Negotiating Complex Contracts

Mark Klein Massachusetts Institute of Technology m_klein@mit.edu Peyman Faratin Massachusetts Institute of Technology peyman@mit.edu

ABSTRACT

Work to date on computational models of negotiation has focused almost exclusively on defining contracts consisting of one or a few independent issues. Many real-world contracts, by contrast, consist of multiple inter-dependent issues. This paper describes a simulated annealing based approach appropriate for negotiating such complex contracts, evaluates its efficacy, and suggests potentially promising avenues for future work.

Categories and Subject Descriptors TBD

General Terms

Algorithms

Keywords

Nonlinear 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 [1] [2]. Real-world contracts, by contrast, are generally much more complex, consisting of a potentially large number of issues that are *inter-dependent*, i.e. where the value of one issue selection to an agent often depends on the selections made for other issues. The value to me of a given DVD player, for example, depends on whether it is a good match with the tuner and speakers I plan to purchase with it.

Does the introduction of issue interdependencies impact which negotiation algorithms are most appropriate? Let us explore this question in the context of mediated single text negotiation, often used for complex negotiations in human settings [4]. 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 continues until all parties are satisfied or abandon the negotiation.

In our experiments, the mediator helps two agents find a mutually acceptable contract consisting of a vector C of 40 issues, each issue assigned the value 0 or 1 corresponding to the presence or absence of a given contract clause. This defines a space of 2^40 , or roughly 10^12 , possible contracts. Each agent accepts or rejects contracts based on its utility, calculated using that agent's 40x40 influences matrix H, wherein each cell represents the utility impact (positive or negative) of the presence of a given pair of issues: the total utility of a contract is the sum of the cell values for every issue pair present in the contract:

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

$$U = \sum_{i=1}^{40} \sum_{j=1}^{40} H_{ij} C_i C_j$$

The influence matrix therefore captures binary dependencies between issues, in addition (in the diagonal cells) to the value of any individual contract clause.

To assess the impact of dependencies we measured the social welfare optimality of contracts produced by hill climber agents working with a mediator, with and without issue dependencies, with the following results:

Mediator Type	Independent issues	Interdependent issues
Single systematic	.88	.78
Double systematic	.98	.87

If issues are independent (i.e. if the non-diagonals of the influences matrices H are all zero), then it is reasonable to simply have the mediator systematically enumerate every possible value (0 or 1) for each issue (the "single systematic" mediator), and have the agents be 'hill climbers', i.e. accept only contracts whose utility is better than the last contract they both accepted. The end result will be contracts in which only mutually acceptable clauses are included. This produces contracts however, whose social welfare averages only 88% of optimal, because we miss opportunities for agents to exchange mutually beneficial concessions. Such tradeoffs can be enabled if the mediator also proposes every possible pair of issue values (the "double systematic" mediator). With this improved procedure we get contracts whose utility averages 98% of optimal. Hill climbing is thus an effective agent strategy with independent issues, given the use of an appropriate mediator.

If we introduce issue inter-dependencies, however, (by populating the non-diagonal cells of the influences matrices with non-zero values) the average contract optimality drops 10% or more. Why does this occur? When issues are independent, the utility functions for an agent will be *linear*, defining a contract space with a single optimum[3]. In such a context hill-climbing is known to be effective. When interdependencies are introduced, however, this leads to *nonlinear* utility functions with multiple optima [3], with the result that hill climbers can get caught on sub-

optimal solutions. Consider a plot of the utility to each agent of the contracts accepted in a typical negotiation, graphed next to the pareto efficient line:



Figure 1. A typical negotiation with two hill climbers

If both agents hill-climb they tend to get stuck in local optima far short of a pareto-efficient contract.

We could of course solve this problem by having the mediator simply propose every possible contract to hillclimbing agents, but with complex contracts, as we have seen, the number of possible contracts is typically intractably large. Our challenge, therefore, is to define negotiation techniques that which allow agents to find 'win-win' contracts in intractably large multi-optima search spaces in a reasonable amount of time.

2. AN ANNEALING APPROACH

We adapted for this purpose a well-known nonlinear optimization technique known as 'simulated annealing' [5]. A set of contracts is generated by the mediator, each one representing a random change to the last accepted contract. Contracts that offer improved utility are always accepted. In addition, each annealer agent has a virtual temperature T such that it accepts lower utility contracts with probability:

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

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 annealers can skip utility valleys on the way to higher optima [3]. We hypothesized that annealers would prove superior to hill climbers in complex contract negotiations.

To test this we compared annealers with hill climbers for interdependent issue contracts. The mediators proposed 2500 contracts, during which period the annealers gradually cooled to a temperature of zero. The results were as follows:

	Agent 2 anneals	Agent2 hill-climbs
Agent 1 anneals	[.86]	[.86]
	.73/.74	.51/.99
Agent1 hill-	[.86]	[.98]
climbs	.99/.51	.84/.84

where the cell values are laid out as follows:

[<social welfare optimality>] <agent 1 optimality >/<agent 2 optimality >

The results reveal that paired annealer agents perform nearoptimally, significantly better than paired hill climbers. That is because the annealers, in contrast to hill climbers, are willing to (temporarily) accept individually worse contracts to help find winwin contracts later on:



Figure 2: A typical negotiation with two annealers.

If an annealer is paired with a hill climber, however, the annealer is 'dragged' by the hill-climber towards the hill climber's optimum, which is unlikely to also be optimal for the annealer:



Figure 3: A typical negotiation with an annealer and hill climber.

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 [6] [7] [8]. In other words, we want the socially most beneficial strategy to also be the individually *dominant* one so that most agents will tend to use it. In our case, however, even though annealing is a socially dominant strategy (i.e. annealer always increase social welfare), *annealing is <u>not</u> 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 than an annealer. 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 [9].

Further analysis reveals that there is no way to avoid this dilemma within a single negotiation. If both agents could know ahead of time what strategy the other agent is going to use, then both agents would select annealing. In an open system environment we can not rely on self-reports for this, since agents are incented to claim they will use annealing 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. Annealers accept a much higher percentage of proposed contracts than hill-climber, especially at first:



Figure 4: Proposal acceptance for hill-climbers and annealers.

The problem with this 'adaptive' approach is that determining the type of an agent based on its voting behavior takes time. But as figure 4 shows, the divergence in acceptance rates between annealers and hill-climbers only becomes clear after at least 100 proposal exchanges or so. By this time, however, *much of the contract utility has already been committed*:



Figure 5: Contract utility over time for a typical negotiation.

so it is too late to fully recover from the consequences of having guessed wrong. In our experiments, for example, between 40% and 60% of the final social welfare had already been achieved in the first 100 proposal exchanges. While adaptive strategies reduce the magnitude of the penalty paid by annealers paired with hill climbers, they can not eliminate it.

Another strategy for reducing the annealers' penalty is for the annealer agent to start at a lower temperature, so that it can not be dragged as far from its own optimum:



Figure 6: Individual Utilities as a Function of Annealer Agent Starting Temperature.

As we can see, if the annealer agent starts at a low enough temperature, its penalty relative to a hill climber can be eliminated, but only at the cost of contracts with lower overall social welfare. This approach also faces the difficulty that it is only possible to determine the appropriate starting temperature empirically, which may be impractical as it requires enacting many repeated negotiations with the same utility functions, but in real life contexts we rarely do so.

Finally, the results show that *hill-climbers reach stability sooner than annealers*. The hill climbers typically reached stability after roughly 100 proposal exchanges, while the annealers approached stable utility values after roughly 800 proposal exchanges. 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, 'hop' among 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.

3. 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 algorithms. This paper presents, as far as we are aware, the first computational negotiation approach suited for multiple interdependent issues. The essence of the 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, a result that is relevant to any collaborative decision making task involving interdependent decisions.

4. NEXT STEPS

4.1. Unmediated Protocols

These same effects probably apply, we believe, to unmediated negotiation protocols, i.e. that involve direct proposal exchanges between agents. We can frame what these protocols do as follows:



Figure 7: Proposal exchanges with independent issues.

Each point on the X axis represents a candidate contract The Y axes represents the utility of each contract to each agent. Both agents have a reservation utility value: only contracts whose utility is above that agent's reservation value will be accepted. 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 mapping a candidate contract to its utility for an agent are linear. 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.

Now consider what happens if we introduce issue dependencies, and the resulting multi-optima utility functions, for each agent:



Figure 8. Proposal exchanges with issue dependencies.

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 8 above, 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 such '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.

One potentially important line of work, therefore, is to explore strategies suitable for direct proposal exchanges with interdependent issues.

4.2. Faster Negotiations

The simulated annealing approach produces better social welfares than hill-climbing but involves larger numbers of proposal exchanges. What can we do about this?

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 we are exploring is 'genetic annealing', which uses abstracted measures of the issue inter-dependency structure to cluster highly-interdependent issue sets into 'genes' that are recombined al la sexual reproduction to more quickly explore the large search spaces involved in contingent contract design.

Introducing (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 hillclimbers *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 behavior, but with inter-dependent issues and large multiple-optima utility functions this becomes intractable and information exchange probably must be done explicitly.

4.3. Addressing the Prisoner's Dilemma

We have shown that there is no way to entirely avoid the prisoner's dilemma within the scope of a single negotiation, though we can reduce the magnitude of the problem using adaptive strategies and cold annealing. Previous work on iterated games [9] has shown however that prisoner's dilemma games, such as the one that emerges in our case, will result in agents choosing the more socially beneficial concession strategy if the games are repeated, i.e. if a given pair of agents engages in multiple negotiations, and the agents take into account what happened with previous negotiations, e.g. conceding only if the other agent has a history of conceding as well ("tit for tat"). In large agent societies, agents may only rarely have a previous negotiation history with each other, but this problem can be resolved through the use of reputation mechanisms that pool reported negotiation experiences over all agents. We would then of course have to account for the possibility of reputation sabotage [10]. Adaptive strategies are a good complement to reputation mechanisms since they reduce the negative consequences of getting misleading reputation information. Another tack is for contractor agents to negotiate with several subcontractors and select the best contract. This will increase the incentive for agents to be annealers, since chronic "tough guys" may find themselves without customers.

ACKNOWLEDGEMENTS

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

REFERENCES

- 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.
- Ehtamo, H., E. Ketteunen, and R. Hamalainen, *Searching for Joint Gains in Multi-Party Negotiations*. European Journal of Operational Research, 2001. 1(30): p. 54-69.
- Bar-Yam, Y., Dynamics of complex systems. 1997, Reading, Mass.: Addison-Wesley. xvi, 848.
- 4. Raiffa, H., *The art and science of negotiation*. 1982, Cambridge, Mass.: Belknap Press of Harvard University Press. x, 373.
- Kalos, M.H. and P.A. Whitlock, *Monte Carlo methods*. 1986, New York: J. Wiley & Sons.
- 6. 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.

- Kramer, J. and J. Magee. Analyzing Dynamic Change in Software Architectures. in ICSE98 Workshop on Principles of Software Evolution. 1998. Kyoto, Japan.
- 8. Sandholm, T.W., *Distributed Rational Decision Making*, in *Multi-Agent Systems*, G. Weiss, Editor. 1998.
- 9. Axelrod, R., *The Evolution Of Cooperation*. 1984: Basic Books.
- 10. Dellarocas, C. Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. in Proceedings of the 2nd ACM Conference on Electronic Commerce. 2000. Minneapolis, MN.