

A research and education initiative at the MIT
Sloan School of Management

Negotiating Complex Contracts

Paper 125

August 2001

Mark Klein
Peyman Faratin
Hiroki Sayama
Yaneer Bar-Yam

For more information,

please visit our website at <http://ebusiness.mit.edu>

or contact the Center directly at ebusiness@mit.edu or 617-253-7054

Negotiating Complex Contracts

Mark Klein
Massachusetts
Institute of
Technology
m_klein@mit.edu

Peyman Faratin
Massachusetts
Institute of
Technology
peyman@mit.edu

Hiroki Sayama
New England
Complex Systems
Institute
sayama@necsi.org

Yaneer Bar-Yam
New England
Complex Systems
Institute
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, linear utility functions, and tractable contract spaces. Many real-world contracts, by contrast, are much more complex, consisting of multiple interdependent issues, nonlinear utility functions, and intractably large contract spaces. This paper describes a simulated annealing based approach appropriate for negotiating such complex contracts, evaluates its efficacy, and suggests some potentially promising avenues for defining more efficient algorithms for negotiating complex contracts.

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]. We can frame what these techniques do as follows (Figure 1):

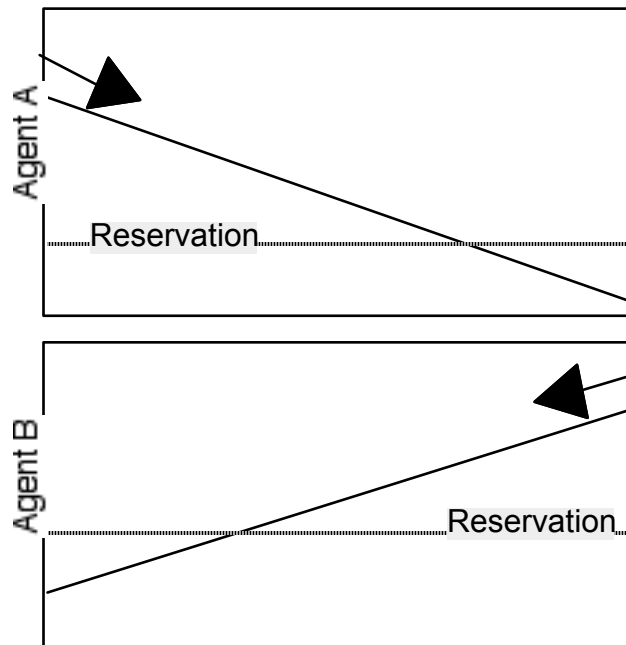


Figure 1: The standard view of negotiation.

Each point on the X axis represents a candidate contract¹. 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* [3]. 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 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 DVD player, for example, depends on whether it is a good match with the tuner and speakers I plan to purchase with it. Such issue interdependencies lead to *nonlinear* utility functions with multiple local optima [3]:




Figure 2: Complex negotiation.

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 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

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

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.

Such contexts, we argue, require radically different negotiation techniques, which allow agents to find 'win-win' contracts in intractably large multi-optima search spaces in a reasonable amount of time. In the following section we describe a negotiation approach that make substantial progress towards achieving these goals.

Mediated Single Text Negotiation

A standard approach to dealing with complex negotiations in human settings is the mediated single text negotiation [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 process continues, generating successively better contracts, until the reservation utility value is met or exceeded for both parties. We can visualize this process as follows (Figure 3):




Figure 3: Single text negotiation.

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 negotiating to find a mutually acceptable contract consisting of a vector S of 100 boolean-valued issues, each issue assigned the value 0 or 1, corresponding to the presence or absence of a given contract clause. This defined a space of 2^{100} , or roughly 10^{30} , possible contracts. Each agent had a utility function calculated using its own 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:

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

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 until the utility values for both agents become stable (i.e. until none of the newly generated contract proposals offer any improvement in utility values for either agent). 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 simulated annealers. The hill climbers used a very simple decision function: they accepted a mutated contract only if its utility to them was greater than that of the last contract they accepted. The annealers were more complicated, implementing a Monte Carlo machine [5]. Each annealer had a virtual ‘temperature’ T , such that it will accept contracts worse than earlier ones with the probability:

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

where ΔU is the utility change between 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. This kind of annealing has proven effective in finding near-optima in large multiple-optima utility functions, because annealers can move freely through the utility function, potentially skipping relatively small valleys on the way to higher optima [3]. This suggests that annealers will be more successful than hill-climbers in finding good contracts through the negotiation process. The reality, as we shall see, turned out to be more complicated.

Our aggregate results comparing hill-climbers with annealers can be summarized as follows:

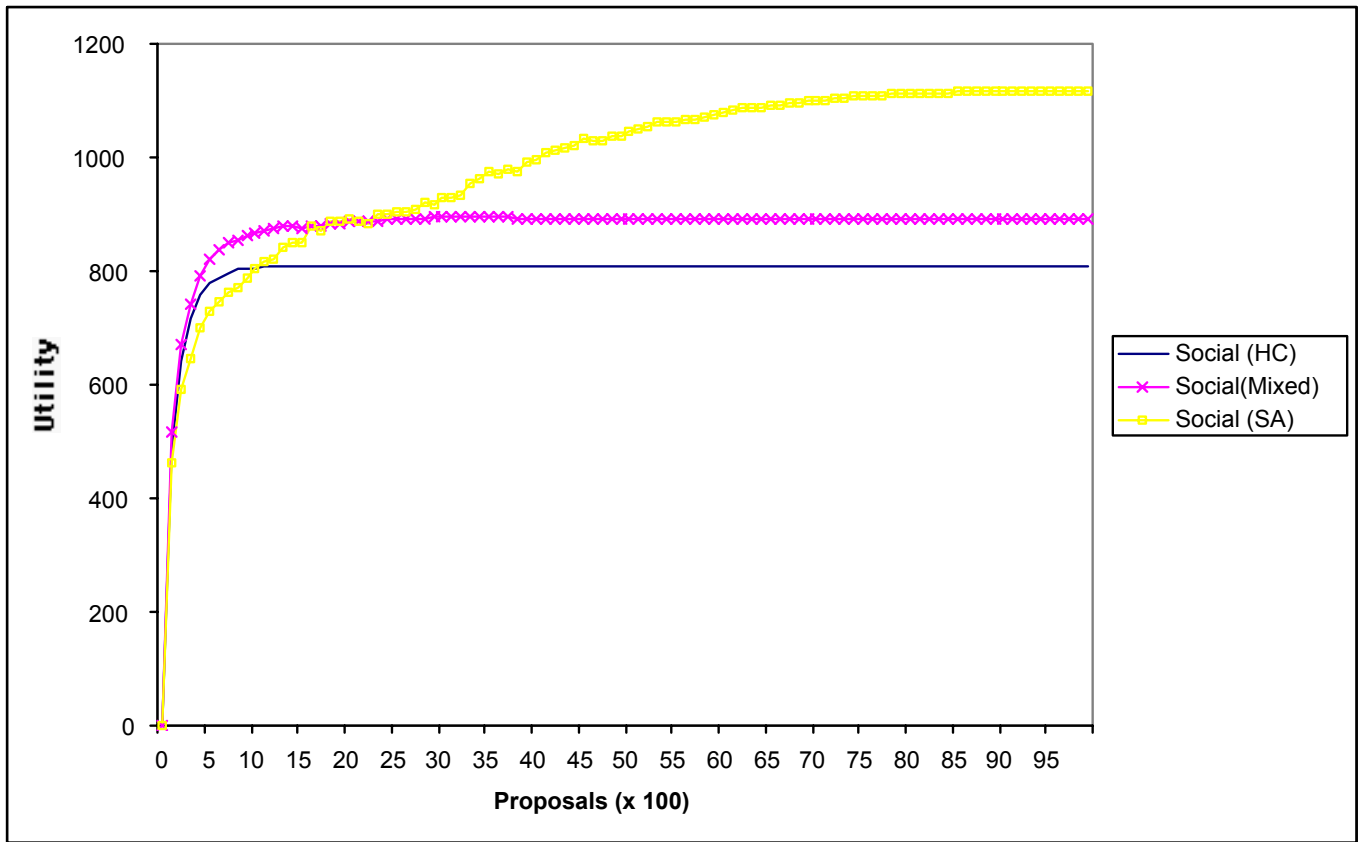


Figure 3: Social welfare values

This figure shows the social welfare (i.e. the sum of the contract utilities for the two negotiating agents) for a pair of annealers, a pair of hill-climbers, and a ‘mixed’ case with both an annealer with a hill-climber, averaged over 100 simulation runs.

The social welfare results revealed two interesting patterns. One is that *the presence of annealer agents always increases social welfare*. The social welfare for the two annealer case was roughly 40% greater than that of the two hill-climber case, and the mixed case produced a smaller but still statistically significant 15% improvement over the hill-climbers. This confirmed the value of the simulated annealing approach.

These results can be understood better by looking at the individual agent payoffs:

	Agent 2 anneals	Agent2 hill-climbs
Agent 1 anneals	[1100] 550/550	[880] 700/180
Agent1 hill-climbs	[880] 180/700	[800] 400/400

where the cell values are laid out as follows:

[<social welfare>
<agent 1 utility>/<agent 2 utility>

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





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

This figure shows the normalized utilities of the accepted contracts for each agent, plotted next to the pareto-efficient line. 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.

If one agent is a hillclimber and the other is an annealer, the hillclimber does well but the annealer fares



very poorly. This pattern can be understood as follows. When an annealer is at a high virtual temperature, it becomes a chronic conceiver, accepting almost anything beneficial or not. The hill-climber thus in effect 'drags' the annealer towards its own local optimum, which is not particularly likely to also be optimal for the annealer:

Figure 5: A typical negotiation with an annealer and hillclimber.

In our experiments, most (roughly 60%) of the negotiations between hillclimbers and annealers reached pareto-optimal contracts, but they tended to unfair and did not maximize social welfare.

The highest social welfare was achieved when both agents are annealers, since they both are willing to accept individually worse contracts in the beginning to help find win-win contracts later on:




Figure 6: A typical negotiation with two annealers.

The contracts found by two annealers are, in addition, much more fair (i.e. the agents achieve roughly equal normalized utilities) than those found by an annealer paired with a hillclimber.

Our analysis reveals a dilemma, however. 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 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. If all agents do this, however, then they forego the higher individual utilities they would get if they both annealed (the “tough guy’s” penalty). 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 [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 all agents would select annealing. In an open system environment we can not rely on self-reports for this, since agents are incited 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. An annealer will tend to accept a much higher percentage of proposed contracts that a hill-climber, especially at higher virtual temperatures (Figure 4, which shows the proposal acceptance percentages plus and minus one standard deviation for hill-climber and simulated annealer agents):

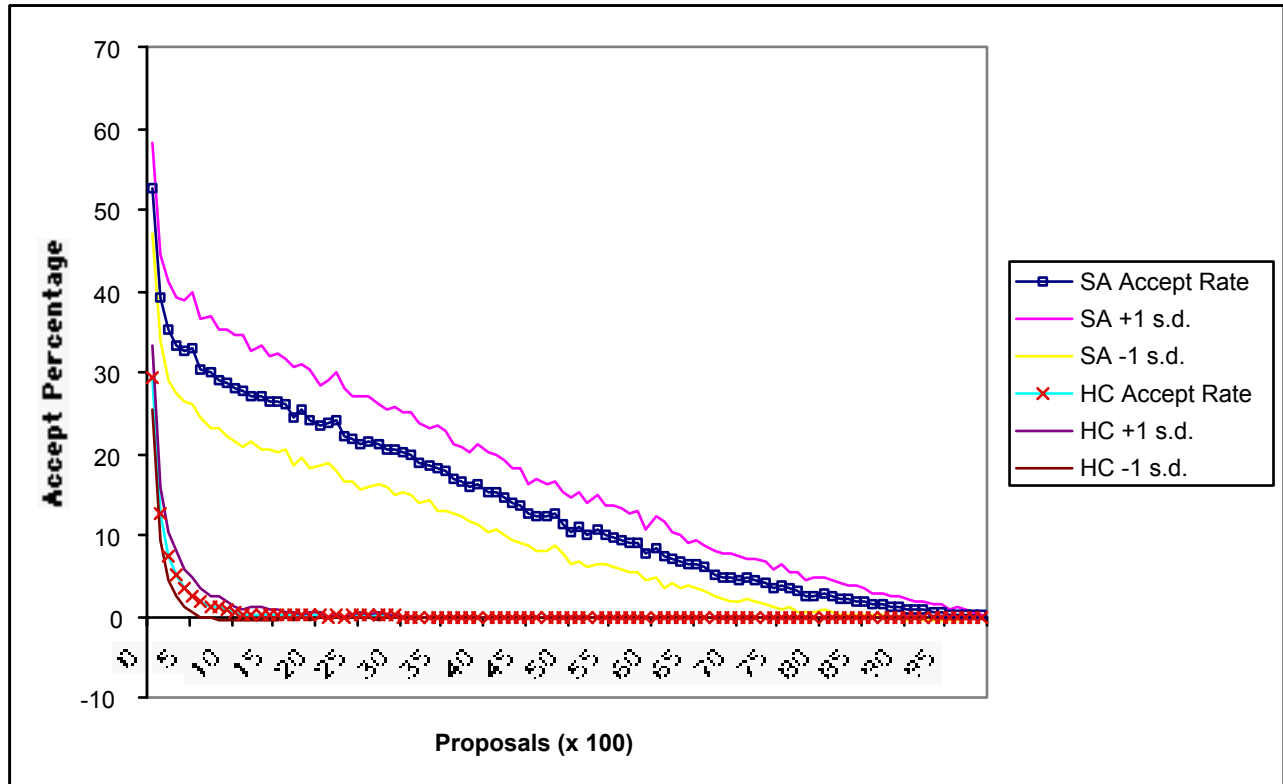


Figure 4: Proposal acceptance percentages for hill-climbers and annealers.

The problem with this ‘adaptive’ approach is that determining the type of an agent based on its voting behavior takes time. Agents must start with a guess concerning the other agent’s strategy and then observe its voting behavior to see what it actually uses. 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*, 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. The early commitment of utility is a result of the topology of nonlinear utility functions. These functions are fractal (i.e. self-similar at different scales) with the highest optima also tending to be the widest. They are thus shaped like mountain ranges, wherein the steepest slope tends to occur earlier, and the slope reduces as one gets closer to the summit.

While adaptive strategies can not eliminate the prisoner’s dilemma, they can however reduce its magnitude. Let us consider a specific example of an adaptive strategy we can call “tit-for-tat annealing” (T4TA). 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 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 the “tough guy’s” penalty. If we include this strategy we get the following payoff matrix:

	Agent2 hill-climbs	Agent 2 anneals	Agent 2 T4TA
Agent1 hill-climbs	[800] 400/400	[880] 700/180	[840] 500/340
Agent 1 anneals	[880] 180/700	[1100] 550/550	[1100] 550/550
Agent 1 T4TA	[840] 340/500	[1100] 550/550	[1100] 550/550

A T4TA agent fares just as well as an annealer when paired with an annealer or another T4TA agent, and has a reduced conceder's penalty when paired with a hill-climber as compared to an annealer.

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. The more proposals considered, however, the more information an agent reveals.

Next Steps: 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 *à 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 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 is 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.

Next Steps: Addressing the Prisoner's Dilemma

We have shown that there is no way to avoid the prisoner's dilemma within the scope of a single negotiation, though we can reduce the magnitude of the effect using adaptive strategies. 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.

Next Steps: Contracts as Processes

Another direction we plan to pursue for our future work involves providing a systematic way of defining the space of possible contracts. Any contract can, we argue, be viewed as the specification for a *process* that spells out which actor does what when. The simplest of contracts may only specify that a good is exchanged for a given monetary consideration. The most complex contracts may spell out, in excruciating detail, what each party should do in a wide range of normal and exceptional circumstances. In all cases, however, the contract represents a mutually agreed-to process.

Previous work [11] has shown that process (and therefore contract) design can be treated as configuration, wherein one first identifies the abstract process one is interested in, and then customizes it by selecting specific processes for each of the substeps in the abstract process. To make this concrete, imagine that we want to subcontract out the task of purchasing goods over the Internet. The first step is to find an abstract process for this that we can customize. One way to do so is by retrieving the appropriate process from a process ontology, such as that stored in the MIT Process Handbook [12]. The Handbook ontology contains over 5000 business process models, ranging from very abstract processes such as ‘allocate resources’ to relatively specific ones. The Handbook model for ‘buy over internet’, for example, consists of the following five substeps:

- Identify needs
- Find sources via internet
- Select supplier
- Place order over internet
- Pay using credit card
- Receive good

Each one of the substeps in this abstract process model has a branch of the process ontology that captures different ways of achieving this step. The branch for the ‘select supplier’ substep, for example, is the following:

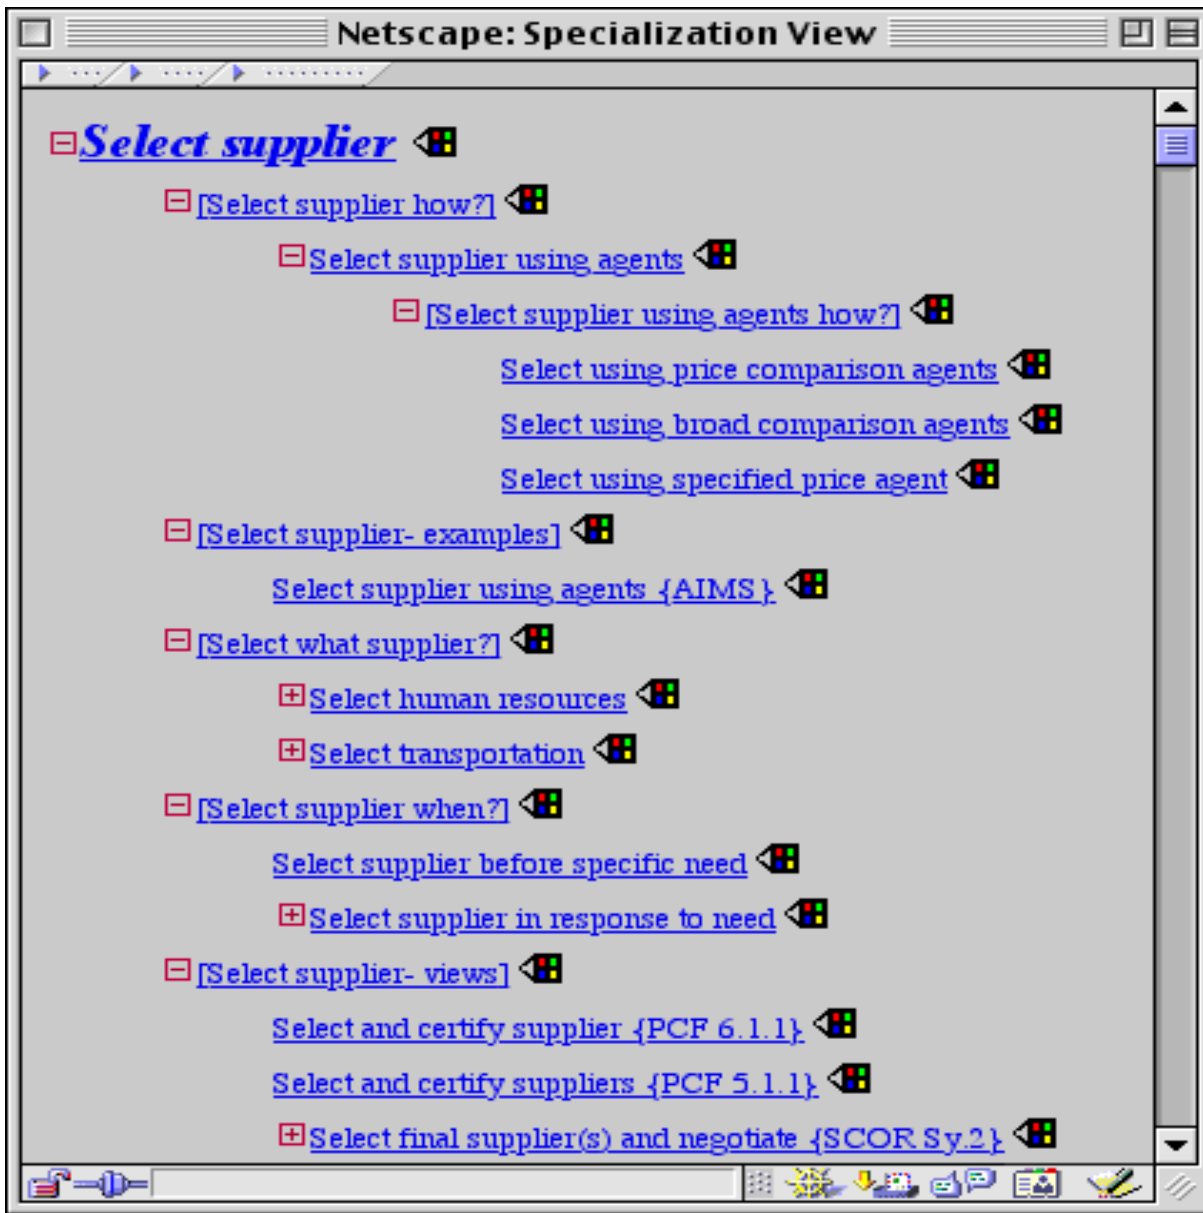


Figure 7: A fragment of the ‘select supplier’ ontology.

One can therefore start specifying the contract by selecting which one of the alternative ‘select supplier’ processes will be used. Repeating this procedure for all the substeps of the abstract ‘buy over internet’ process (as well as specifying any necessary attribute values for the selected substep processes) should result in a complete specification of the normative aspects of the subcontract.

The next step in defining the contract is to consider how the execution of the contract can fail (i.e. its *exceptions*) and how those exceptions can be handled. This is an important problem. MAS are increasingly being viewed as a way to rapidly connect entities that have never worked together before, for such applications as disaster recovery operations, open electronic marketplaces, virtual supply chains, and international coalition military forces [13] [14] [15] [16] [17]. Such ‘open’ systems introduce a wide range of potentially devastating exceptions including infrastructure failures, agents renegeing on commitments, denial of service attacks, emergent dysfunctions such as chaotic behavior [18] [19] and so on [20] The vast majority of MAS work to date, however, has considered only well-behaved agents running on reliable infrastructures in relatively simple domains [21] [22].

One possible approach to handling contract execution exceptions is to simply require that agents charge less for their services in proportion to how likely they are to fail at their assigned tasks. The problem with this approach is that it removes from subcontractor agents any direct incentive to avoid failures or reduce their impact on the contractor (they have already paid the penalty up-front) and it can be difficult for the contracting agent to assess the subcontractor failure likelihood and thereby the appropriate discount. An alternative approach is for the contract to specify the actions, including for example early notification or penalties, to be taken for each important exception. The appropriate penalties are easier to determine because the contractor need only estimate the impact (but not the probability) of an exception, and it provides the subcontractor with incentives to avoid high-impact failures. This thus results in contracts that increase social welfare.¹

The Process Handbook ontology has been extended to support such contractual specification of exception handling behavior, by the addition of a taxonomy of generic exception types, such that all processes are linked to their characteristic exceptions, and all exceptions are linked in turn to processes suitable for handling (anticipating and avoiding, or detecting and resolving) them [23]. Imagine, for example, that we have selected a Dutch (descending price) auction process for supplier selection. By consulting the Process Handbook ontology we can see that such auctions are prone, among other things, to the ‘price collision loop’ exception wherein two agents give the same bid for a good, leading the auctioneer to raise the price and try again, leading to the same bid conflict, ad infinitum. The Handbook ontology describes a range of handlers for this exception, such as randomly selecting a winner, disqualifying one or both conflicting bidders, and so on. Once handlers have been selected for all exceptions that can potentially occur with the selected business process (and any handler process attributes have been specified as necessary), we have finished specifying the exception-related aspects of the contract.

It is straightforward to map contract configuration as described above into a negotiation framework:

- for every substep S in the abstract process being specified, create the issue “how do we achieve substep S?” whose candidate values are all the possible processes suitable for achieving that substep
- for every exception E in the process being defined, create the issue “how do we handle exception E?” whose candidate values are all the possible processes suitable for handling that exception
- for every process attribute A in the abstract process, substeps and exception handlers, create the issue “what is the value for attribute A?” whose candidate values depend on the attribute being considered

The issues involved in designing contingent contracts are thus highly interdependent. The output of one substep in a process, for example, will often have to match the input of the next substep, so the utility of one choice is highly dependent on the other choice. The value of a given exception handler can be dependent upon choices made for the normative steps and for the other exception handlers.

This approach relies on treating contract formation as configuration from a pre-defined design space, but this arguably is realistic for many important real-world domains such as supply chains where the abstract processes and most commonly used alternatives are relatively stable and well-known.

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 that can deal effectively with the nonlinear utility functions. 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

¹ We are grateful to Benjamin Grosz of the MIT Sloan School of Management for pointing out this tradeoff.

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.

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

1. 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.
2. 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.
3. 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.
5. 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.
7. 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.
11. Bernstein, A., M. Klein, and T.W. Malone. *The Process Recombinator: A Tool for Generating New Business Process Ideas*. in *Proceedings of the International Conference on Information Systems (ICIS-99)*. 1999. Charlotte, North Carolina USA.
12. Malone, T.W., et al., *Tools for inventing organizations: Toward a handbook of organizational processes*. Management Science, 1999. **45**(3): p. 425-443.
13. Jennings, N.R., K. Sycara, and M. Wooldrige, *A Roadmap of Agent Research and Development*. *Autonomous Agents and Multi-Agent Systems*, 1998. **1**: p. 275-306.
14. Wooldridge, M., N.R. Jennings, and D. Kinny. *A Methodology for Agent-Oriented Analysis and Design*. in *Proceedings of the Third Annual Conference on Autonomous Agents (AA-99)*. 1999. Seattle WA USA: ACM Press.
15. *Proceedings of the International Workshop on Knowledge-Based Planning for Coalition Forces*. 1999: Edinburgh, Scotland.
16. Fischer, K., et al. *Intelligent agents in virtual enterprises*. in *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM'96)*. 1996. Blackpool, UK.
17. Tsvetovaty, M.B., et al., *MAGMA: An agent-based virtual marketplace for electronic commerce*. Applied Artificial Intelligence, 1997. **11**(6): p. 501-524.
18. Youssefmir, M. and B. Huberman. *Resource contention in multi-agent systems*. in *First International Conference on Multi-Agent Systems (ICMAS-95)*. 1995. San Francisco, CA, USA: AAAI Press.
19. Serman, J.D., *Learning in and about complex systems*. 1994, Cambridge, Mass.: Alfred P. Sloan School of Management, Massachusetts Institute of Technology. 51.
20. Klein, M. and C. Dellarocas. *Exception Handling in Agent Systems*. in *Proceedings of the Third International Conference on AUTONOMOUS AGENTS (Agents '99)*. 1999. Seattle, Washington.
21. Hägg, S. *A Sentinel Approach to Fault Handling in Multi-Agent Systems*. in *Proceedings of the Second Australian Workshop on Distributed AI, in conjunction with Fourth Pacific Rim International Conference on Artificial Intelligence (PRICAI'96)*. 1996. Cairns, Australia.
22. Klein, M., J.A. Rodriguez-Aguilar, and C. Dellarocas, *Using Domain-Independent Exception Handling Services to Enable Robust Open Multi-Agent Systems: The Case of Agent Death*. 2000, Massachusetts Institute of Technology: Cambridge MA USA.

23. Klein, M. and C. Dellarocas, *A Knowledge-Based Approach to Handling Exceptions in Workflow Systems*. Journal of Computer-Supported Collaborative Work. Special Issue on Adaptive Workflow Systems., 2000. 9(3/4).