a hill-climber (see Table 1). If all agents do this, however, they forego the higher individual utilities they would get if they both annealed. Individual rationality thus drives the agents toward the strategy pairing with the lowest individual and social welfare. This is thus an instance of the Prisoner's Dilemma.

Researchers have shown that we can avoid this dilemma if we assume repeated interaction between agents,[4] but we would prefer to have a negotiation protocol that makes it individually rational to engage in 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 annealers to hill-climbers if they determine, by observing their opponents' proposal acceptance rates, that their opponents are being hill-climbers. We found, however, that it takes too long to determine the other agent's type. By the time it has become clear, much of the contract utility has been committed, and it's too late to recover from the consequences of having started out as an annealer.[5]

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 that both negotiating agents accepted. Our refined protocol endows the mediator 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.

In our initial implementations, each agent gave a simple accept or reject vote for each proposal from the mediator, but this resulted in final social-welfare values significantly lower than what we earlier achieved using annealing agents. In the next round of experiments, we modified the agents so that they provide additional information to the mediator in the form of vote strengths: each agent annotates their accept or reject vote as being strong or weak. The agents are 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. The mediator accepts a proposal if the score is non-negative—that is, if both agents vote to accept it or if a strong accept by one agent overrides a weak reject from the other. The mediator can also accept rejected contracts (those with a negative aggregate score) using the annealing scheme described earlier. This approach works surprisingly well, achieving final social-welfare values that average roughly 99 percent of optimal even though the agents give the mediator only two bits of information. We found, in fact, that increasing the number of possible vote weights doesn't increase final social welfare. This is because the strong and weak vote annotations are sufficient to allow the system to pursue social-welfare-increasing contracts that cause a utility decrement for one agent.

Incentives for truthful voting

Any voting scheme introduces the potential for strategic nontruthful voting by the agents, and our scheme is no exception. Imagine that one agent always votes truthfully, while the other exaggerates so that its votes are always "strong." As you might expect, this would bias negotiation outcomes to favor the exaggerator (see Table 2).

As you can see, even though exaggerating substantially decreases social welfare, it is individually rational to do so, thus recreating the Prisoner's Dilemma we encountered earlier. The underlying problem is simple: exaggerating agents can induce the mediator to accept proposals that are advantageous to them (if the other agent weakly rejects them), while preventing the other agent from doing the same. So, we need an enhancement to the negotiation protocol that motivates truthful voting, preserves equity, and maximizes social welfare.

However, simply limiting the number of strong votes each agent can use doesn't work. If the limit is too low, we effectively lose the benefit of vote weight information, ending up with lower social-welfare values. If the strong-vote limit is high enough to avoid this, then all an exaggerator has to do is save all its strong votes until the end of the negotiation, at which point it can drag the mediator toward 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. An override occurs when the mediator accepts a contract supported by one agent (the "winner") over the other agent's objections. Overrides drag a negotiation toward contracts favorable to the winner, so it makes sense to make the total number of overrides equal for each agent. But this isn't enough, because exaggerators always win disproportionately more than truth-tellers do.

The solution, we found, came from enforcing a running parity between the number of overrides given to each agent throughout the negotiation, so that neither agent can get more than a given advantage. This approach at least maintains rough equity no matter when (or whether) either agent chooses to exaggerate. Table 3 shows the results of this approach when the override disparity is limited to 3. The parity-enforcing mediator makes being truthful the individually rational strategy.

Table 3. The optimality of the negotiation outcomes for truthful versus exaggerating agents with a parity-enforcing mediator. The top value in each cell represents how close the final contract's social-welfare value is to optimal; the two values in each cell represent how close the final contract is to optimal for Agents 1 and 2, respectively.

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

When agents are truthful, we find that this approach achieves social welfare just slightly below that achieved by a simple annealing mediator, while offering a significantly ($p < .01$) 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 randomly, or exaggerating just near the negotiation's end. Being truthful 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? 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 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 an override even when the utility increment it derives is relatively small. At the end of the day, the truthful agent has spent its tokens more wisely and to better effect.

The unmediated single-text protocol

The protocol that we've just considered worked well in the contexts we studied but has the disadvantage of requiring a mediator. One issue concerns trust. Because the annealing mediator is empowered to selectively ignore agent votes, it might do so in a way that favors one agent over another (although the parity-enforcing token mechanism does somewhat reduce this problem's potential impact).

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. So, the mediator generates a high proportion of rejected contracts, which is partly why our experimental runs each involved so many (2,500) proposals. The negotiating agents could provide the mediator with information about their utility functions so that the mediator could propose contracts more "intelligently." However, this is problematic for several reasons, including the typical reluctance of self-interested agents to reveal their utility functions to a party that might not be worthy of their trust.

Fortunately, we can define an effective unmediated version of the annealing protocol. Agents each start with a given number of tokens (two each, in our experiments) and a mutually agreed-on starting temperature T. A random contract is generated, and one negotiating agent is randomly selected to propose a small (single-issue) variant of the contract—presumably the variant that most increases the contract's utility for that agent. The other agent then votes on the proposed variant. The proposals and votes indicate the strength of the agents' preference for the proposed contract using the scheme described earlier (that is, strong reject, weak reject, weak accept, strong accept). The contract is provisionally accepted with the probability

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

where the aggregate score (U) is calculated as for the annealing mediator, and the outcome is determined using the roll of fair, mutually observable dice. If the decision to accept a proposal represents the override of one agent's reject vote, the winning agent needs to give one of its tokens to the overridden agent. An override isn't 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. Proposers may pass but may not repeat proposals. The temperature T declines at a mutually agreed-on rate during this process. This protocol thus reproduces the key elements of the annealing mediator protocol—a time-dependent annealing regime plus tokens—without requiring a mediator. Our experiments show that this protocol produces results just as good as the annealing mediator (averaging 99 percent of optimal) while requiring fewer proposal exchanges (averaging about 200 exchanges per negotiation).

Contributions

This article presents, as far as we are aware, the first negotiation protocol specifically for complex contracts. Although some researchers have studied multi-issue negotiation,[6–8] they treated the issue utilities as independent, so each agent's utility functions were linear, with single optima. As we have seen, however, introducing 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 try to maximize the utility received by the buyer. If no bid is satisfactory, the buyer modifies its published utility function and tries again. This introduces a search process. This approach's problem is that it doesn't provide any guidance for how the parties involved should control their search through the vast space of possibilities.

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 a key to negotiating good complex contracts. Conceding isn't individually rational in the face of agents that might choose not to concede, but we can resolve this problem by either introducing a mediator that stochastically ignores agent preferences or introducing dice into the negotiation protocol. In both cases, we can use the exchange of tokens during an override to encourage the truthful voting that enables win-win outcomes.

There are many other promising avenues for future work in this area. The high social welfare achieved by our approach partially reflects the fact that the agent's utility functions, 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] To address this challenge, adapting nonlinear optimization techniques such as genetic algorithms into the negotiation context might be necessary.

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 depend heavily on each other, letting the mediator focus attention on tightly coupled issue clumps and ignore other less influential issues until later. Agents might be encouraged to tell the truth about this to ensure that negotiations can complete in an acceptable amount of time.

Finally, we'd like to derive formal incentive compatibility proofs (that is, concerning when agents are encouraged 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 don't hold with complex contracts. ■

Acknowledgments

Funding from the DARPA Control of Agent-Based Systems (CoABS) program and the National Science Foundation's Computation and Social Systems program supported this work.

References

1. Y. Bar-Yam, *Dynamics of Complex Systems*, Addison-Wesley, 1997, pp. xvi, 848.
2. H. Raiffa, *The Art and Science of Negotiation*, Belknap Press of Harvard Univ. Press, 1982, pp. x, 373.
3. J.S. Rosenschein and G. Zlotkin, *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*, MIT Press, 1994, pp. xxi, 229.
4. R. Axelrod, *The Evolution of Cooperation*, Basic Books, 1984.
5. M. Klein, P. Faratin, and Y. Bar-Yam, "Using an Annealing Mediator to Solve the Prisoners' Dilemma in the Negotiation of Complex Contracts," *Proc. Agent-Mediated Electronic Commerce Workshop* (AMEC-IV), Springer-Verlag, 2002.
6. P. Faratin, C. Sierra, and N.R. Jennings, "Using Similarity Criteria to Make Negotiation Trade-Offs," *Proc. 4th Int'l Conf. Multi-Agent Systems*, IEEE CS Press, 2000.
7. C.M. Jonker and J. Treur, "An Agent Architecture for Multi-Attribute Negotiation," *Proc. Int'l Joint Conf. Artificial Intelligence* (IJCAI 01), Morgan Kaufmann, 2001, pp. 1195–1201.
8. S.S. Fatima, M. Wooldridge, and N.R. Jennings, "Multi-issue Negotiation under Time Constraints," *Proc. 2002 Int'l Conf. Autonomous Agents and Multi-Agent Systems* (AAMAS 02), ACM Press, 2002.
9. M. Bichler and J. Kalagnanam, "Bidding Languages and Winner Determination in Multi-Attribute Auctions," to be published in *European J. Operational Research*.
10. S.A. Kauffman, *The Origins of Order: Self-Organization and Selection in Evolution*, Oxford Univ. Press, 1993.

The Authors

Mark Klein is a principal research scientist at the MIT Center for Coordination Science and an affiliate at the MIT AI Lab and the New England Complex Systems Institute. His research focuses on better understanding coordination and how it can help solve important problems concerning the design and management of human organizations and software systems. He received his PhD in computer science from the University of Illinois. He is an IEEE member. Contact him at the Center for Coordination Science, Massachusetts Inst. of Technology, 77 Massachusetts Ave., NE 20-336; Cambridge, MA 02142-1347; m.klein@mit.edu.

Peyman Faratin is a postdoctoral associate at MIT's Laboratory for Computer Science. He does research on negotiation protocols and algorithms for multiagent systems. He received his PhD in computer science from the University of London. He is an IEEE member. Contact him at the Laboratory of Computer Science, MIT, 77 Massachusetts Ave., NE43-534, Cambridge, MA 02139; peyman@mit.edu.

Hiroki Sayama is an assistant professor in the Department of Human Communication at the University of Electro-Communications in Japan. He conducts research on the emergent dynamics of complex systems, including evolutionary and collaborative multiagent systems. He received his DSc in information science from the University of Tokyo. He is an IEEE member. Contact him at the Dept. of Human Communication, Univ. of Electro-Communications, 1-5-1 Chofugaoka, Chofu, Tokyo 182-8585, Japan; sayama@hc.uec.ac.jp.

Yaneer Bar-Yam is president of the New England Complex Systems Institute. He is also chairman of the International Conference on Complex Systems, managing editor of *InterJournal*, and author of *Dynamics of Complex Systems* (Perseus Books, 1997). He studies patterns of collective behavior, evolution, complexity, and multiscale representations in application to physical, biological, social, and engineered systems. He received his PhD in physics from MIT. Contact him at the New England Complex Systems Inst., 24 Mt. Auburn St., Cambridge, MA 02138; yaneer@necsi.org.