Negotiation Over Decommitment Penalty

UMass Amherst Department of Computer Science

Technical Report # UM-CS-2011-009

Bo An
Dept. of Computer Science
University of Massachusetts, Amherst, USA
ban@cs.umass.edu

Victor Lesser
Dept. of Computer Science
University of Massachusetts, Amherst, USA
lesser@cs.umass.edu

ABSTRACT

Leveled-commitment contracting is a negotiation protocol in which an agent is allowed to be freed from a contract at the cost of paying a penalty to the other contract party. While there has been some related work on analyzing agents' strategic behavior in leveled-commitment contracting and applying decommitment contracting in practical resource allocation problems, decommitment penalties are exogenously set in the absence of mutual consensus of contract parties. Alternatively, this paper considers the role of negotiation in deciding decommitment penalties. In our model, agents negotiate over both the contract price and the amount of decommitment penalty in the contracting game and then decide whether to decommit from contracts in the decommitment game. This paper first analyzes agents' strategic behavior in both the contracting game and the decommiting game. We then examine the efficiency of negotiating over the penalty through experiments in dynamic contracting scenarios. Experimental results show that setting penalties through negotiation achieved higher social welfare than other exogenous penalty setting mechanisms for a range of contracting strategies.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems

General Terms

Algorithms, Economics, Experimentation

Keywords

Negotiation agents, leveled-commitment, penalty, equilibrium

1. INTRODUCTION

In automated negotiation systems for self-interested agents, contracts have traditionally been binding and do not allow agents to efficiently deal with future events in the environment. Sandholm and Lesser [12] proposed leveled-commitment contracts which allow an agent to be freed from an existing contract at the cost of simply paying a penalty to the other contract party. A self-interested agent will be reluctant to decommit because the other contract party might decommit, in which case the former agent gets freed from the contract, does not incur a penalty, and collects a penalty from the other party. Despite such strategic decommiting, leveled-commitment increases the expected payoffs of all contract parties and can enable deals that are impossible under full commitment [12]. This approach has been applied in a number of different applications [1, 3, 9, 10].

In leveled-commitment contracting, both contract parties strategically choose their level of commitment based on the contract price and decommitment penalty which are determined prior to the start of the decommiting game. The efficiency of leveled-commitment contracting depends on how the contract price and decommitment penalty are set. In Sandholm *et al.*'s model of leveled-commitment contracts [12–14], both the contract prices and decommitment penalties are assumed to be known to the contract parties before the decommiting game. Although algorithms are provided to optimize the social welfare of the equilibrium outcome [13], the optimization is not for the favor of each contract party. In existing applications (e.g., [1, 3, 9, 10]) of automated negotiation with decommitment, decommitment penalties are set by third parties (e.g., system designers) and are either fixed or a function of contract prices. In contrast to previous approaches of setting decommitment penalties exogenously, this paper proposes to study how to set the contract price and decommitment penalty through negotiation between the contract parties.

We propose to study negotiating simultaneously over contract prices and decommitment penalties for several reasons. First, it is difficult for system designers to decide *optimal* contract prices and decommitment penalties to maximize the social welfare, especially when there are multiple agents and agents have incomplete information. It is also intractable to compute agents' rational equilibrium strategies in many practical sequential games. Furthermore, it is not appropriate to assume that system designers have complete knowledge about agents in the system. Finally, a selfish agent may feel it is advantageous for it to decide the contract price and penalty by itself. When agents are allowed to negotiate over penalties, each agent has a larger strategy space which gives it more options for how to react to the current situation and it may be able to achieve a utility which cannot be achieved when it is not allowed to negotiate over penalty. At the same time, agents may reach a higher social welfare that is unable to be reached when agents are not allowed to negotiate over penalty.

This paper first analyzes agents' strategic behavior in the bilateral contracting game prior to the decommiting game to make agreements on a contract and a decommiting penalty. One selfish contract party may prefer another pair of contract price and decommitting penalty to the contract price and decommitment penalty which maximize the social welfare. The leveled-commitment contracting we propose includes two games: a contracting game where the two parties bargain over contract price and decommitment penalty and a decommiting game in which the two agents make strategically decommiting decisions. During the decommiting game, agents will make optimal decommiting decisions while taking into account the contract price and decommitment penalty previously agreed upon. Therefore, in the contracting game, each agent will try to make the best contract price and penalty which will maximize its utility in the decommiting game. In this paper, we analyze agents' equilibrium strategies in finite horizon contracting games where both agents don't know exactly but only statistically what offers they will receive from other agents in the future.

The other focus of this paper is experimentally investigating the efficiency of setting decommitment penalties through negotiation as compared to choosing decommitment penalties exogenously. When each agent strategically negotiates with other agents to maximize its expected utility in the decommiting game, agents may get worse social welfare as compared with not negotiating over penalties. We experimentally compare negotiating over penalties with other ways of setting penalties in the literature, e.g., fixed penalty, and penalty as a percentage of the contract price. In our experiment, a number of contractors and contractees negotiate and decommit dynamically where agents have incomplete information. Selfish agents can either carry out game theoretic partial lookahead or be myopic. Experimental results from various environments show that when decommitment penalties are decided through negotiation, agents achieved higher social welfare than other approaches of setting decommitment penalties.

The rest of this paper proceeds as follows: Section 2 studies the contracting game when agents are uncertain about their outside options. Section 3 experimentally evaluates the efficiency of setting penalties through negotiation. Section 4 concludes this paper.

2. NEGOTIATING OVER PENALTY

As in [12–14], we consider a contracting setting with two risk neutral agents who attempt to maximize their own expected payoff: contractor **b** who pays to get a task done, and contractee **s** who gets paid for handling the task. The setting can be interpreted as modeling a variety of scenarios, for example bargaining between a buyer and a seller in e-commerce. In our model, **b** and **s** negotiate over contract price and decommitment penalty before additional offers (outside offers) from other agents become available. Then they strategically choose to decommit or not when their outside offers are available.

2.1 Leveled-commitment contracting

We study a setting where the future of agents involves uncertainty. We model this as agents' potentially receiving outside offers as in [12–14]. The contractor's outside offers could come from some other contractees which can provide the service requested by the contractor. The contractor can make agreements with those contractees in the future. The contractor's best (lowest) outside offer v is characterized by a probability density function $f(v)$. The contractee's best (highest) outside offer w is characterized by a probability density function $g(w)$. $f(v)$ and $g(w)$ are assumed statistically independent and are common knowledge [12–14]. That is, both agents have symmetric information as they both don't know the value of v and w.

The contractor's options are either to make a contract with the contractee or to wait for v. Similarly, the contractee's options are either to make a contract with the contractor or to wait for w. The two agents could make a full commitment contract at some price. Alternatively, they can make a leveled-commitment contract which is specified by a contract price, ρ, and a decommitment penalty q. If one agent decommits from the agreement, it needs to pay the penalty q to the other agent.[1] When the decommitment penalty q is very large, a leveled-commitment contract is equivalent to a full contract as no agent will choose to decommit. Therefore, full commitment contracts are a subset of leveled-commitment contracts.

One implicit assumption is that during the contracting game, the contractor can only bargain with one contractee and the contractee can also only negotiate with one contractor. The other assumption is that the bargaining game finishes before outside options become available. Bargaining protocols can be used to control the length of negotiation. Moreover, even if agents are allowed to conduct infinite time negotiation, negotiation often stops soon since bargaining agents usually have deadline constraints and often face bargaining costs.

The leveled-commitment contracting consists of two stages. In the first stage, which we call the *contracting game*, the agents make agreements on both a contract price and a decommiting penalty. In the second stage, which we call the *decommiting game*, the agents decide on whether to decommit or not. Clearly, the equilibrium of the decommiting game affects the agents's preferences over contract prices and decommitment penalties in the contracting game. There is no decommiting game if agents make a $\mathbf{null}$ contract (i.e., no agreement is made) in the contracting game.

2.1.1 Contracting game

We consider the widely used alternating-offers protocol [11]. Formally, contractor $\mathbf{b}$ and contractee $\mathbf{s}$ can act at times $t \in \{0, 1, \ldots, T\}$. $T = \min(T_{\mathbf{b}}, T_{\mathbf{s}})$ where $T_{\mathbf{b}}$ and $T_{\mathbf{s}}$ are the negotiation deadline of **b** and **s**, respectively. The player function $\iota : \mathbb{N} \to \{\mathbf{b}, \mathbf{s}\}$ returns the agent that proposes at time t and is such that $\iota(t) \neq \iota(t+1)$, i.e., a pair of agents bargain by making offers in alternate fashion. As noted, there is no bargaining in the previous literature [12–14].

Possible actions $\sigma_{\mathbf{a}}^t$ of of agent $\mathbf{a} \in \{\mathbf{b}, \mathbf{s}\}$ at any time point $t \geq 0$ are defined as follows. If $\iota(t) = \mathbf{a}$, $\mathbf{a}$ can choose to 1) **offer** $[\rho, q]$ where ρ is contract price and q is decommitment penalty; or 2) $\mathbf{exit}$. If $\iota(t) \neq \mathbf{a}$, $\mathbf{a}$ can choose to 1) $\mathbf{accept}$, 2) $\mathbf{reject}$, or 3) $\mathbf{exit}$. If $\sigma_{\mathbf{a}}^t = \mathbf{accept}$, the bargaining outcome is $([\rho, q], t)$ if $\sigma_{\hat{\mathbf{a}}}^{t-1} = \mathbf{offer}\,[\rho, q]$ where $\hat{\mathbf{a}}$ is $\mathbf{a}$'s trading partner. If $\sigma_{\mathbf{a}}^t = \mathbf{exit}$ the bargaining stops and the outcome is $\mathbf{null}$ contract. Otherwise the bargaining continues to the next time point.

Each agent $\mathbf{a} \in \mathcal{A} = \{\mathbf{b}, \mathbf{s}\}$ is associated with a bargaining cost $\delta_{\mathbf{a}} > 0$ for each round. Such a bargaining cost could be incurred by agents' reasoning, communication, and waiting. We assume that the bargaining cost is relatively small as compared with the cost for finishing a task. Agents' bargaining costs are common knowledge.

2.1.2 Decommiting game

The decommitting game happens only when the two agents make a leveled-commitment contract $[\rho, q]$. In the decommiting game, each agent has exactly one chance to decommit and there are different decommiting mechanisms depending on who decommits first [12–14]: 1) contractee has to reveal its decision first; 2) contractor has to reveal its decision first; and 3) agents reveal their decisions simultaneously. We consider the first decommiting mechanism, i.e., contractee takes actions first and contractor moves next.

[1] Our analysis can be easily extended to handle the setting where the penalties for the contractor and the contractee are different. Setting different penalties for contractor and contractee only makes it difficult to solve the decommitment game, which has been thoroughly analyzed in the work by Sandholm *et al*. [12]. For the contracting game, a new variable will be added but the analysis will be the same.

The other mechanisms can be analyzed analogously.

2.2 Optimal contracts

Agents' bargaining strategies in the contracting game are affected by the outcome of the decommiting game: each agent wants to make the optimal contract that maximizes its expected utility in the decommiting game. There may be multiple optimal contracts or the optimal contract may be the $\mathbf{null}$ contract.

We follow the same analysis as in [12] to compute agents' optimal contracts. Assume that the contract made during the contracting game is $[\rho, q]$. In a sequential decommiting game where the contractee has to decommit first, if the contractee has decommited, the contractor's best move is not to decommit as $q \geq 0$. In the subgame where the contractee has not decommited, the contractor's best move is to decommit if $-v - q > -\rho$, i.e., the contractor decommits if its outside offer, v, is below a threshold $v^* = \rho - q$. So, the probability that it decommits is $p_{\mathbf{b}} = \int_{-\infty}^{v^*} f(v)dv$.

The contractee gets $w - q$ if it decommits, $w + q$ if it does not but the contractor does, and ρ if neither decommits. Thus the contractee decommits if $w - q > p_{\mathbf{b}}(w+q) + (1-p_{\mathbf{b}})\rho$. When $p_{\mathbf{b}} < 1$ the inequality above shows that the contractee decommits if its outside offer exceeds a threshold $w^* = \rho + q(1+p_{\mathbf{b}})/(1-p_{\mathbf{b}})$. So, the probability that it decommits is $p_{\mathbf{s}} = \int_{w^*}^{\infty} g(w)dw$.

Given agents' equilibrium strategies under contract $c = [\rho, q]$, $\mathbf{b}$'s expected payoff $\pi_{\mathbf{b}}(c, f(v), g(w))$ ($\pi_{\mathbf{b}}(c, f, g)$ for short) is

$$p_{\mathbf{s}}\int_{-\infty}^{\infty}(q-v)f(v)dv + (1-p_{\mathbf{s}})\Big[\int_{-\infty}^{v^*} -(v+q)f(v)dv - \int_{v^*}^{\infty}\rho f(v)dv\Big]$$

The expected payoff $\pi_{\mathbf{s}}(c, f, g)$ of contractee $\mathbf{s}$ is

$$\int_{w^*}^{\infty} g(w)(w-q)dw + \int_{-\infty}^{w^*} g(w)\big[p_{\mathbf{b}}(w+q) + (1-p_{\mathbf{b}})\rho\big]dw$$

If agents fail to make a contract, an agent can wait for its best outside offer. Thus, agents' expected utilities under the $\mathbf{null}$ contract are $\pi_{\mathbf{b}}(\mathbf{null}, f, g) = \int_{-\infty}^{\infty} -f(v)vdv = -E(v)$ and $\pi_{\mathbf{s}}(\mathbf{null}, f, g) = \int_{-\infty}^{\infty} g(w)wdw = E(w)$.

We assume that agents are *individually rational* (IR), i.e., no agent will accept a contract worse than the $\mathbf{null}$ contract. A contract c is IR if it is individually rational for both agents. Formally, the set $\mathcal{C}(f, g)$ of IR contracts based on agents' beliefs $f(v)$ and $g(w)$ are

$$\{c | \pi_{\mathbf{b}}(c, f, g) \geq -E(v), \pi_{\mathbf{s}}(c, f, g) \geq E(w)\}$$

We assume that $\mathcal{C}(f, g)$ is not empty. The contract $c^*_{\mathbf{b}}(f, g)$ ($c^*_{\mathbf{s}}(f, g)$) which maximizes the contractor's (contractee's) expected utility is the contractor's (contractee's) optimal contract. Formally,

$$\pi_{\mathbf{b}}(c^*_{\mathbf{b}}(f, g), f, g) = \max_{c \in \mathcal{C}(f,g)} \pi_{\mathbf{b}}(c, f, g)$$
$$\pi_{\mathbf{s}}(c^*_{\mathbf{s}}(f, g), f, g) = \max_{c \in \mathcal{C}(f,g)} \pi_{\mathbf{s}}(c, f, g)$$

Therefore, the utility $\mathbf{a}$ can get is in the range $[\pi_{\mathbf{a}}(\mathbf{null}, f, g), \pi_{\mathbf{a}}(c^*_{\mathbf{a}}(f, g), f, g)]$. For any value x in the range $[\pi_{\mathbf{a}}(\mathbf{null}, f, g), \pi_{\mathbf{a}}(c^*_{\mathbf{a}}(f, g), f, g)]$, we assume that there is always a contract c such that $\pi_{\mathbf{a}}(c, f, g) = x$. That is, the contract space is continuous. Furthermore, we assume that selfish agents are benevolent in the sense that when an agent is indifferent among several offers, it always proposes the offer giving its trading partner the highest payoff. Based on this analysis developed previously by Sandholm *et al.* [12–14], we now extend it to a contracting game which agents negotiate over contract prices and decommitment penalties.

2.3 Agents' Equilibrium Strategies

In this setting, agents have symmetric information and thus the appropriate solution concept for the bargaining game is the subgame perfect equilibrium. In subgame perfect equilibrium, agents' strategies are in equilibrium in every possible subgame. Such a solution can be found by backward induction [7].

Initially, it is determined the time point T where the game rationally stops: it is $T = \min(T_{\mathbf{b}}, T_{\mathbf{s}})$. The equilibrium outcome of every subgame starting from $t > T$ is $\mathbf{null}$, since at least one agent will have exited from the negotiation. Therefore, at $t = T$ agent $\mathbf{a} = \iota(T)$ would propose its best offer which is acceptable to $\hat{\mathbf{a}}$ as $\hat{\mathbf{a}}$ would accept any offer which gives it a utility not worse than $\mathbf{null}$, namely, any offer c such that $\pi_{\hat{\mathbf{a}}}(c, f, g) \geq \pi_{\hat{\mathbf{a}}}(\mathbf{null}, f, g)$. Thus, the optimal offer $c^*(T)$ of $\iota(T)$ is $c^*_{\iota(T)}(f, g)$. From $t = T$ back to $t = 0$ it is possible to find the optimal offer agent $\iota(t)$ can make at t, if it makes an offer, and the offers that it would accept. $c^*(t)$ denotes the optimal offer of agent $\iota(t)$ at t. $c^*(t)$ is the offer which generates the highest utility for agent $\iota(t)$ under the condition that $c^*(t)$ is acceptable to agent $\iota(t+1)$ at time t. Agent $\iota(t+1)$ accepts $c^*(t)$ at time t if 1) accepting $c^*(t)$ is no worse than exiting from negotiation, and 2) accepting $c^*(t)$ is no worse than rejecting the offer and making $c^*(t+1)$ at time $t+1$ (if it's possible). The equilibrium strategy of any sub-game starting from $0 \leq t \leq T$ prescribes that agent $\iota(t)$ offers $c^*(t)$ at t and agent $\iota(t+1)$ accepts it at t.

Backward propagation is used to provide a recursive formula for $c^*(t)$: 1) If $t = T$, $c^*(t) = c^*_{\iota(t)}(f, g)$, i.e., the proposing agent at T can propose its optimal contract; 2) If $t < T$, the proposing agent will propose its best offer which is acceptable to agent $\iota(t+1)$. Formally, $c^*(t)$ is such that $\pi_{\iota(t+1)}(c^*(t), f, g) = \max\{\pi_{\iota(t+1)}(\mathbf{null}, f, g), \pi_{\iota(t+1)}(c^*(t+1), f, g) - \delta_{\iota(t+1)}\}$. Throughout this paper, we assume that such $c^*(t)$ exists and it generates a utility no worse than exiting from negotiation . The values of $c^*(t)$ can be calculated recursively from $t = T$ back to $t = 0$ applying backward induction.

Finally, agents' equilibrium strategies can be defined as follows:

- $\mathbf{a} = \iota(t)$: $\mathbf{a}$ proposes $c^*(t)$ at time $t \leq T$ and exits negotiation at time $t > T$.
- $\mathbf{a} \neq \iota(t)$: At time $t < T$, $\mathbf{a}$ accepts the offer c from $\iota(t)$ if $\pi_{\mathbf{a}}(c, f, g) \geq \pi_{\mathbf{a}}(c^*(t+1), f, g) - \delta_{\mathbf{a}}$ and rejects otherwise. At time $T \leq t \leq T_{\mathbf{a}}$, $\mathbf{a}$ accepts the offer c from $\iota(t)$ if $\pi_{\mathbf{a}}(c, f, g) \geq \pi_{\mathbf{a}}(\mathbf{null}, f, g)$ and rejects otherwise. At time $t > T_{\mathbf{a}}$, exits negotiation.

Therefore, at equilibrium, two agents will reach an agreement at time $t = 0$ and the agreement contract is $c^*(0)$. Agents' equilibrium strategies depend on many factors, e.g., outside offers in the future, bargaining costs, and order of proposing.

The social welfare of the equilibrium contract is $\pi_{\mathbf{b}}(c^*(0), f, g) + \pi_{\mathbf{b}}(c^*(0), f, g)$ which may be lower than the optimal social welfare $\max_{c \in \mathcal{C}(f,g)}(\pi_{\mathbf{b}}(c, f, g) + \pi_{\mathbf{s}}(c, f, g))$. When the decommitment penalty is set exogenously, the system designer can compute the optimal penalty which maximizes the social welfare of the equilibrium contract by solving the contracting game. In more realistic scenarios in which there are usually more than two agents and agents have more uncertainties, it may be intractable to compute the optimal penalty to optimize the social welfare and it is important to investigate the benefit of negotiating over penalty.

3. EFFICIENCY OF NEGOTIATING OVER PENALTY

In the previous sections, we provide agents' equilibrium strategies in the contracting settings in which selfish agents strategically make offers and respond to received offers. While selfish agents are only interested in maximizing their individual objective functions, the overall efficiency of the outcomes of decentralized strategic interactions, can be worse than social optimal solutions formed by a central authority maximizing aggregate social welfare. It's desirable to compare the efficiency of the negotiation based mechanism for setting penalties with some existing methods for deciding penalties [1, 3, 4, 9, 10]. While the previous sections consider a simple two agent contracting scenario, in this section, we consider more realistic bargaining scenarios where there are multiple agents which have incomplete information about others. We consider the case of contracting between multiple agents since most practical scenarios there is more than two agents. Therefore, doing experiments in realistic environments makes our results more convincing. If there are only two agents, we can mathematically compute the contracting outcome. In contrast, in the multi-agent environment, it's intractable for each agent to mathematically derive its equilibrium strategy [7] and thus we evaluate the benefits of negotiating over penalty through experimentation. Different contracting settings are experimentally compared using agents of different types: myopic agents or agents that can carry out game-theoretic partial lookahead whose design took into account the important factors affecting agents' equilibrium strategies analyzed in Section 2.3. Bargaining based penalty is also compared with some representative approaches for settings penalties exogenously [1, 3, 4, 9, 10].In all of the settings, setting decommitment penalties based on bargaining outperform setting penalties exogenously.

To make our comparison fair, we tested out approach in a wide range of negotiation environments. For space limitations, we provided average results over a wide range of scenarios. However, the results for each scenario class indicated that our approach led to better performance. With respect to practicality, we consider the negotiation management component [1] for Collaborating, Autonomous Stream Processing systems (CLASP) [5], which has been designed and prototyped in the context of System S project [8] within IBM Research to enable sophisticated stream processing. There are multiple sites running the System S software, each with their own administration and goals. Different sites may have different processing capabilities, so cooperation among these sites can frequently be of mutual benefit. For each resource, there are multiple resource providers and resource competitors. There is a cost associated with a site's providing a resource and the costs of different sites could be different. Each site has a limited number of contracting opportunities since a site has to acquire resources by a deadline.

3.1 Contracting Protocol

In the experiments, the agents are divided into two subsets: contractors and contractees. Each contractor has one task to finish and has a cost associated with the task. Each contractee has no task initially and also has a cost to handle a task. A contractor can either complete its task by itself or contract out its task to a contractee who could also handle the task.

The negotiation protocol is sequential, i.e., only one contractor and one contractee negotiate in each round. When a pair of contractor and contractee was chosen to conduct negotiation at one round, the contractor makes a proposal to the contractee and the proposal includes both a contract price and a decommitment penalty. The contract price is chosen by the contractor. The decommitment penalty depends on the penalty setting mechanism and it could be either set exogenously or decided by the contractor. The contractee can either accept the proposal or reject the proposal. The two agents make a new contract if the contractee accepts the proposal from the contractor. If an agent has already a contract at the start of the negotiation, it has to decommit from the existing contract if it makes a new contract with the other agent. If the contractor has no agreement at the end of negotiation, it can either handle the task by itself or consider contracting out the task in the future negotiations.

3.2 Agent design

For each type of agents (contractors or contractees), we developed two kinds of agents: myopic and partially lookahead. In our experiments, there may be a large number of agents and agents have incomplete information (e.g., agents don't know the costs of other agents, agents don't know the negotiation order), it is computationally intractable (even impossible) to solve the sequential bargaining game by searching the game tree. Alternatively, following the *negotiation decision functions* paradigm [6], we design agents which take a partially lookahead strategy in which each agent approximate future events and outside options. The other kind of agents are myopic agents which never look ahead.

3.2.1 Myopic agents

Contractee: A myopic contractee accepts an offer if and only if it can gain some immediate payoff by accepting the offer. Specifically, if $\mathbf{s}$ has no contract, $\mathbf{s}$ will accept an offer $c = [\rho, q]$ if its immediate gain is higher than 0, i.e., the offering price ρ is no less than cost $c_{\mathbf{s}}$. If $\mathbf{s}$ has a contract $c' = [\rho', q']$ before receiving a new offer $c = [\rho, q]$, $\mathbf{s}$ will accept the offer $c = [\rho, q]$ if $\rho - q' > \rho'$.

Contractor: A myopic contractor $\mathbf{b}$ gradually increases its offering price when it fails to make a contract. That is, a contractor takes a time-dependent strategy when it has no contract [6]. Its offering price in its i^{th} negotiation is $\int_{c_{\mathbf{s}}^{\min}}^{c_{\mathbf{s}}^{\max}} x f_{\mathbf{s}}(x)dx + (c_{\mathbf{b}} - \int_{c_{\mathbf{s}}^{\min}}^{c_{\mathbf{s}}^{\max}} x f_{\mathbf{s}}(x)dx) \cdot i/\tau$ where τ is the number of negotiation opportunities, $\int_{c_{\mathbf{s}}^{\min}}^{c_{\mathbf{s}}^{\max}} x f_{c_{\mathbf{s}}}(x)dx$ is a contractee's expected cost, $c_{\mathbf{b}}$ is the cost of $\mathbf{b}$, $f_{\mathbf{s}}(x)$ is the probability density function of the cost of a contractee, and $c_{\mathbf{s}}^{\min}$ and $c_{\mathbf{s}}^{\max}$ are the minimum and maximum cost of a contractee, respectively. Therefore, a contractor without a contract will increase its offer with the decrease of its contracting opportunities.

If $\mathbf{b}$ already has one contract $c' = [\rho', q']$, it will make a new contract $c = [\rho, q]$ such that $\mathbf{b}$'s utility will be immediately increased by $\varepsilon > 0$ if the offer is accepted by the contractee, i.e., $\rho' - q' - \rho = \varepsilon$. In our experiments, ε is randomly selected from $[0, c_{\mathbf{s}}^{\max}/5]$.

If the penalty is exogenously decided, the penalty q of an offer is decided based on the offering price. If the penalty is decided through negotiation, contractor $\mathbf{b}$ decides the penalty considering the offering price: $q = (c_{\mathbf{b}} - \rho)/c_{\mathbf{b}} \cdot c_{\mathbf{s}}^{\max}$. That is, a lower price corresponds to a high penalty, which is intuitive as a selfish agent does not hope a contract with a low contract price to be decommited. Note that the myopic strategy is independent of information about the number of contractors and contractees.

3.2.2 Partially lookahead agents

We design a lookahead bargaining strategy based on 1) the competition between contractors and contractees, and 2) agents' multiple opportunities to make a contract. Both market competition and trading opportunities affect an agent's outside offers and bargaining power. Furthermore, trading opportunities can be used to model bargaining costs in the sense that the number of an agent's trading opportunities decreases as time goes on.

As an example, consider the strategies of contractor $\mathbf{b}$. There are two situations: $\mathbf{b}$ has one contract before making an offer to $\mathbf{s}$ or

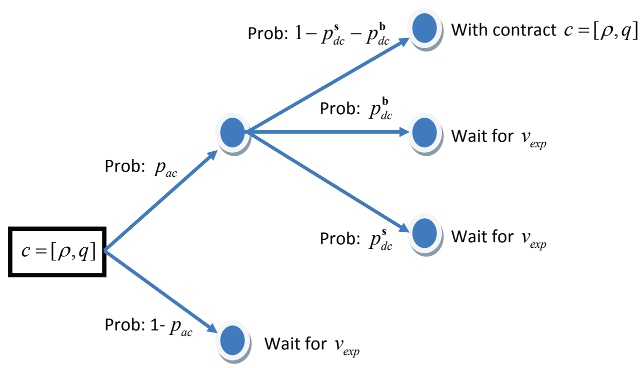

Figure 1: Game tree after b makes an offer to s (b has no contract before making the offer).

not. Let's consider the case that $\mathbf{b}$ has no contract at the start of making an offer to $\mathbf{s}$. Assume that $\mathbf{b}$ proposes a contract $c = [\rho, q]$ (the game tree after this proposing action is shown in Fig. 1). Then there are two possibilities: $\mathbf{s}$ accepts the proposal with probability p_{ac} and $\mathbf{s}$ rejects the proposal with probability $1 - p_{ac}$. If $\mathbf{s}$ rejects the proposal, $\mathbf{b}$ can still make an agreement with other contractees in the future and the expected contract price is v_{exp}. If $\mathbf{s}$ accepts the proposal, the probability that $\mathbf{s}$ will decommit from the contract is $p^{\mathbf{s}}_{dc}$. If $\mathbf{s}$ accepts the proposal, the probability that $\mathbf{b}$ will decommit from the contract is $p^{\mathbf{b}}_{dc}$. After the contract is decommited, $\mathbf{b}$ may contract out the task to other contractees in the future. Therefore, the expected final cost $u_{\mathbf{b}}([\rho, q])$ of making offer $c = [\rho, q]$ is

$$(1-p_{ac})v_{exp}+p_{ac}\big(p^{\mathbf{s}}_{dc}(v_{exp}-q)+p^{\mathbf{b}}_{dc}(v_{exp}+q)+(1-p^{\mathbf{s}}_{dc}-p^{\mathbf{b}}_{dc})\rho\big)$$

The calculation of the expected final cost of an offer $c = [\rho, q]$ depends on values of p_{ac}, v_{exp}, $p^{\mathbf{s}}_{dc}$, and $p^{\mathbf{b}}_{dc}$, which are approximated based on the expected contracting price and the number of contracting opportunities. Agents' expected contracting price is approximated by considering the market competition between contractors and contractees. Consider that one contractor and one contractee are conducting one-shot negotiation: contractor proposes first and contractee accepts or rejects the contractor's offer. Then the contract price is the contractee's cost if the contractee's cost is known to the contractor. Consider that some buyers and sellers are negotiating with each other. With more contractees, the contracting price will decrease due to the competition between contractees. Analogously, with more contractors, the contracting price will go up since the competition among contractors will increase. While making a proposal, each strategic agents needs to consider the expected contracting price, which is estimated by considering market competition between contractors and contractees. The probability distribution for any contractor's cost is $F_{\mathbf{b}}(.)$, where $F_{\mathbf{b}}(y)$ denotes the probability that the cost of a contractor b is no greater than y. The probability density function of $F_{\mathbf{b}}(y)$ is denoted by $f_{\mathbf{b}}(y)$.

Let $F^k_{\mathbf{b}}(y)$ be the probability distributions of the k^{th} highest cost among all the contractors. The probability density function of $F^k_{\mathbf{b}}(y)$ is denoted by $f^k_{\mathbf{b}}(y)$. $F^1_{\mathbf{b}}(y)$ is equal to the product of the probabilities that the cost is higher than or equal to y for each contractor. $F^2_{\mathbf{b}}(y)$ is equal to $F^1_{\mathbf{b}}(y)$ plus the probability that the highest maximum reserve price is greater than y, and the second highest maximum reserve price is less than or equal to y. These probabilities can be calculated as follows (assume that there are n contractors and m contractees):

$$F^1_{\mathbf{b}}(y) = \big(F_{\mathbf{b}}(y)\big)^n$$
$$F^2_{\mathbf{b}}(y) = F^1_{\mathbf{b}}(y) + C^1_n\big(1 - F_{\mathbf{b}}(y)\big)^{2-1}\big(F_{\mathbf{b}}(y)\big)^{n-1}$$
$$F^k_{\mathbf{b}}(y) = F^{k-1}_{\mathbf{b}}(y) + C^{k-1}_n\big(1 - F_{\mathbf{b}}(y)\big)^{k-1}\big(F_{\mathbf{b}}(y)\big)^{n-k+1}$$

The corresponding probability density functions are:

$$f^1_{\mathbf{b}}(y) = n\big(F_{\mathbf{b}}(y)\big)^{n-1}$$
$$f^2_{\mathbf{b}}(y) = f^1_{\mathbf{b}}(y) - C^1_n f_{\mathbf{b}}(y)\big(F_{\mathbf{b}}(y)\big)^{n-1} + C^1_n (n-1)f_{\mathbf{b}}(y)\big(1 - F_{\mathbf{b}}(y)\big)^{2-1}\big(F_{\mathbf{b}}(y)\big)^{n-2}$$
$$f^k_{\mathbf{b}}(y) = f^{k-1}_{\mathbf{b}}(y) - C^{k-1}_n(k-1)f_{\mathbf{b}}(y)\big(1 - F_{\mathbf{b}}(y)\big)^{k-2}\big(F_{\mathbf{b}}(x)\big)^{n-k+1} + C^{k-1}_n f_{\mathbf{b}}(y)\big(1 - F_{\mathbf{b}}(y)\big)^{k-1}\big(F_{\mathbf{b}}(y)\big)^{n-k+1}$$

In the same way, we can get $G^k_{\mathbf{s}}(y)$, the probability distribution of the k^{th} lowest cost among all the contractees. The probability density function of $G^k_{\mathbf{s}}(y)$ is denoted by $g^k_{\mathbf{s}}(y)$.

We provide a heuristic approach to estimate the expected agreement price c_{ex}. When the number n of contractors is higher than the number m of contractees, the contracting price is the $(m+1)^{st}$ highest cost of the contractors. Otherwise, the contracting price is the n^{st} lowest cost of the contractees. Formally,

$$c_{ex} = \begin{cases} \int_{c^{\min}_{\mathbf{b}}}^{c^{\max}_{\mathbf{b}}} f^{m+1}_{\mathbf{b}}(y)ydy & \text{if } n > m \\ \int_{c^{\min}_{\mathbf{s}}}^{c^{\max}_{\mathbf{s}}} f^{n}_{\mathbf{s}}(y)ydy & \text{if } n \leq m \end{cases}$$

where $c^{\min}_{\mathbf{a}}$ and $c^{\max}_{\mathbf{a}}$ are the minimum and maximum costs of agent $\mathbf{a}$, respectively.

Approximation of v_{exp}, p_{ac}, $p^{\mathbf{s}}_{dc}$, and $p^{\mathbf{b}}_{dc}$: We let v_{exp} be c_{ex}. While receiving an offer $[\rho, q]$, the contractee $\mathbf{s}$ will take both offering price ρ and penalty q into account. If ρ is too small, $\mathbf{s}$ will not accept the offer. If the penalty is too high and $\mathbf{s}$ accepts the offer, $\mathbf{s}$ will not have a chance to decommit from the contract in the future. Therefore, p_{ac} is defined as $\min((\rho-q/2)/c_{ex}, 1)$. If $\mathbf{s}$ accepts $\mathbf{b}$'s offer $[\rho, q]$, $\mathbf{s}$ will decommit from the contract if another contractor can pay no less than $\rho + q$. Given that the expected contracting price c_{ex}, the probability that $\mathbf{s}$ decommits from the contract is 0 if $\rho + q \geq c_{ex}$. Otherwise, $p^{\mathbf{s}}_{dc} = 1/2$ as it's possible that the contractor decommits before the contractee decommits and we assume that both agents have the same probability of decommiting first. If $\mathbf{s}$ has no negotiation opportunities in the future or $\rho - q < c_{ex}$, $p^{\mathbf{b}}_{dc} = 0$. Otherwise, $p^{\mathbf{b}}_{dc} = 1/2$ as it's possible that $\mathbf{s}$ decommits first.

Given the definition of the expected cost of each proposal, the optimal proposal $c^* = [\rho^*, q^*]$ is the proposal minimizing its expected cost. That is, for any offer $c = [\rho, q]$, it follows that

$$u_{\mathbf{b}}([\rho^*, q^*]) \leq u_{\mathbf{b}}([\rho, q])$$

Therefore, $\mathbf{b}$ will search all possible values of ρ and q to find out the best offer if the penalty setting mechanism is "bargaining". As the search space is infinite, in our experiments we only search all integer values. If the decommitment penalties are determined exogenously, $\mathbf{b}$ only searches the possible values of ρ.

The situation that $\mathbf{b}$ has one contract while making an offer can be analyzed in the same way. When $\mathbf{b}$ has one contract, it can still use the existing contract if its new offer is rejected.

A contractee acts as follows: If it has no contract while receiving an offer $[\rho, q]$, it will accept the offer if and only if its utility while accepting the offer is higher than that while not accepting the offer. If it rejects the offer, it can make a contract with the expected contracting price in the future. If it accepts the offer, there are two

Table 1: Variables

Variables	Values
Number of contractors (n)	$[2, 20]$
Number of contractees (m)	$[2, 20]$
Contracting opportunities of each contractor (τ)	$[1, 8]$
Cost of a contractor	$[50, 100]$
Cost of a contractee	$[0, 50]$
Fixed penalty	$\{0, 10, 20, 40\}$
Penalty rate (percentage of contract price)	$\{0.1, 0.3\}$

Table 2: Average cost ratios

Strategy	All Myopic	All Lookahead	Random match
Bargaining	2.161	3.109	2.775
Fixed penalty-0	3.837	3.844	3.778
Fixed penalty-10	2.618	3.573	3.252
Fixed penalty-20	2.529	3.573	3.262
Fixed penalty-40	2.653	3.627	3.357
Price rate-0.1	3.355	3.573	3.518
Price rate-0.3	2.541	3.547	3.174

situations in the future: one agent decommits from the contract or no agent decommits from the contract. $\mathbf{s}$ can compute its expected utility for accepting the offer in the same way as computing the expected cost when the contractor is making an offer. The situation when $\mathbf{s}$ has a contract while receiving an offer $[\rho, q]$ can be analyzed analogously.

3.3 Experimental settings

Experiment parameters: Extensive experiments were conducted in a variety of scenarios subjected to parameters in Table. 1. The experiment parameters reflect the design of the Collaborating, Autonomous Stream Processing systems [1, 5]. The number of contractors and contractees were randomly selected from $[2, 20]$. Each contractor may have multiple opportunities to make contracts with contractees. The number of contracting opportunities of each contractor is in the range of $[1, 8]$. Each agent is assumed to know the number of agents of different types and the number of contracting opportunities of each contractor. Initially, each agent was randomly assigned a cost for handling its task. The contractors' costs were drawn uniformly in the interval [50, 100] and each contractee's cost is in the interval [0, 50]. The contractors never negotiate with each other. The contractees also never negotiate with each other.

We compare the contracting results when agents' decommitment penalties are determined by negotiation and results where decommitment penalties are determined exogenously. While there are many methods to exogenously set decommitment penalties [1, 3, 4, 9, 10], the following two approaches are the most widely used: 1) fixed penalty independent of contract prices and 2) penalty as a percentage of contract prices. We compare our negotiation based approach with the above two approaches. For fixed penalty, the penalty is chosen from $\{0, 10, 20, 40\}$. When the decommitment penalty is a percentage of a contract price, the rate is chosen from $\{0.1, 0.3\}$, i.e., $q = 0.1\rho$ or $q = 0.3\rho$. Thus, there are 7 approaches to set penalties: 4 fixed penalty values, 2 penalty functions in which a penalty is a fraction of contract price, and the bargaining approach. For the bargaining approach, the contractor can choose penalties freely.

Extensive stochastic simulations were carried out for all the combinations of variables in Table. 1. For each setting, we randomly generated 100 instances and for each instance, there is only one contractor and one contractee negotiating with each other at each time period. The order in which the agents meet for negotiation is random. That is, in each round of the bargaining game, one contractor and one contractee are randomly picked to negotiate with each other. For example, in a setting where there are two contractors and two contractees, and each contractor has two negotiation chances. Agents' costs in each instance are randomly generated. The negotiation order in each instance is randomly generated and the maximum number of negotiation rounds is determined by the number of contractors and the number of each contractor's negotiation opportunities. For each instance, we try the 7 different approaches to set decommitment penalties and 3 different strategy combinations: 1) all agents choose the myopic strategy, 2) all agents choose the partial lookahead strategy, and 3) each agent randomly chooses a strategy.

Performance measure-cost ratio: The social welfare sw is simply $-sc$ where sc is the sum of agents' costs. After each experiment, we measure the ratio of the social welfare of the solution obtained through negotiation to the optimal social welfare. For each instance, the optimal social welfare can be calculated in polynomial time: recursively make a contract between the contractor with the highest cost and the contractee with the lowest cost until all contractors or contractees have contracts. The average cost ratio for all instances is calculated for each setting. The lower cost ratio, the better. For each setting, we compare the cost ratio when agents use different strategy combinations and different approaches to set penalties.

3.4 Observations

Observation 1: Table 2 summarizes the average cost ratios in all settings when the contractor/contractee ratio is within the range $[1/3, 3]$. We found that on average, negotiating over penalty achieved lower cost ratio as compared with exogenous methods for setting penalties, no matter which strategies were used by agents. When agents can negotiate over penalties, they have a much larger strategy space. Although agents are optimizing the utilities of their own, agents still achieved higher social welfare than setting penalties exogenously. Furthermore, when the decommitment penalty is 0, the cost ratio is higher than any other exogenous methods for setting decommitment penalties. When there is no decommitment penalty, each agent can decommmit from contracts for free and thus contract breach may happen forever. Throughout our experiment, when there is no penalty, the number of decommitments is 4 times as large as that when the decommitment penalty is 10. When the penalty is too high (e.g., fixed penalty 40), decommitment penalty is not going to happen frequently and the initial contract of each agent is likely to be its final contract. In Table 2, the cost ratio for very high fixed penalty (e.g., 40) is higher than that for a moderate fixed penalty (e.g., 20).

The cost ratio when all agents taking myopic strategies is lower than the cost ratio when agents take lookahead strategies or randomly determine choose lookahead strategies or myopic strategies (called *random match*). Furthermore, agents with random strategies achieved lower cost ratio than agents with lookahead strategies. Myopic contractees will accept an offer if it can gain some immediate payoff. In contrast, contactees perform lookahead will take market competition into account and will accept an offer if and only it can gain some utility increase from the long-run (see Section 3.5 for further analysis). Therefore, a contract is more likely to happen when a contractee uses the myopic strategy as compared with using the lookahead strategy. In our experiments, when all agents take the myopic strategy, the average number of contract-

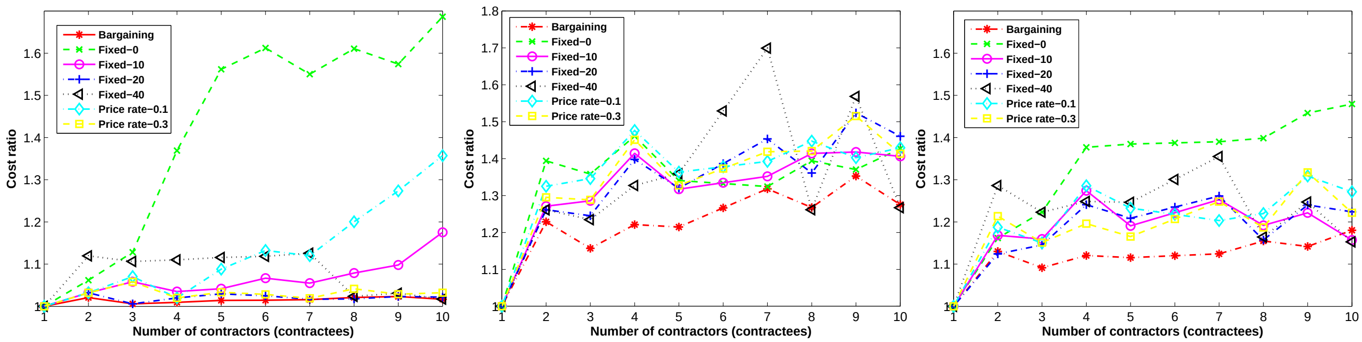

(a) Myopic agents (b) Lookahead agents (c) Random match

Figure 2: Number of contractors and cost ratio (contractor/contractee ratio: 1)

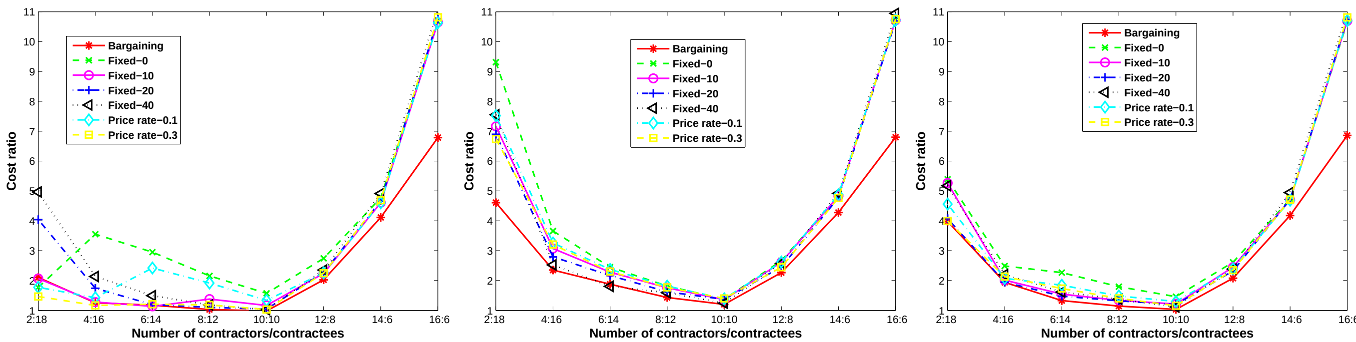

(a) Myopic agents (b) Lookahead agents (c) Random match

Figure 3: contractor/contractee ratio and cost ratio

s per contractor is 27% higher than that when all agents take the lookahead strategy.

Observation 2: Fig. 2 shows the cost ratio with different number of contractors when the number of contractors are equal to the number of contractees. In all the settings, negotiating over penalty achieved lower cost ratio as compared with exogenous methods for setting penalties. It's observed that the cost ratio increases with the increase of number of agents, especially when the penalty is low and agents use the myopic strategy.

Observation 3: It can be observed from Fig. 3 that with different contractor/contractee ratios, negotiating over penalty achieved lower cost ratio as compared with exogenous methods for setting penalties. The cost ratio decreases with the increase of contractor/contractee ratio when the contractor/contractee ratio is low. However, the cost ratio increases with the increase of contractor/contractee ratio when the contractor/contractee ratio is higher than 1. In general, due to the existence of decommitment penalties, a contractee having already a contract is less likely to accept a new offer as compared with when it has no contract. Thus, generally the first contract made by each contractee has a high probability of being its final contract. When there is big difference between the number of contractors and contractees, the initial contracts made by all agents could be "bad" as compared with the socially optimal solution. For example, assume that there are 10 contractees and 1 contractor. If the contractor is chosen to negotiate with a contractee with relatively high cost as compared with other contractees. Assume that they make a contract and no decommitment happens. The solution may be much worse than the socially optimal solution in which the contractor makes a contract with the contractee with the lowest cost. In contrast, when the number of contractors is equal to the number of contractees, it's more likely that each contractor makes one contract, which is the socially optimal solution.

Observation 4: Fig. 4 shows the change of cost ratio with negotiation opportunities. With different number of negotiation opportunities, negotiating over penalty achieved lower cost ratio as compared with exogenous methods for setting penalties. When all agents are myopic, the number of negotiation opportunities has no significant impact on the cost ratio. But when agents perform lookahead, the cost ratio decreases with the increase of negotiation opportunities when the number of negotiation opportunities is low. This is probably because each contractee will take market competition into account while deciding whether to accept an offer.

Observation 5: When decommitment penalties are set through negotiation, lookahead agents search the whole contract space (including contract prices and penalties) and find the best contract. We analyzed the relations between contract price, penalties, and remaining negotiation opportunities. We found that: 1) the contract price of any contractor without a contract increases with the decrease of remaining negotiation opportunities; and 2) when the contract price of a contract is low, the decommitment penalty is high. Those observations correspond to the heuristics for setting the penalty of myopic agents. Fig. 6(a) shows that when each contractor has 8 negotiation opportunities, an agent without a contract gradually increases its offering price with the decrease of negotiation opportunities. Fig. 6(b) shows that the offering penalties by lookahead contractors are lower with the decrease of negotiation opportunities. Fig. 6(c) shows the relation between proposed penalties and prices. When the ratio of offering penalty to offering price decreases with the increase of offering price, which is intuitive as contractors don't want to lose favorable contracts.

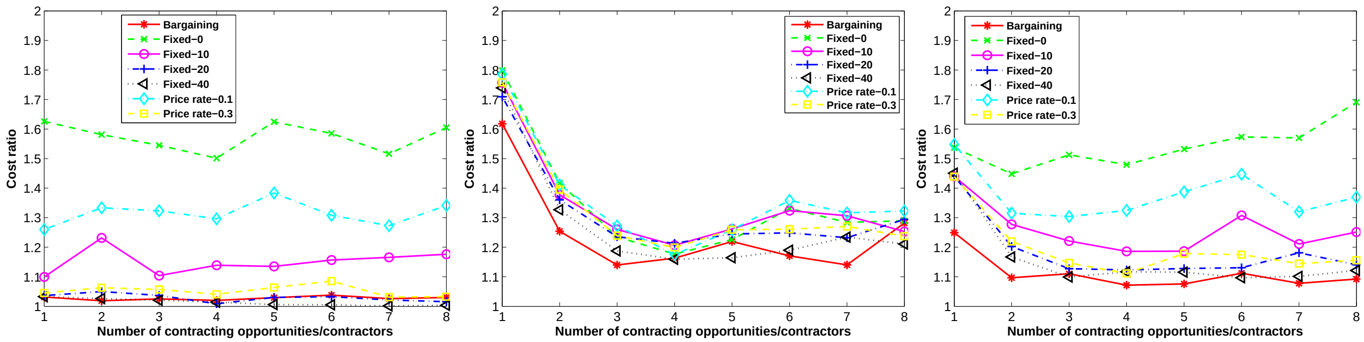

(a) Myopic agents (b) Lookahead agents (c) Random match

Figure 4: Number of contracting opportunities and cost ratio

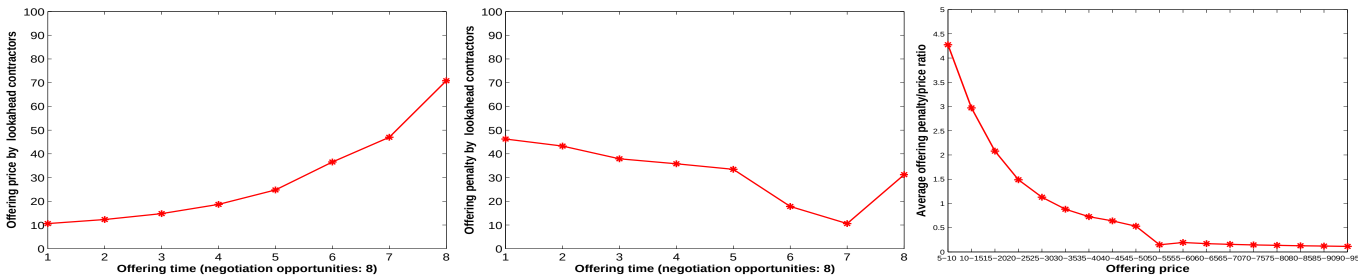

(a) Offering price and offering time (b) Offering penalty and offering time (c) Offering price and offering penalty

Figure 5: Offers by contractors perform lookahead

3.5 Further analysis

We assume that a lookahead agent knows 1) the number of contractors and contractees, and 2) the number of negotiation opportunities. We examined how this assumption will affect the efficiency of leveled-commitment contracting. When the true number of negotiation opportunities is half of the contractors' belief about their negotiation opportunities, the cost ratio of lookahead agents decreased by 6% but setting penalty through negotiation still achieved a lower cost ratio than other methods for setting penalties. Note that the strategy of myopic agents doesn't depends on agents' knowledge of the number of contractors and contractees. Therefore, even if we remove the first assumption, negotiation over penalty is still a better option than setting penalties exogenously.

In our previous experiments, contractors using the *same* penalty setting method compete with each other and make contracts with contractees. Then the social welfare for different penalty setting approaches were compared with each other. From market designers' perspective, it's better to allow agents to negotiate over penalty rather than setting a fixed penalty (function). If each selfish contractor has an opportunity to select a method to set its penalty, is it better to comply to the penalty suggested by the market designer? To answer this question, we let contractors with *different* penalty setting methods compete with each other and compare their final costs. Table. 3 shows the average cost of contractors with different penalty setting methods. For each instance, we first created 7 contractors with the same cost and each contractor uses a different method to set its penalty, then the 7 contractors compete with each other and make contracts with 7 contractees. We can see that no matter which strategy is used by agents, contractors not using a penalty suggested by the system designer achieved lower cost.

For the experimental results in Table. 3, all contractors in each experiment use the same strategy and lookahead contractors achieved higher cost than myopic contractors. This is mainly due to the indirect competition between contractors. As compared with lookahead agents, myopic agents are "cooperative". When strategic agents compete with each other, agents' performance could be much worse than agents behave in a cooperative way. We also evaluated agents' performance when agents using different negotiation strategies compete with each other. We observed that, when myopic contractors compete against lookahead contractors in each experiment, the average cost of myopic contractors is higher than that of lookahead contractors. Fig. 6 compares the average cost of myopic contractors and lookahead contractors when they compete with each other in each experiment where there are equal number of contractors and contractees. We examined the contractors' cost while changing 1) the strategies of contractees and 2) the probability that each contractor uses the myopic strategy. When the probability that each contractor uses the myopic strategy is 0, all contractors use the lookahead strategy. In contrast, when the probability that each contractor uses the myopic strategy is 1, all contractors are myopic. We can see that independent of the strategies of sellers, lookahead contractors achieved lower average cost than that of myopic contractors. In addition, contractors' cost when contractees use the lookahead strategy is higher than contractors' cost when contractees use the myopic strategy.

4. CONCLUSION

This paper complements previous work on leveled-commitment contracting by integrating a strategic contracting game with the leveled decommiting game. Agents are allowed to make leveled-commitment contracts. This paper analyzes the optimal contracts in the decommiting game and agents' equilibrium strategies in finite horizon negotiation with bargaining costs. We also experimentally evaluated the efficiency of negotiating over penalty as

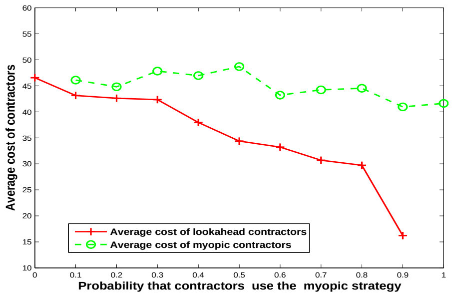

(a) Contractees use the myopic strategy

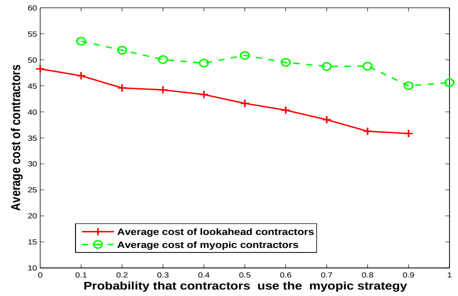

(b) Contractees randomly choose their strategies

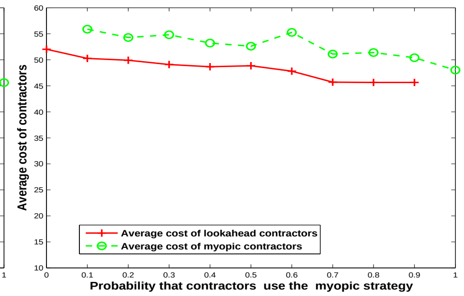

(c) Contractees use the lookahead strategy

Figure 6: Average cost of contractors when contractors with different strategies compete with each other

Table 3: Average cost of contractors

Strategy	Myopic	Lookahead	Random match
Bargaining	40.24	42.05	41.23
Fixed penalty-0	65.27	63.84	64.13
Fixed penalty-10	45.52	47.48	47.80
Fixed penalty-20	43.72	45.45	44.24
Fixed penalty-40	44.61	48.44	45.57
Price rate-0.1	50.43	49.34	49.08
Price rate-0.3	46.08	47.93	46.87

compared with some other approaches for setting decommitment penalties in the literature. Experimental results show that setting penalties through negotiation achieved higher social welfare than other mechanisms for setting penalties, no matter which strategies are chosen by agents. We are currently using this framework in an actual distributed multi-agent resource allocation application and the preliminary results are encouraging [2].

While we developed partial lookahead agents by approximating the future outside options, we plan to increase the spectrum of lookahead by allowing agents search the game tree. However, as it is often intractable to compute a rational strategy in game trees, it would be desirable to develop some anytime search algorithms concerning computational constraints. We will also explore some other myopic strategies. In the future work, we will also compare negotiating over penalty with other exogenous approaches like penalty as a function of contracting time. In addition, our experiments thus far focused on scenarios with moderate complexity, but we wish to investigate much larger problems where there are more agents. The most practical motivation of bargaining theory is designing agents in environments close to the real world. By allowing agents to negotiate over penalty, agents are given more flexibility to optimize their contracts. While there has been some work applying leveled-commitment contracting in the practical e-commerce, they all use exogenous methods to set decommitment penalties [1,3,9,10]. Designing agents negotiating over penalty e-commerce applications is another focus of our future research.

5. REFERENCES

[1] B. An, F. Douglis, and F. Ye. Heuristics for negotiation schedules in multi-plan optimization. In *Proc. of the Seventh Int. Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 551–558, 2008.

[2] B. An, V. Lesser, D. Irwin, and M. Zink. Automated negotiation with decommitment for dynamic resource allocation in cloud computing. In *Proc. of the 9th International Joint Conference on Autonomous Agents and Multi-Agent Systems*, May 2010.

[3] B. An, V. Lesser, and K. M. Sim. Decommitment in multi-resource negotiation. In *Proc. of the Int. Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 1553–1556, 2008.

[4] M. Andersson and T. Sandholm. Leveled commitment contracts with myopic and strategic agents. *Journal of Economic Dynamics & Control*, 25:615–640, 2001.

[5] M. Branson, F. Douglis, B. Fawcett, Z. Liu, A. Riabov, and F. Ye. CLASP: collaborating, autonomous data stream processing systems. In *Proc. of the ACM/IFIP/USENIX 8th International Middleware Conference*, pages 348–367, 2007.

[6] P. Faratin, C. Sierra, and N. R. Jennings. Negotiation decision functions for autonomous agents. *Int. Journal of Robotics and Autonomous Systems*, 24(3-4):159–182, 1998.

[7] N. Gatti, F. D. Giunta, and S. Marino. Alternating-offers bargaining with one-sided uncertain deadlines: an efficient algorithm. *Artificial Intelligence*, 172(8-9):1119–1157, 2008.

[8] N. Jain, L. Amini, H. Andrade, R. King, Y. Park, P. Selo, and C. Venkatramani. Design, implementation, and evaluation of the linear road bnchmark on the stream processing core. In *Proc. of the 25th ACM SIGMOD International Conference on Management of Data*, pages 431–442, 2006.

[9] T. D. Nguyen and N. R. Jennings. Coordinating multiple concurrent negotiations. In *Proc. of the 3rd Int. Conf. on Autonomous Agents and Multi-Agent Systems*, pages 1064–1071, 2004.

[10] T. D. Nguyen and N. R. Jennings. Managing commitments in multiple concurrent negotiations. *Electronic Commerce Research and Applications*, 4(4):362–376, 2005.

[11] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.

[12] T. Sandholm and V. Lesser. Leveled commitment contracts and strategic breach. *Games and Economic Behavior*, 35(1-2):212–270, 2001.

[13] T. Sandholm, S. Sikka, and S. Norden. Algorithms for optimizing leveled commitment contracts. In *Proc.of the 16th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 535–541, 1999.

[14] T. Sandholm and Y. Zhou. Surplus equivalence of leveled commitment contracts. *Artificial Intelligence*, 142(2):239–264, 2002.