

Leveled Commitment Contracts with Myopic and Strategic Agents

Martin R. Andersson and Tuomas W. Sandholm¹

{mra, sandholm}@cs.wustl.edu
Department of Computer Science,
Washington University
One Brookings Drive
St. Louis, MO 63130-4899

Abstract

In automated negotiation systems consisting of self-interested agents, contracts have traditionally been binding, *i.e.*, impossible to breach. Such contracts do not allow the agents to efficiently deal with future events. This deficiency can be tackled by using a *leveled commitment contracting protocol* which allows the agents to decommit from contracts by paying a monetary penalty to the contracting partner. The efficiency of such protocols depends heavily on how the penalties are decided. In this paper, different leveled commitment protocols and their parameterizations are empirically compared to each other and to several full commitment protocols. In the different experiments, the agents are of different types: self-interested or cooperative, and they can perform different levels of lookahead.

Surprisingly, self-interested myopic agents reach a higher social welfare quicker than cooperative myopic agents when decommitment penalties are low. The social welfare in settings with agents that performed lookahead did not vary as much with the decommitment penalty as the social welfare in settings that consisted of myopic agents. For a short range of values of the decommitment penalty, myopic agents performed almost as well as agents that performed lookahead. In all of the settings studied, the best way to set the decommitment penalties was to choose low penalties, but ones that were greater than zero. This indicates that leveled commitment contracting protocols outperform both full commitment protocols and commitment free protocols.

Introduction

Systems which include *automated negotiation* are playing an increasingly important role in our society. This is due both to *technology push* and *application pull*. New technology has made global communication possible with the help of standardized protocols *e.g.* IP, WWW, Java, HTML, KQML which have made it possible for independent hardware platforms and software to communicate with each other. With the expansion of the Internet and novel methods for electronic payments, electronic commerce is also becoming widespread (Kalakota & Whinston 1996). Another

¹Supported by NSF CAREER award IRI-9703122 and NSF grant IRI-9610122.

application of automated contracting systems is automated markets for electric power (Sandholm & Ygge 1997; Ygge & Akkermans 1996) which are a reality in both the United States and in Europe.

Contracts in automated negotiation systems consisting of *self-interested* agents have traditionally been binding *i.e.* impossible to breach. Such contracts do not allow the agents to act efficiently upon future events because contracts might become unfavorable to one or both of the agents after the contracting. If the agents were allowed to breach contracts they could accommodate changes in the environment more efficiently and the social welfare would improve.

In systems with cooperative agents, decommitting from contracts without reprisals can be accepted even after the other party has partly completed the task of the contract (Sen & Durfee 1994; Smith 1980). On the other hand, in the case of self-interested agents, there is a need to compensate the party who is the victim of a decommitment.

Contingency contracts have been suggested to be used between self-interested agents when they do not have knowledge (or only have probabilistic knowledge) about future events (Raiffa 1982). In these contracts the obligations of the contract are made conditional on future events. Contingency contracts can increase the payoff of both parties, so contracts not possible with full commitment protocols may become beneficial for both parties. However, it may be impossible to anticipate and enumerate all future events. Monitoring all events after the contract is made can also be impractical. If some events are observable by only one of the parties, another problem arises: one party can have an incentive to lie about the events in order to be better off himself.

Recently, a *leveled commitment protocol* has been proposed that allows self-interested agents to decommit from a contract by simply paying a decommitment penalty to the partner of the contract (Sandholm & Lesser 1996). In that protocol, the decommitment penalties are decided at the time of contracting and the penalties do not need to be the same for the two contracting parties. It has been shown through formal game theoretic analysis that this leveled commitment feature increases the Pareto efficiency of contracts and can make contracts individually rational to both parties even in cases where full commitment contracts can-

not.

The leveled commitment contracts which only can handle one task could emulate full commitment contracts that are capable of handling several tasks in one single contract (Andersson 1998; Andersson & Sandholm 1998; Sandholm 1996; 1998).

With leveled commitment protocols there is no need for an agent to conduct a feasibility check before contracting. If the contract cannot be performed because of lack of resources *etc.* the agent can decommit from the engagement. Also the agents do not need to perform a complete computation of the marginal costs before taking on a contract but can complete the computation after contracting.²

The concept of breaching contracts in the real world is analyzed by Posner (Posner 1977). The main ideas are that the party that breaches must compensate the victim for lost profit and that the penalties for breaching contracts should be set so that the social welfare is maximized.

Strategic thinking behind contracting and one-sided decommitment among self-interested agents has been studied (Park & Durfee & Birmingham 1996) where the decision making mechanism is modeled as a Markov-process. They assume that the agents expect none of their bids to be accepted which makes the approach unsound from the perspective of game theory.

Diamond and Maskin (Diamond & Maskin 1979) have studied systems in which both agents can decommit from a contract by paying a decommitment penalty to the other party of the contract. Those penalties can be set in different ways: they can be compensatory or privately decided (*i.e.* liquidated; not necessarily decided by the parties of the contract – maybe imposed by a court) in the contract. The compensatory decommitment penalties are favored because of efficiency *i.e.* they provide whenever possible an increasing mutual welfare between the agents that enter a new contract.³ Another argument for the compensatory penalties is that they are exactly the penalties that two rational parties would agree on for privately decided penalties.⁴

²This allows the agent to act faster and with less constraints in the contracting process than if it always had to perform a feasibility check and a thorough marginal cost calculation before contracting. The system also becomes more efficient if one agent (the one that takes on the task) conducts a full marginal cost calculation, than if all the agents would perform such a calculation.

³However, the social welfare may decrease because of the inefficiency arising from the contract that is breached.

⁴One reason for having over-compensating penalties is that one agent can make himself more trustworthy and the expected utility of the other agent will increase enough to make the contract possible (Posner 1977). Hence, over-compensating penalties can increase the space of possible contracts. However, it would limit the space of possible decommitments.

The next section discusses leveled commitment protocols. Then the example problem domain is described which is followed by a presentation of the types of agents and protocols included in the study. Next is the notation of the different protocols presented followed by a description how the data was evaluated. Then the results are presented and discussed.

Leveled Commitment Protocols

The penalties could be chosen freely by the agents in which case the agents would try to optimize the penalties in its favor (Sandholm & Lesser 1996). The negotiation will be more complex if the agents can decide the penalties themselves compared to when they are set by the protocol because there would be more variables to agree on in order for all parties to accept a contract. If the penalties are set by the protocol the negotiation becomes easier but the result may not be optimal (*e.g.* fixed penalties does not guarantee that it is profitable for the agents to decommit in all situations when a decommitment is mutually profitable for the agents involved in the contract).

Another method is to relate the decommitment penalty to the price of the contract. This could be done by choosing the penalty as a percentage (or a more complex function) of the price. Another way is to make the penalties compensate the victim of the breach for its lost profit. Because the victim would have an incentive to lie about the expected profit a mechanism for calculating the lost profit would be necessary. The state of both agents might have changed since the contract was made so the expected lost profit at contracting time and breaching time may differ. In the extreme the lost profit for the victim can be negative at breaching time that is also the victim of the breach benefits from being freed from the contract obligations.

A breach close to the execution deadline of the contract or late in a negotiation is likely to be more costly for the victim since it can be hard to find someone else to contract with within a short amount of time. In order to prevent such occurrences the decommitment penalties can be increased over time.

The leveled commitment protocols were studied in order to conclude which mechanism should be used for setting the decommitment penalties among qualitatively different agents. Several environments were studied and in each of them 16 protocols with many parameterizations were tested.

Description of the Problem Domain

To investigate the performance of different mechanisms of setting the decommitment penalty for leveled commitment protocols the agents are divided into two subsets. One subset consisted of *contractors* which each had one task with a cost associated to it. A contractor considers to contract out its task to a *contractee* which would be able to handle the task at a potentially cheaper cost. The other subset consists of the con-

contractees which do not have any tasks initially but resources to handle a maximum of one task at a cost specific to each combination of contractees and tasks. There is no other differences between a contractor and a contractee other than that the former initially has a task and is willing to give it away and the latter is willing to take on a task. The fallback position of a contractor is the cost of handling the task it has at the start of the negotiation while the fallback of the contractees are zero (*i.e.* they do not have any tasks at the start of negotiation and therefore no pending expenses).

Types of Agents

Four types of agents are included in the study. Each type is designed from two properties: the amount of *lookahead* the agent performs before it accepts or rejects a proposed contract and how *self-interested* the agent is.

Agents With and Without Lookahead

The agents can either perform full lookahead or none at all. If they perform full lookahead they compute the payoff of all possible future events and agree to the contract only if the expected payoff of agreeing is greater than rejecting *i.e.* they act strategically as game theoretic agents would. The other option of lookahead is to do none at all. If the agents perform no lookahead they act myopically and only consider the immediate payoff of the contract under negotiation.

Individual Rational and Cooperative Agents

The agents can be *self-interested* (*SI*) and only agree to contracts which increase their own payoff (expected or immediate depending on whether they perform lookahead or not). They could also be explicit social welfare maximizers (*SWF-maximizer*) *i.e.* *cooperative agents* which consider the summed payoff of all agents in the system when deciding to accept or reject a contract. That is a *SWF-maximizing* agent can agree to a contract even if it makes itself worse off as long as the total social welfare in the system increases.

Types of Leveled Commitment Protocols

The leveled commitment contracts are defined as follows:

Definition. 1 A leveled commitment contract is a tuple $\langle C, \Gamma \rangle$, where C is the underlying full commitment contract and Γ is the set of decommitment penalties. Let A_C be the set of agents involved in the contract C . Then Γ will consist of one decommitment penalty for each pair of agents in A_C , so that $|\Gamma| = \frac{|A_C|^2 - |A_C|}{2}$.

This definition has the nice feature of separating the leveled commitment framework from the obligations of the contract called the underlying contract C . This means that the leveled commitment protocol can be applied to any type of full commitment contract. If

an agent wants to decommitment from the underlying contract it has to pay the decommitment penalties stated in Γ to all agents involved in the contract.⁵

Four different mechanisms of deciding the decommitment penalties are studied: *fixed*; *percentage* of contract price; increasing and decided at the time of *contracting*; and increasing and decided at the time of *breaching*. Two more properties (how to compute the price of the contract and the sequence in which the agents meet each other in the negotiation) are also varied for different protocols in the study. All of these alternatives are now discussed.

Decommitment Penalties as a Fixed Value (FIX-protocol)

In the *FIX-protocol* all contracts have the same fixed decommitment penalty decided prior to the start of the negotiation and that penalty is used throughout the negotiation. In the study experiments with six different values of the fixed decommitment penalty were conducted. A summary of the different values can be found in Table 1.

Decommitment Penalties as a Percentage of Contract Price (PER-protocol)

The percentage decommitment penalties used in the *PER-protocol* consist of a fraction of the price of the contract. It is referred to as percentage since the most efficient parameterizations are fractions less than one. The same percentage is used for all contracts throughout the negotiation. The percentages in the study can be found in Table 1.

FIX-protocol	PER-protocol	CON-protocol BRE-protocol
0	0.1	0.25
5	0.25	0.5
10	0.4	0.75
15	0.5	1.0
30	0.75	2.0
50	1.0	4.0

Table 1: *Parameterizations in the study of leveled commitment protocols in different environments. (Note that 10% is written as 0.1 in the table.)*

Decommitment Penalties Decided at Time of Contracting (CON-protocol)

In the *CON-protocol* the decommitment penalty varies during the negotiation. For a certain contract the decommitment penalty is fixed at the time of contracting. The decommitment penalty is linearly increasing from zero to the highest penalty used. This highest penalty is set by the protocol to be a percentage of the contract price which is applied to a contract

⁵The concept of leveled commitment is not specific to task allocation problems, although the empirical studies of this paper focus on task allocation. All the leveled commitment protocols in this study have a full commitment contract that only can transfer one task from one agent to another as the underlying full commitment contract.

agreed on at the second to last round and then decommitted in the last round. The percentages for setting the highest penalty in the study can be found in Table 1.

Decommitment Penalties Decided at Time of Breaching (BRE-protocol)

In BRE-protocols the decommitment penalty also varies with the time of the negotiation as in the CON-protocol. The difference is that the decommitment penalties are fixed at the time of breaching instead of the time of contracting. The decommitment penalty is linearly increasing over time and a contract which is decommitted in the last round has the highest penalty. This highest penalty is decided as a percentage of the contract price and Table 1 shows the percentages used in the study.

Methods of Computing the Contract Price

The price of a contract is set such that the profit of the contract is divided equally between the agents involved in the contract. However the profits are calculated in two different ways: including and not including decommitment penalties:

- The price is set so that the profits from the fallback positions of the agents are divided equally between the agents. The current payoffs are not considered when the price is calculated and neither are the potential decommitment penalties.
- The price is set so that the increase from the current mutual profit of the agents is equally divided between them including discounts for decommitment penalties that are to be paid.

Sequencing of Contracts

The order in which the agents meet for negotiation is either stochastic or deterministic. In the stochastic model one contractor and one contractee are randomly picked to negotiate with each other. In the deterministic model the order is decided prior to negotiation (contractor 1 meets contractee 1, contractor 1 meets contractee 2, . . . contractor 2 meets contractee 1, . . .). The negotiation protocol is sequential that is only two agents (one contractor and one contractee) negotiate at a time.

If the agents are individually rational a contract is accepted if the payoffs (immediate or with lookahead) after the contract will be greater for both the contractor and the contractee compared to the payoffs before a potential acceptance of the contract. If one of the agents is indifferent that is the contract does not increase its payoff the other agent decides whether or not to perform the contract. However if both the agents are indifferent the contract is rejected.

Studied Parameterizations

The parameterizations of the leveled commitment protocol that were studied can be found in Table 1. The

Notation of a leveled commitment protocol in a certain domain: Cab.Dd. Cab characterizes the protocol and Dd the domain.	
$C \in \{\text{FIX, PER, CON, BRE}\}$	The type of penalty.
$a \in \{w, d\}$	The type of price computation: based on fallback positions (w) or current profit (d).
$b \in \{n, r\}$	The contract sequencing: random (r) or not random (n).
$D \in \{\text{SI, SWF}\}$	Agent type: willingness to cooperate: self-interested (SI) or cooperative (SWF).
$d \in \{l, m\}$	Agent type: lookahead depth: lookahead (l) or no lookahead (myopic) (m).

Table 2: Notation of the leveled commitment protocols and the domain.

percentages for CON- and BRE-protocols are considerably greater than the percentages for the PER-protocol. This is because in the CON- and BRE-protocols these great percentages are only used in the last negotiation round since the penalty is increased linearly from zero to the specified percentage with time of negotiation.

Conventions in Naming the Protocols and Domains

The notation for the different combinations of features of the leveled commitment protocols is summarized in Table 2. One or more properties may be left out if it is clear from the context which properties are referred to. If there is a star ("*") in place of a property all possible types of that property are considered.

Evaluation

In our experiments two contractors negotiated with two contractees over the contractors' initial tasks. Initially each agent was randomly assigned a cost for handling each task. The contractors' costs were in the interval [100, 200] and the contractees' in the interval [0, 100]. Since the contractors never can handle a task cheaper than a contractee they never have to negotiate with each other. The problem was solved for 100 randomly initiated problem instances with five negotiation rounds in each. The number of negotiation rounds was assumed common knowledge among the agents.

To compare the different protocols the ratio bound (ratio of the welfare of the obtained local optimum for a given type of contracts to the welfare of the global optimum) was used. The mean ratio bounds (over the 100 problem instances) were calculated for all the different leveled commitment protocols and agents. The 95% confidence intervals were also computed from which

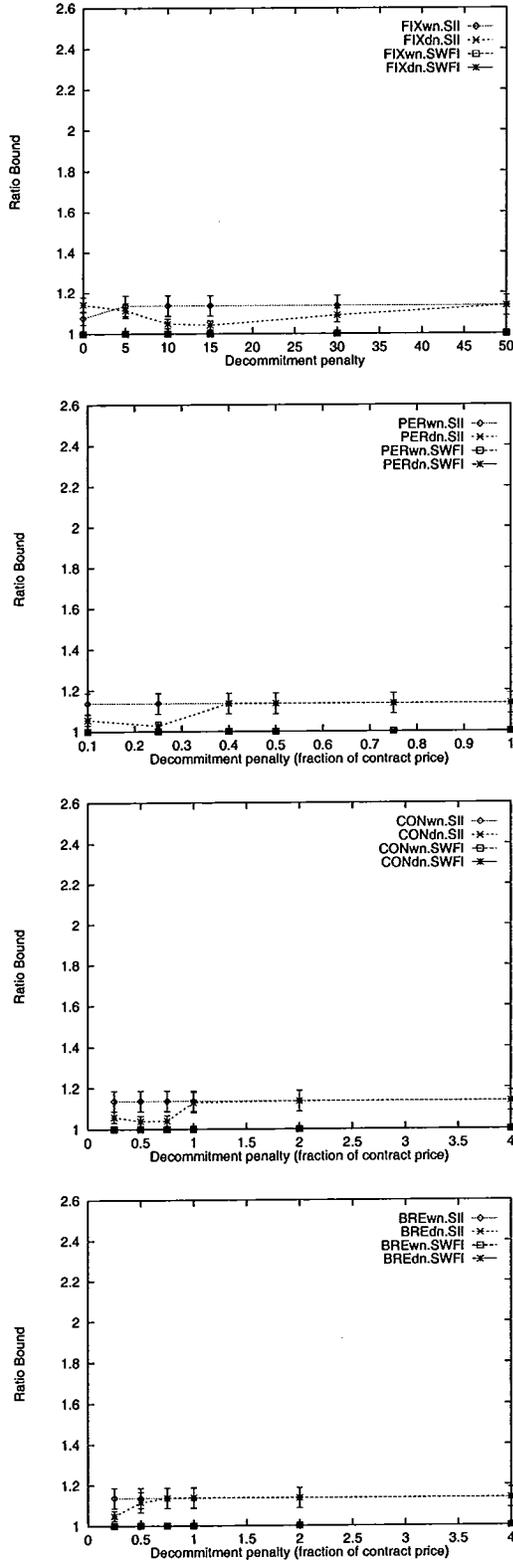


Figure 1: Agents that perform lookahead using deterministic protocols.

the results could be statistically analyzed.⁶

Results

The results are presented in the following order: First, the agent types are compared, then the methods of sequencing the contracts, the different ways of deciding on a contract price, and the best protocols for each agent type are discussed.

Comparison of Agent Types

Over all agents which performed lookahead reached a higher social welfare than myopic agents. By definition the SWFI-agents (SWF-maximizers conducting full lookahead) always reach the global optimum that could be reached considering the sequence of the contracts. The impact of the sequencing of the contracts will be discussed later. SWFI-agents contracting with a deterministic protocol always reached the global optimum (Figure 1).

Comparing the SIM- and SWFM-agents using a deterministic protocol, the SIM-agents outperformed the SWFM-agents in the region of low decommitment penalties (Figures 1 and 2). They did that for all the eight deterministic protocols. We can see that the *dn-protocols performs better than the *wn-protocols in the area of decommitment penalties where SIM-agents perform better than SWF-agents. This might be surprising, but without lookahead not even cooperative agents will reach the global optimum for sure, in a limited time. Technically, the reason why this happens is that a decommitment, which is not SWF-maximizing, is conducted by the SIM-agents. For some instances, if the negotiation contained more rounds, the SWFM-agents reached the same social welfare as the SIM-agents.

In several cases the myopic agents performed almost as well as the agents with full lookahead (Figures 1-4). There is a clear trade-off between reaching the globally optimal solution and computation. In these experiments (with small problem sizes) the agents do not gain much by performing a full lookahead compared to myopic agents with well set decommitment penalties. It is considerably more complex to perform a full lookahead than no lookahead, and for large problem instances a full lookahead is not even possible. On the other hand, the decommitment penalties do not have to be chosen so carefully if the agents perform a full lookahead. That is because the agents can evaluate the future events and acting upon that knowledge upfront, reducing the risk in a commitment.

Comparison of Methods of Sequencing the Contracts

Comparing the stochastic and the deterministic sequencing methods, the deterministic method always

⁶When using random sequenced leveled commitment protocols, the expected outcome was computed: a randomization device was not used to compute one of the possible outcomes for one instance, but all possible outcomes were computed and the expected value found.

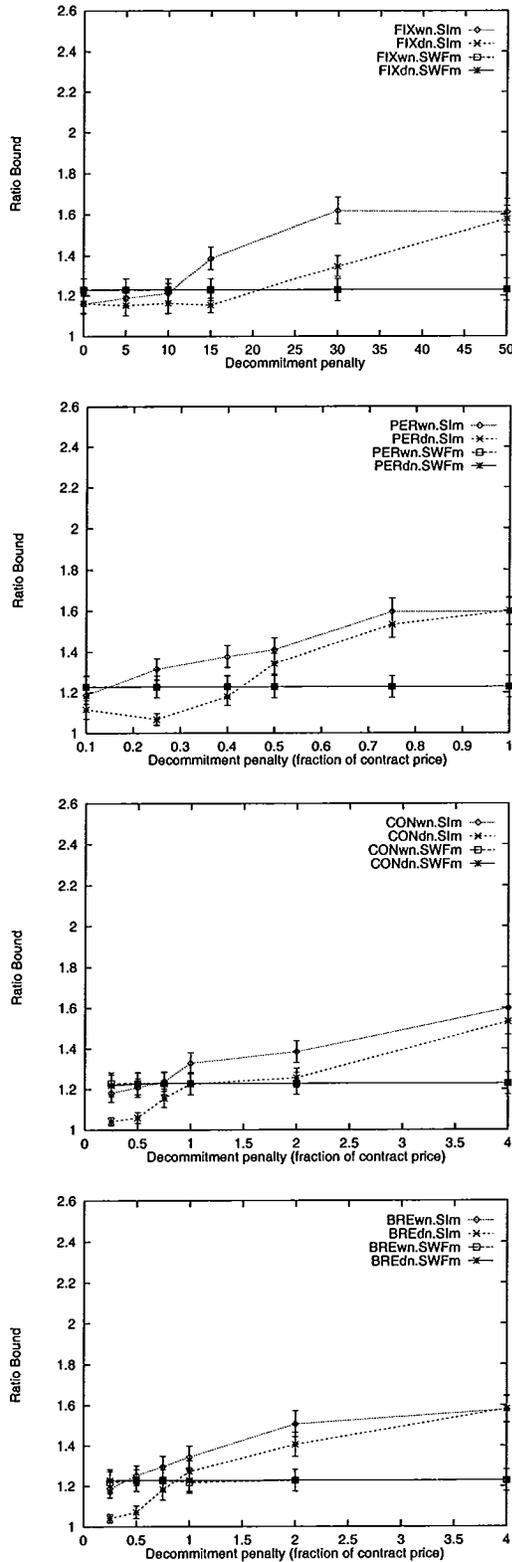


Figure 2: Myopic agents using deterministic protocols.

yields a lower ratio bound (Figures 1-4). That can be explained by the fact that the best possible ratio bound which is achievable with a stochastic method is greater than one. That is because the ratio bound is averaged over all possible outcomes including those that never can reach the global optimum. The extreme example is when the same contractor and contractee meet each other in every round of the negotiation. The best result achievable with the protocols using the stochastic method was always reached by the SWF-agents (SWF-maximators conducting full lookahead). For the sequences where the stochastic protocol could perform well it did (*i.e.* where all the agents participated in the negotiation).

Comparison of Methods of Computing the Contract Price

Of the two methods of computing the contract price (the method that considered the current profit and not only the original fallback positions) never reached a higher ratio bound than the other method for all protocols if the optimal parameterization for each protocol was used (Table 3). For other parameterizations the optimal method of setting the price varied. The method that considers the current profit performs well for low penalties with the deterministic protocols while the method based on the fallback positions performs well in the case of low penalties and a stochastic protocol.

Comparison of Methods and Parameterizations of Setting the Decommitment Penalty

A summary of the protocol that achieved the lowest ratio bound for each agent is found in Table 3. Two sets of best protocols are extracted one among the stochastic protocols and one among the deterministic protocols. For all the protocols the optimal choice of parameters was to use a low decommitment penalty (or a low percentage of contract price) which was greater than zero. Neither zero penalties nor high penalties performed well.

		Stochastic sequencing	
		Lookahead	Myopic
SI	SWF	BREd	PERd
	SWF	All	CONd

		Deterministic sequencing	
		Lookahead	Myopic
SI	SWF	PERd	BREd
	SWF	All	CONd/BREd

Table 3: Summary of the optimal choice of protocols for each agent and each parameterization.

Conclusions

In automated negotiation systems with self-interested agents it has traditionally not been possible to breach accepted contracts. Because of that the agents have been lacking the ability to act efficiently in a dynamic environment since they cannot accommodate future

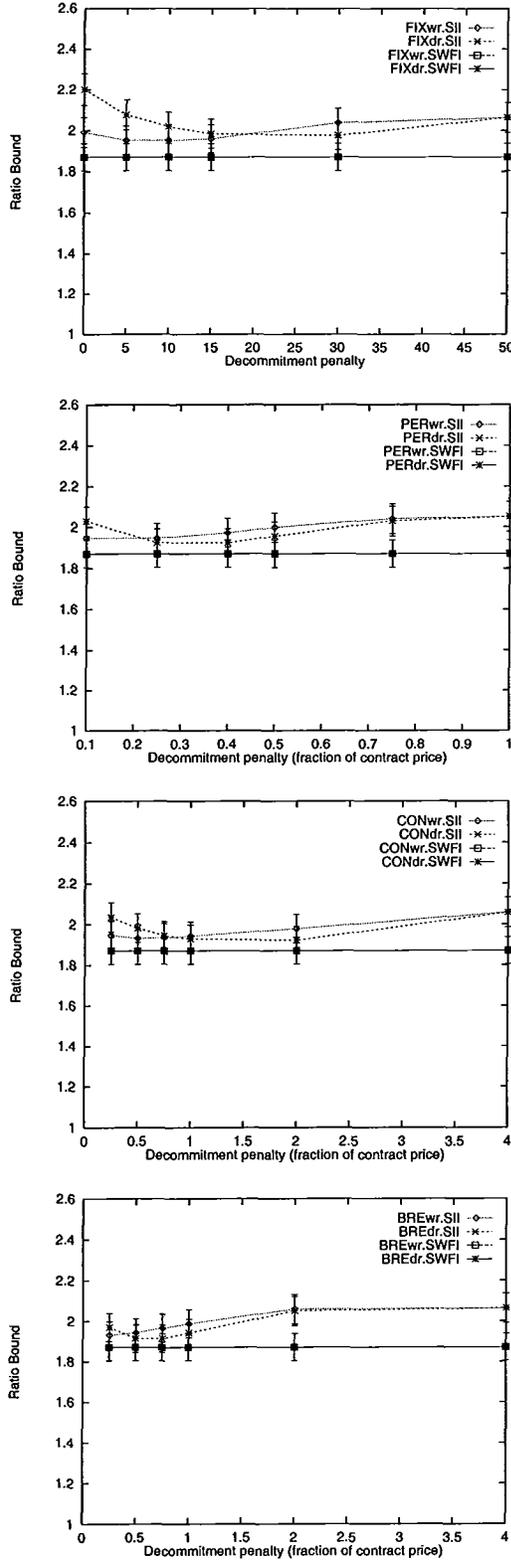


Figure 3: Agents that perform lookahead using stochastic protocols.

events efficiently. Contingency contracts have been suggested to solve this problem but they are not practical in all environments. Another alternative is to renegotiate but that incurs extra negotiation overhead and requires all parties of the contract to accept the new contract. Later leveled commitment protocols were introduced and it has been shown that they are more efficient than full commitment protocols. In a leveled commitment protocol the agents can decommit from a contract by paying a penalty to the partner(s) of the contract.

The efficiency of leveled commitment protocols depends drastically on how the decommitment penalties are decided. In this work we have investigated several different methods of setting them. If it would be possible for the agents to choose the penalties freely they would try to optimize the penalties in their favor. As a result the negotiation would be more complex: there would be more variables to agree on in order for all parties of the contract to accept. If the penalties are set by the protocol on the other hand complexity would be eliminated from the negotiation. For example the penalties could be fixed at a certain level by the protocol but this may lead to suboptimal results. Another method is to relate the decommitment penalty to the price of the contract. The penalty can be either a percentage or a more complex function of the contract price. Another approach is to choose the penalties so that they compensate the victim of the breach for its lost profit. In that case the agent would have an incentive to lie about the expected profit so a non-manipulable mechanism for calculating the lost profit would be necessary.

A breach close to the execution deadline of the contract or late in a negotiation is likely to be more costly to the victim of the breach since it could be hard to find someone to contract with within a short amount of time. In order to prevent such behavior the decommitment penalties can be increased over time.

Surprisingly self-interested myopic agents reach a higher social welfare quicker than cooperative myopic agents when decommitment penalties are low. The social welfare in the settings with agents that performed lookahead did not vary as much with the decommitment penalty as the social welfare in settings that consisted of myopic agents. For a short range of values of the decommitment penalty the myopic agents performed almost as well as the agents that performed lookahead.

In all of the settings studied the best way to set the decommitment penalties was to choose low penalties but ones that were greater than zero. Concerning the solution quality allowing decommitting for free is not optimal. The best strategy is to have a low decommitment penalty and a low rate of increase of the decommitment penalties.

In future research we will study longer negotiations for agents with a full lookahead to compare the opti-

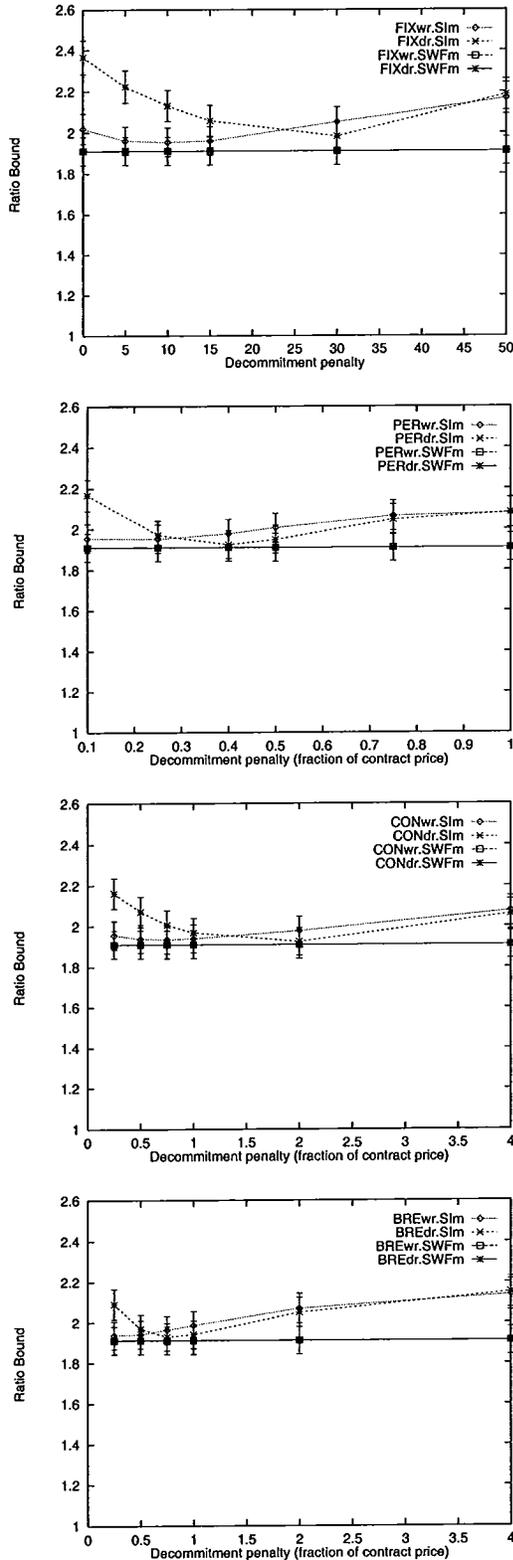


Figure 4: Myopic agents using stochastic protocols.

mal protocols for them with the ones for myopic agents. Optimal in the sense that they achieve a optimal social welfare with self-interested agents. Another topic of interest is to study domains with different types of agents performing different amounts of lookahead.

Yet another important part of our future work will be to come up with better mechanisms for deciding the decommitment penalties. In this work they have been decided by the protocol but the agents could also be allowed to conduct a negotiation over the decommitment penalties. Effective negotiation protocols for these decommitment penalties must also be developed.

References

- Andersson, M. R., and Sandholm, T. W. 1998. Leveled commitment contracting among myopic individually rational agents. In *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS-98)*.
- Andersson, M. R. 1998. Performance of leveled commitment protocols for automated negotiation: An empirical study. Master's thesis, Royal Institute of Technology, Stockholm, Sweden.
- Diamond, P. A., and Maskin, E. 1979. An equilibrium analysis of search and breach of contract, I: Steady states. *Bell Journal of Economics* 10:282-316.
- Kalakota, R., and Whinston, A. B. 1996. *Frontiers of Electronic Commerce*. Addison-Wesley Publishing Company, Inc.
- Park, S.; Durfee, E. H.; and Birmingham, W. P. 1996. Advantages of strategic thinking in multiagent contracts. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)*, 259-266.
- Posner, R. A. 1977. *Economic Analysis of Law*. Little, Brown and Company, 2nd edition.
- Raiffa, H. 1982. *The Art and Science of Negotiation*. Cambridge, Mass.: Harvard Univ. Press.
- Sandholm, T. W., and Lesser, V. R. 1996. Advantages of a leveled commitment contracting protocol. In *Proceedings of the National Conference on Artificial Intelligence*, 126-133.
- Sandholm, T. W., and Ygge, F. 1997. On the gains and losses of speculation in equilibrium markets. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, 632-638.
- Sandholm, T. W. 1996. *Negotiation among Self-Interested Computationally Limited Agents*. Ph.D. Dissertation, University of Massachusetts, Amherst.
- Sandholm, T. W. 1998. Contract types for satisficing task allocation: I theoretical results. In *AAAI Spring Symposium Series: Satisficing Models*, 68-75.
- Sen, S., and Durfee, E. 1994. The role of commitment in cooperative negotiation. *International Journal on Intelligent Cooperative Information Systems* 3(1):67-81.
- Smith, R. G. 1980. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers* C-29(12):1104-1113.
- Ygge, F., and Akkermans, J. M. 1996. Power load management as a computational market. In *Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS-96)*, 393-400.