Hybrid Computational Models for the Mediated Negotiation of Complex Contracts

Gregory Todd Jones
Georgia State University College of Law
Interuniversity Consortium on Negotiation and Conflict Resolution
Max Planck Institute for Research on Collective Goods
gtjones@gsu.edu

Abstract

The negotiation of even the most straightforward real-world contracts tends to be quite complex. A contract with only 25 distinct issues with two alternatives each presents the parties with more than 33 million possible contracts, far too many to be evaluated exhaustively within feasible time constraints. Furthermore, contract issues that exhibit high levels of interdependence result in highly nonlinear utility functions with the possibility of many local optima. This paper employs hybrid computational models, integrating both simulated annealing and tabu list optimization, to aid in the design of social heuristics and institutional mechanisms that may serve to improve the effectiveness of human negotiators.

Contact:
Gregory Todd Jones
Interuniversity Consortium on Negotiation and Conflict Resolution
Georgia State University College of Law
Atlanta, GA 30303

Tel: 404-651-2377
Fax: 404-463-9789
Email: gtjones@gsu.edu

Key Words: Automated Negotiation; Complex Contracts; Simulated Annealing; Tabu Lists

Acknowledgement: Much of this work was completed while a Visiting Research Scholar at the Max Planck Institute for Research on Collective Goods, Bonn, Germany. Their kind hospitality is gratefully acknowledged. I would like to thank Doug Yarn, Reidar Hagtvedt, Christoph Engel, Arndt Bröder, Martin Beckenkamp, Stephan Tontrup and the members of the Max Planck Institute for Research on Collective Goods Rationality Working Group for stimulating discussions and many helpful comments.

Hybrid Computational Models for the Mediated Negotiation of Complex Contracts

Gregory Todd Jones

The negotiation of even the most straightforward real-world contracts tends to be quite complex. A contract with only 25 distinct issues with two alternatives each presents the parties with more than 33 million possible contracts,[1] far too many to be evaluated exhaustively within feasible time constraints. Further, while most work related to the analysis of negotiation assumes these issues to be independent and therefore the utility functions used in evaluating possible contracts to be linear [Cheng, Chan & Lin, 2005; Faratin, Sierra, & Jennings, 2002; Ehtamo, Kettunen, & Hamalainen, 2001], it is more typical that contract issues exhibit high levels of interdependence that result in nonlinear utility functions with the possibility of multiple local optima [Klein, Faratin & Bar-Yam; Bar-Yam, 1997; for a thorough treatment of multi-attribute utlity, see generally Keeny & Raiffa, 1993]. As a consequence, even simple negotiations frequently result in sub-optimal, Pareto-inferior agreements [Raiffa, 2003].

At their essence, all but the most trivial negotiations introduce a search problem not unlike those faced by computer scientists when combinatorial explosion forces the consideration of a constrained set of possible solutions to complex problems.[2] This paper seeks to address the challenges of intractably large contract spaces and utility functions with multiple local optima by relying on two well known computational models for nonlinear optimization, simulated annealing and tabu search optimization. Several scholars, particularly Faratin and Klein [for a representative work among more than a dozen similar publications, see Klein & Faratin, 2003] have applied simulated annealing, a probabilistic algorithm for locating approximations to global optima in large search spaces, to the mediated negotiation of complex contracts. The algorithm owes its name, as well as its substance, to the process of annealing in metallurgy, in which materials are heated and gradually cooled under controlled conditions in order to favorably impact such characteristics as strength and hardness [Kirkpatrick, Gelatt & Vecchi, 1983; Cerny, 1985]. However, with simulated annealing-based negotiation strategies alone, there remains the danger that small sections of the search space can be cycled through, leaving parties trapped in the local optima that we seek to avoid. In response to this difficulty, this paper proposes hybrid models that integrate tabu lists as well. Tabu lists are a simple form of tabu search optimization methodology which address the problems associated with local optima by employing memory structures to force exploration of regions of the search space that may otherwise go unexplored [Glover & Laguna, 1997]. In addition to these novel models, our work can be distinguished from previous research by virtue of its motivation. While in most previous research, including Faratin's and Klein's, "bargaining automated agents are programmed with rules-of-thumb distilled from intuitions about good behavioral practice in human negotiations," [Binmore & Vulkan, 1999, p. 3] with the objective of developing more effective autonomous agents, our objective is just the opposite. We are using computational simulations to aid in the design of social heuristics and institutional mechanisms that will make humans more effective in social exchange.

Defining the Problem of Contract Complexity

The exponential complexity of multi-issue contracts can defy a negotiated agreement in at least two ways. First, when pre-existing circumstances or preliminary negotiations (Contract 1 in Figure 1) place one party (Agent B) at a optima that is above her reservation utility while the other party (Agent A) remains below his reservation utility, and where the first party (Agent B) will only consider alternative proposals that are strictly better than the current proposal, a protocol known as "hill climbing," there is no opportunity to reach alternative contracts (Contract 2) that would be acceptable to both parties.

[1] The actual number is $2^{25} = 33,554,432$.

[2] Indeed, the negotiation problem would be made much worse if there were more than two parties such that the formation of coalitions was possible. The number of possible coalitions that could be formed among 3 individuals is 5, among 5 individuals it is 52, among 10 individuals it is 115,975 and among 20 individuals it is 51,724,156,235,572 [Krippendorff, 1986]. The role of coalitions is not formally dealt with here.

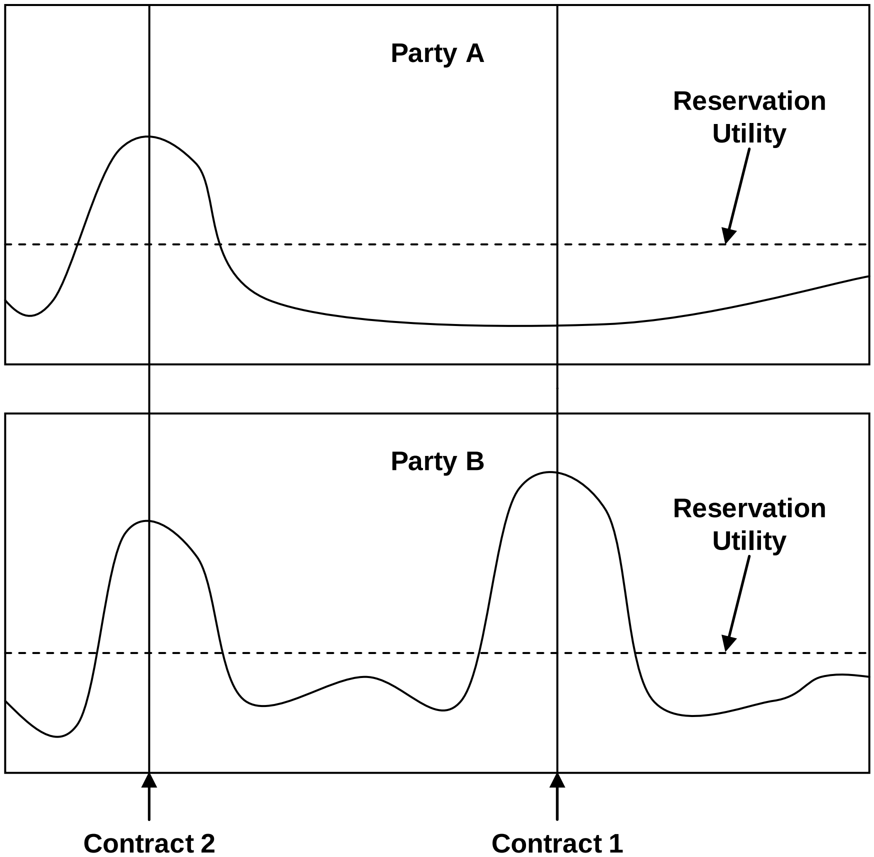

Figure One – Failure to Reach Agreement

Similarly, where both parties find themselves at local optima and both are above their respective reservation utilities (Contract 1 in Figure 2), if even one of the parties engages in hill climbing, pareto superior agreements (Contract 2) can never be reached.

Simulated Annealing

Simulated annealing, a generic probabilistic meta-algorithm for global optimization, deals with this complexity and the prospect of local optima in the following way. Using the context of negotiation to illuminate, we start with an arbitrarily high "temperature" and evaluate alternative contracts by comparing the utility of the alternative with the current contract "on the table." Where the utility of the alternative is greater, this contract becomes the new working solution. Where the alternative represents a reduction in utility, we are more likely to accept the contract if the difference in utility is small and the temperature is high. The "temperature" of the negotiation declines steadily according to some prescribed cooling schedule, which in our simulations is linear. As a result, we are more likely to accept large utility reductions early in the negotiations and less and less likely to accept even small reductions as the negotiations wear on. At the limit, annealers become hill climbers. Where *P(δE, T)* is the probability of accepting a working solution given a difference in utility and a temperature,

$$P(\delta E, T) = 1 \quad \text{where the working solution is better than current solution}$$

$$= e^{-\left[\frac{\delta E}{T}\right]} \quad \text{otherwise}$$

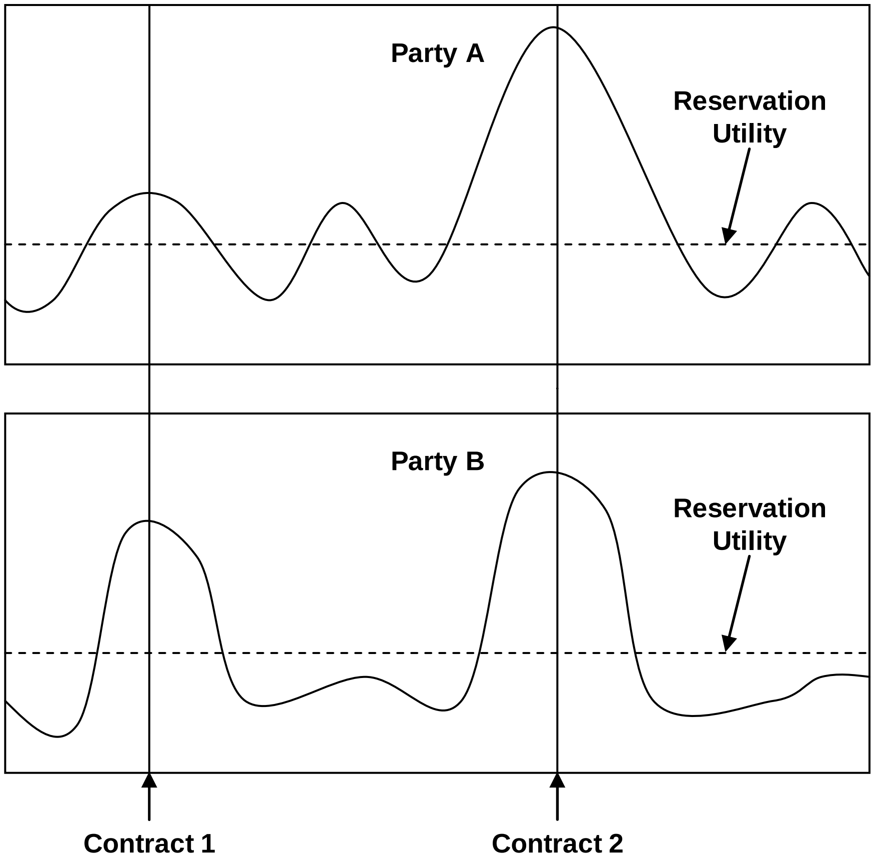

Figure Two – Pareto Inferior Agreement

Tabu Lists

While simulated annealing can make it more likely for a party to accept an inferior contract as a working solution, particularly early in the negotiation, there remains the danger that parties will trade acceptances of proposals that mirror the same issues in cycles that continue until the temperature cools to the point that even these proposals are rejected and the parties are left with no agreement, or at best, a pareto inferior agreement. To address this possibility, our simulations implement hybrid protocols that integrate the very simple idea of tabu lists. A tabu list is a mathematical optimization method employing memory structures. The search procedure moves from solution x to solution x', in the neighborhood of x, $N^*(x)$. The tabu list contains solutions recently attempted (less than n proposals ago, where n is the tabu tenure), and proposals on the tabu list are excluded from $N^*(x)$. Put more simply, once we have considered a particular contract and put it aside, we will not return to the consideration of that contract for a specified number of proposals. This prevents the possibility that parties become trapped in "back-and-forth" proposals that limit consideration to a small area of the solution space.

The Simulations

To explore these computational models, we developed a "in silica" negotiation laboratory (see Figure Three) for a two-party contract negotiation with 25 interdependent issues, presenting the possibility of highly non-linear utility functions and numerous local optima. The software makes it quite easy to simulate complex negotiations with varied parameters for initial temperature, cooling schedule, tabu list length, and the extent to which parties cognitively "bundle" issues together during the course of the negotiations.

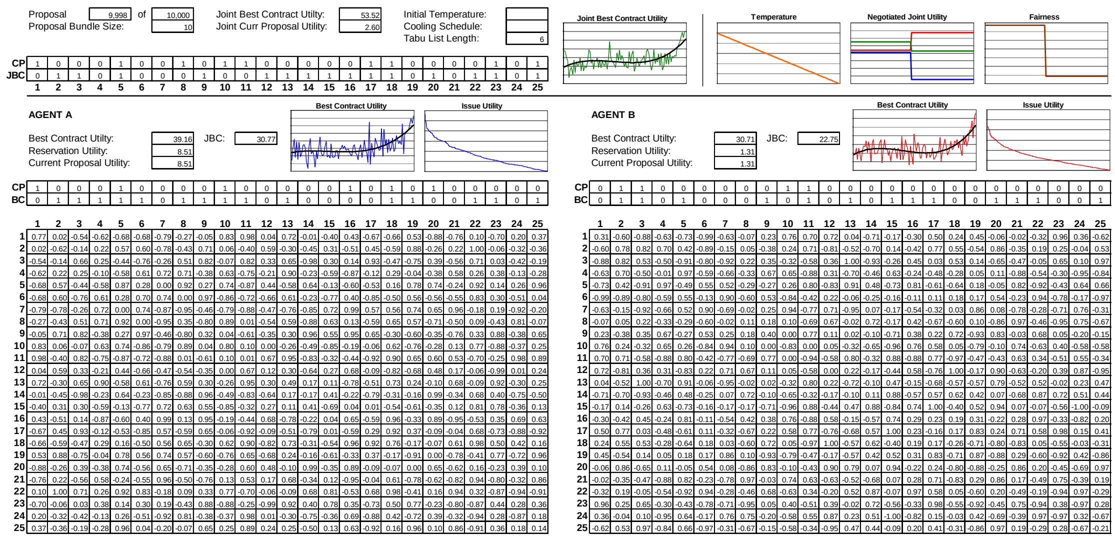

Figure Three – "In Silica" Negotiation Laboratory

Experiments proceed in three stages. First, utility schedules for the 25 issues are independently randomized for each party, with each utility derived from a uniform distribution between 1 and -1 (utilities can be negative as well as positive). Note that the diagonal of these matrices represent the independent utility of each contract issue, while the symmetrical, off-diagonal values represent the interdependent utility of each issue pair. Using an annealing algorithm, the optimal contract for each agent is determined for purposes of comparison and based on modifiable assumptions of salience and cognitive bundling, an initial proposal is generated from Agent A.

Finally, the two agents, each with modifiable negotiation strategies, engage in a negotiation consisting of a pre-determined number of alternating proposals. Figure Four reports mean joint utility and mean agent utilities after 30 negotiations in each of four conditions: 1) A hill climbs and B hill climbs, 2) A employs a hybrid annealing strategy and B hill climbs, 3) A hill climbs and B employs a hybrid annealing strategy, and 4) both A and B employ a hybrid annealing strategy. Rearranging these results (see Figure Five) makes it clear that, like Klein, Faratin & Bar-Yam, our simulations produce a prisoner's dilemma where: 1) hill climbing is dominant, that is, each agent is privately better off hill climbing regardless of what the other agent does, and 2) there is a deficient equilibrium at mutual hill climbing, that is, aggregate social welfare is maximized with mutual annealing and is at its lowest with mutual hill climbing, but there is no individual incentive for annealing.

		% of Optimal Utility		
A	B	Joint	A	B
Hill	Hill	0.435	0.386	0.251
Anneal	Hill	0.556	0.051	0.777
Hill	Anneal	0.694	0.895	0.229
Anneal	Anneal	0.947	0.570	0.704

Figure Four – Hill Climbing vs. Hybrid Annealing

A \ B	Hill	Anneal
Hill	0.386 0.251 0.435	0.895 0.229 0.694
Anneal	0.051 0.777 0.556	0.570 0.704 0.947

% Optimal Utility: A / B / Joint

Figure Five – A Prisoner's Dilemma

Discussion & Conclusions

Klein, Faratin & Bar-Yam go on to try and develop institutional mechanisms that will avoid the deficient equilibrium of the prisoner's dilemma with particular application to autonomous agents. However, this is not our concern. Preliminary results of our research have demonstrated that the hybrid computational models considered are successful in maximizing social welfare. It is no accident that both simulated annealing and tabu lists can be translated into very straightforward suggestions for how to conduct human negotiations. Annealing suggests that mediators should "Encourage the consideration of contracts that may be inferior to previous proposals." And tabu lists counsel that "If a state of a particular issue has been recently proposed, consider other issues." Our intent was to use the computational simulations to aid in the design of social heuristics that may improve the outcomes of human negotiations.

Future research will examine more closely the role of cooling schedules and the length of tabu lists. We also plan to add the possibility of both qualitative and quantitative issues to the boolean issues that are now considered. Higher order issue interdependence and mutli-party (n>2) negotiations that give rise to coalitions would also both be of interest. Most importantly, however, the next phase of our research will be conducted in human subjects laboratories to determine if the insight provided by the computational simulations is translatable to the vulgarities of actual human exchange.

References

[Bar-Yam, 1997] Bar-Yam, Y., 1997, *Dynamics of Complex Systems*, New York: Perseus.

[Binmore & Vulkan, 1999] Binmore, K. & Vulkan, N., 1999, "Applying Game Theory to Automated Negotiation," *Netomics* 1: 1-9.

[Cerny, 1985] Cerny, V., 1985, "A Thermodynamical Approach to the Traveling Salesman Problem: An Efficient Simulation Algorithm," *Journal of Optimization Theory and Applications* 45:41.

[Cheng, Chan & Lin, 2005] Cheng, C., Chan, C.H. & Lin, K., 2005, "Intelligent Agents for E-Marketplace: Negotiation with Issue Trade-offs by Fuzzy Inference Systems," *Decision Support Systems* (in press).

[Ehtamo, Kettunen & Hamalainen, 2001] Ehtamo, H., Kettunen, E. & Hamalainen, R.P., 2001, “Searching for Joint Gains in Multi-Party Negotiations,” *European Journal of Operational Research* 130:54.

[Faratin, Sierra & Jennings, 2002] Faratin, P., Sierra, C. & Jennings, N.R., 2002, “Using Similarity Criteria to Make Issue Tradeoffs in Automated Negotiations,” *Artificial Intelligence* 142:205.

[Glover & Laguna, 1997] Glover, F. & Laguna, M., 1997, *Tabu Search*, Norwell, MA: Kluwer.

[Keeny & Raiffa, 1993] Keeny, R.L. & Raiffa, H., 1993, *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*, Cambridge: Cambridge University Press.

[Kirkpatrick, Gelatt & Vecchi, 1983] Kirkpatrick, S., Gelatt, C. D. & Vecchi, M. P., 1983, “Optimization by Simulated Annealing,” *Science* 220:671.

[Klein, Faratin & Bar-Yam] Klein, M., Faratin, P. & Bar-Yam, Y., “Using an Annealing Mediator to Solve the Prisoners’ Dilemma in the Negotiation of Complex Contracts,” (unpublished working paper).

[Klein & Faratin, 2003] Klein, M. & Faratin, P., “Negotiating Complex Contracts,” *Group Decision and Negotiation* 12: 111-125.

[Krippendorff, 1986] Krippendorff, K., 1986, *A Dictionary of Cybernetics* (unpublished report, The Annenberg School of Communications, University of Pennsylvania).

[Raiffa, 2003] Howard Raiffa, 2003, *Negotiation Analysis: The Science and Art of Collaborative Decision Making*, New York: Belknap.