Monopolizing Markets by Exploiting Trust

Sandip Sen
University of Tulsa
sandip@utulsa.edu

Dipyaman Banerjee
University of Tulsa
dipu@utulsa.edu

ABSTRACT

We investigate trust-based relationships in electronic supply chains where trust is a measure of consistency in meeting negotiated contract deadlines. We consider scenarios where contractors assign contracts to contractees who are significantly more successful in meeting deadlines. The task deadlines are drawn from a known distribution. We present a probabilistic analysis that enables contractees to strategically bid on only certain tasks to earn the trust of their contractors. Once a contractee achieves a high level of trust, it can then exploit that trust to increasingly corner a larger portion of the market share of all tasks. We present a trust exploitation scheme that monopolizes the market against greedy contractees who bid on all announced tasks. We also show that such market monopolization is not possible in the presence of a trusted, but non-exploiting, contractee. The exploiter can, however, effectively "starve" the non-exploiting trusted agent.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Coherence and coordination, Multiagent systems, Intelligent agents*

General Terms

Algorithm, Performance, Experimentation

Keywords

trust, monopoly, markets, supply chains

1. INTRODUCTION

Supply chains form distributed, dynamic networks of agile, adaptive entities. As a result, supply chain management where rational agents represent interests of individual entities and organizations have been an area of active research [16, 21, 23]. From the multiagent systems perspective, a supply chain is often modeled using a decentralized network of agents where each agent can do a part of a task. Typical approaches try to optimize the profitability of entities with emphasis on pricing and scheduling [7].

Researchers in business and management science, however, have recognized that a key component of decision-making in real-world supply chains is the consideration of trust between the contracting organizations [12, 24]. Though the use of information technology and computer systems to facilitate trust-building between organizations have been proposed in mid-80s [14], little work has been done in this area.

We believe an agent based supply chain can effectively capture the needs and constraints faced by real-life B2B systems, and hence, can become an integral part of fielded applications only when such higher level issues like trust are adequately addressed. In our previous work we focused on the aspect of developing stable, long term, trusted partnerships that allow organizations to have more accurate predictions of the time and cost requirements to meet market demands [18]. This work measured trust as the likelihood of completing contracted tasks by their deadlines. Agents that were sufficiently more trusted were always awarded contracts over others. Task deadlines were drawn from a known distribution, and strategic agents bid only on tasks for which they were sufficiently more confident of meeting deadlines. Over time, such an agent would be recognized as more trustworthy by the contractor(s) and would be preferred over other agents who bid for all tasks. We refer to such trust-building agents as *trust-building (TB)* agents.

In this paper, we investigate how a trusted contractee agent can exploit the trust of its contractor(s) to eliminate the competition and monopolize market share. From a self-interested viewpoint, trust is like any other resource that can be leveraged to gain competitive advantage. We propose that once an agent has build sufficient trust, the next logical step would be to exploit that to incrementally take over the entire market. This process mirrors the development of a monopolist in a marketplace though practical scenarios are more complex and involve aspects beyond trust. We refer to agents who exploit their trust as *trust exploiting (TE)* agents.

We further evaluate the performance of such a monopolistic agent in the presence of a non-exploiting TB agent who has no intention to monopolize. The presence of the TB agent presents a major impediment to the TE agent's goal of monopolizing the market directly. We motivate, implement, and evaluate two variations of TE agent that tries to corner as much of the market share as possible even if

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.
Copyright 2006 ACM 1-59593-303-4/06/0005 ...$5.00.

monopolization is precluded.

We use a contracting framework for connecting the suppliers to the manufacturers: manufacturers announce contracts for tasks with given deadlines; suppliers bid on these tasks; and the contract is allocated to one of the bidders based on its past performance history. We use a simplified contracting model where contractors prefer to assign tasks to contractees who have significantly higher success rates in meeting deadlines. We assume that the trust preferences of the contractor, the task deadline distribution and the performance distribution of the contractees are known. We then develop a precise bidding strategy for trust-building contractees. The motivation is to bid only on those tasks for which they have a high likelihood of meeting deadlines. However, not bidding on tasks also reduces the success rate of completing tasks. We provide a probabilistic analysis to handle this tradeoff. We then augmented this strategy to incorporate strategic shifts in the bidding range, once trust has been established, to monopolize the market. Further, we investigate the different scenarios which may occur when another trusted agent is present in the market. Finally we provide experimental verification of our formal analysis.

We emphasize that we are not pursuing, in this paper, a game-theoretic approach to equilibrium analysis in the context of trust-based contracting. While such an analysis is critical and necessary under the assumption of all bidders being rational, we are more interested in approximately modeling real world situations. In typical market environments, participants have a varying degree of knowledge, longevity, resources, strategic reasoning capability, etc. In particular, we model market situations where almost of all of the contractors use a myopic or greedy approach to bidding for tasks. This myopic behavior can stem from short market lifetime, financial exigencies, resource constraints, lack of knowledge or reasoning capability, etc. The research question we are pursuing is how should a rational contractor, with long-term market presence, approximate knowledge about market conditions, and no limiting resource constraint, etc. choose its bidding strategy to maximize its long-term utility. While we follow intuitive guidelines while developing and exploiting trust, the technical contribution of this work lies in formally developing precise bidding strategies with accurate performance predictions followed by experimental verification. Whereas trust and reputation in agent based systems is a very active research area in multiagent systems [5, 6, 10, 15, 20, 22, 25], to date there exists very little work on formal derivation of strategies to build or exploit trust.

2. CONTRACTING MODEL

We assume a trust-based model for task allocations by a contractor to bidders in a marketplace.

2.1 Task and performance distributions

We assume a task distribution $\mathcal{T}$ from which the deadline of different tasks are drawn. $\mathcal{P}_j$, defined over the closed interval $[l, h]$, represents the distribution from which the actual time taken by supplier j to process tasks is drawn. For continuous distributions, this means $\int_l^h \mathcal{P}_j(x)dx = 1$ and $\int_l^h T(y)dy = 1$. The corresponding discrete distributions, where T and $\mathcal{P}_j$ represent probability mass functions, are $\sum_{x=h}^{l} \mathcal{P}_j(x) = 1$ and $\sum_{y=h}^{l} T(y) = 1$ respectively.

2.2 Use of trust in selecting contractees

Each contractor maintains a trust rating, t_c for each possible contractee, c, which is the proportion of time that contractee could meet an assigned task deadline. For any two bidders, $\underline{b}, \overline{b} \in B_i$, we say $\overline{b}$ γ-dominates $\underline{b}$ (equivalently, $\underline{b}$ is γ-dominated by $\overline{b}$) if

$$\gamma * t_{\underline{b}} < t_{\overline{b}}, \tag{1}$$

where $\gamma > 1$ is a trust constant, chosen by the contractor. The basic motivation or incentive for the contractor to use a trust measure for selecting suppliers is to ensure reliability. If a supplier fails to deliver the required materials or products by the negotiated deadline, the manufacturer may not be able to meet its own production schedules and delivery deadlines and hence may face stiff penalties.

Each task announcement, T_i, contains an associated deadline $d(T_i)$ drawn from the distribution $\mathcal{T}$. Each supplier decides to bid or not bid on an announced task. As we are concentrating only on the trust evaluation of the bidders in awarding tasks, a bid from a bidder j can just be a default message signaling js interest in the task. The contractor of a task assigns it to a contractee randomly chosen from the set of non-γ-dominated bidders in the set of bidders for the task, B_i.

Note that this is a deliberate simplification of the real-world contracting process. In particular, we have left out the consideration of the price in the bid while awarding contracts! We believe that real-world stable supply chains incorporate trust considerations as a key metric for awarding contracts. We acknowledge that even in such situations, trust is one of several key parameters, including price, that determines contract awards. In this study, however, we wanted to focus exclusively on trust considerations as there has been very little work in agent based systems on the use of trust in supply chain contracting. Another way of situating this work is to assume that the contractor is choosing, based on trust, between all bidders with the same price. The issues we are studying can also be analyzed effectively with only one manufacturer and several suppliers and contractees, and hence we do not use a full supply chain. The techniques presented, however, readily transfers to more elaborate supply chains.

2.3 Goal of contractee

Let there be N contractees. The goal of a contractee j is to maximize its success in procuring and delivering contracts. The objective criterion to be maximized, the *success rate*, is[1]:

$$\int_{bl}^{bh} r(j, y)T(y) \left(\int_l^y \mathcal{P}_j(x)dx \right) dy, \tag{2}$$

where $r(j, y)$ is the probability that a task of length y is assigned to supplier j by the contractor (depends on the set of bidders for the task and their respective trust values), the integration within the parenthesis represents the likelihood of meeting the deadline y and bl, bh are respectively the minimum and maximum deadlines for tasks on which j will bid. The *success rate* represents the expected number of assigned tasks successfully delivered by their deadline.

[1]Most of the analysis presented in the following uses notations corresponding to continuous distributions. The equivalent expressions for discrete distributions can be similarly derived.

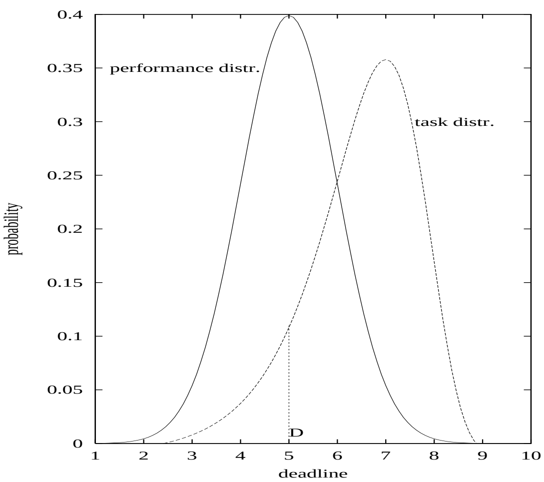

Figure 1: Example continuous task deadline and performance distributions.

3. BUILDING TRUST

We assume that the contracting model is common knowledge between the contractor and the suppliers and that each supplier j is cognizant of the task deadline distribution $\mathcal{T}$, its performance distribution $\mathcal{P}_j$ and the trust constant γ. We now present two types of bidding strategies for a contractee j in response to an announcement for a task t_i with deadline $d(T_i)$:

Greedy (G): This strategy bids for a task whenever there is a non-zero probability of meeting the deadline, i.e., $\int_l^{d(T_i)} \mathcal{P}_j(x)dx > 0$.

Trust-building (TB): This strategy focuses on building trust, which is measured as the likelihood of meeting the deadline of an assigned task. This strategy is elaborated in the rest of this section.

The number of contractees using the G and TB strategy are known and given by N_G and N_T respectively. Let $N = N_G + N_T$.

The key consideration for a TB agent is the choice of the minimum deadline threshold, D, for bidding such that it is viewed by the contractor to be sufficiently more trustworthy, in the sense of γ-domination, than a G contractee. We assume that performance distribution $\mathcal{P}$ is the same for all agents. The Figure 1 presents a typical situation with task and performance probability distributions and the deadline D below which TB agents will not bid.

The average expected success likelihood or *trustworthiness* of an agent j who wins, i.e., is awarded, all tasks in the region $[bl, bh]$ is given by

$$\bar{P}(bl, bh) = \frac{1}{\mathcal{T}_{bl,bh}} \int_{bl}^{bh} T(y) \left(\int_l^y \mathcal{P}_j(x)dx \right) dy, \qquad (3)$$

where $\mathcal{T}_{x,y} = \int_x^y T(z)dz$, is the cumulative probability of tasks arriving with deadlines in the region $[x, y]$. The above expression is obtained using Equation 2 with $r(j, y) = 1$, i.e., agent j wins the contract with certainty. If an agent wins only a fraction f of tasks in that region, the corresponding average success likelihood is $f\bar{P}(bl, bh)$.

To facilitate the presentation of the analysis of strategic bidding for tasks, we consider the steady state case, where TB agents always win the contract when they bid (as they have been recognized to be more trustworthy than greedy contractees), and at other times one of the G agents win the contract. First we note that the TB agents should be more inclined to bid on tasks at the higher end of the task deadline distribution as they are more likely to be able to meet the corresponding deadlines. Hence, the upper limit of the range of task deadlines a TB agent will bid, bh_{TB} is equal to h, the maximum task deadline. Let the minimum deadline for which a TB agent bids at steady state be D_{ss}. Therefore, at steady state, the TB agents will win all the tasks for deadlines in the range $[D_{ss}, h]$ and G agents will win all contracts in the region $[l, D_{ss})$. From Equation 1, the following inequality holds if the TB agents were to γ-dominate the greedy agents at steady state:

$$\bar{P}(D_{ss}, h) > \gamma\bar{P}(l, D_{ss}). \qquad (4)$$

Equations 3 and 4 can be used to calculate D_{ss}.

The assumption implicit in Equation 4, that the TB agents win whenever they bid, is not valid initially, when they are yet to be recognized as γ-dominant. Hence D_{ss} is not the appropriate choice for the initial minimum task deadline to bid for, D_I, by TB agents. To calculate D_I, we assume that tasks are initially assigned randomly between all bidders. So, while all the tasks with deadlines in the range $[l, D_I)$ will be assigned to G contractees, tasks in the region $[D_I, h]$ will be assigned to G versus TB bidders in the ratio of N_G to N_{TB}. The choice of D_I should be such that it allows a TB agent to have at least γ times higher trust than G agents when tasks are being assigned randomly between bidders:

$$\frac{N_{TB}}{N}\bar{P}(D_I, h) > \gamma \left(\frac{\mathcal{T}_{l,D_I}\bar{P}(l, D_I) + \mathcal{T}_{D_I,h}\frac{N_G}{N}\bar{P}(D_I, h)}{\mathcal{T}_{l,h}} \right). \qquad (5)$$

The left hand side of the inequality represents the proportion of tasks expected to be successfully delivered by TB agents when tasks are randomly assigned between all bidders and the TB agents bid only in the interval $[D_I, h]$. The term within the parenthesis on the RHS of the inequality denotes the proportion of tasks successfully delivered by G agents in this period. The terms $\mathcal{T}_{l,D_I}$ and $\mathcal{T}_{D_I,h}$ are used to normalize the trustworthiness in the regions $[l, D_I]$ and $[D_I, h]$ respectively. To calculate D_I we simplify Equation 5

$$\bar{P}(D_I, h) > \frac{N\gamma\mathcal{T}_{l,D_I}}{N_{TB} - \gamma N_G \mathcal{T}_{D_I,h}}\bar{P}(l, D_I), \qquad (6)$$

where the simplification uses the fact that $\mathcal{T}_{l,h} = 1$ and that

$$N_{TB} > \gamma N_G \mathcal{T}_{D_I,h}. \qquad (7)$$

The inequality in Equation 6 can be satisfied for a range of D_I values. The TB agent uses the minimum value in the range which also satisfies the inequality in Equation 7.

In this paper we use equal number of TB and G agents and hence Equation 6 simplifies to

$$\bar{P}(D_I, h) > \frac{2\gamma\mathcal{T}_{l,D_I}}{1 - \gamma\mathcal{T}_{D_I,h}}\bar{P}(l, D_I), \qquad (8)$$

and Equation 7 simplifies to $\mathcal{T}_{D_I,h} < \frac{1}{\gamma}$.

The goal of the contractee is to maximize its success rate as described in Section 2.3. Hence, a contractee must evaluate whether it will have a better success rate, in the long run, if it decides to be a TB or a G agent. If all agents were of type G, each contractee will be awarded tasks with equal probability, i.e., $\forall j, r(j,y) = \frac{1}{N}$. With N_{TB} TB agents, the probability of such an agent winning a contract when it bids, and after it becomes γ-dominant, is $\frac{1}{N_{TB}}$. Hence, in addition to satisfying the inequality in 6 and 7, the D_I value chosen must satisfy the following inequality

$$\int_{D_I}^{h} T(y) \left(\int_{l}^{y} \mathcal{P}_j(x)dx \right) dy > \frac{N_{TB}}{N} \int_{l}^{h} T(y) \left(\int_{l}^{y} \mathcal{P}_j(x)dx \right) dy. \quad (9)$$

Given T and $\mathcal{P}$ distributions, we designate by γ_{max} the maximal γ value for which there exists a D_I that satisfies all these conditions.

3.1 Bidding strategy of Trust Exploiting Agents

We now present the strategy of the *trust exploiting (TE)* agent. Whereas the TB agent picks a D_I and bids only for tasks above it, TE will progressively bid over larger parts of the task distribution, i.e., for smaller task deadlines, until it monopolizes the market. The key to understanding how such a monopoly can be achieved is to recognize that as a TE agent shifts its deadline to bid, D_I, to the left, i.e., for tasks with shorter deadlines, its trust level drops but so does that of greedy agents who now win only on even more riskier tasks. The following discussion assumes that the TE agents first enter a market with G agents and have achieved γ-domination by bidding for tasks in the range $[D_I, h]$. At this stage, the trust ratings of TE and G agents are $t_{TE} = \bar{P}(D_I, h)$ and $t_G = \bar{P}(l, D_I)$ respectively. Actually, t_{TE} should be much more than the value required to γ-dominate t_G because D_I was calculated assuming half the tasks in the safer task deadline range $[D_I, h]$ go to G agents, which is not true after TE agents achieve γ-domination. Hence it is possible for TE agents to use a lower task deadline, D_{new} to bid on and still maintain γ-domination over the G agents, i.e., $D_{new} < D_I$ and $\bar{P}(D_{new}, h) > \gamma\bar{P}(l, D_{new})$. After sticking with this new deadline for sometime to stabilize new trust values, the TE agents can again expand its bidding range, and continue to do so until it bids over the entire range, effectively monopolizing all task contracts! This process of expanding the bid range after achieving γ-domination over the G agents is described in Algorithm 1.

Algorithm 1 Selecting increasingly larger bidding range to achieve market monopoly.

while $D_I > l$ **do**
  Find D_{new} such that $\bar{P}(D_{new}, h) > \gamma\bar{P}(l, D_{new})$.
  $D_I \leftarrow D_{new}$
  Bid in the range $[D_I, h]$ for a sufficient period to settle on new trust values
end while

The situation is more challenging for the TE agent if it enters the market when a TB agent is already γ-dominating G agents. The first task for the TE agent to gain a foothold is to γ-dominate the TB agent. The corresponding bidding deadline for tasks, D_{TE} can be calculated from the inequality, $\bar{P}(D_{TE}, h) > \gamma\bar{P}(D_I, D_{TE})$ (we assume here that newcomers to the marketplace are given some tasks to evaluate their trustworthiness). Once it has achieved that domination, there are two possibilities for further exploitation:

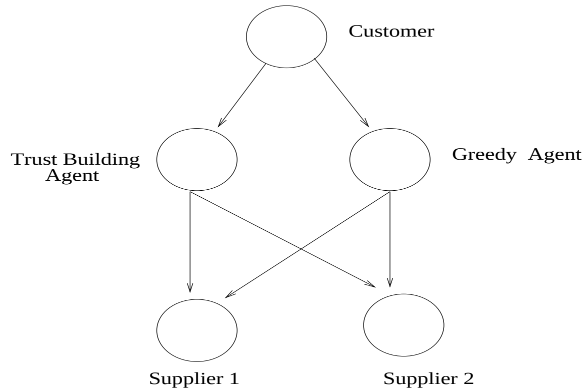

Figure 2: A simple supply chain.

TE1: TE1 expands to bid for all tasks. As a result it will lose γ-domination over TB. When this happens, it switches back to original D_{TE}. Thus the TE1 agent switches its low-end bid deadline between D_{TE} and l.

TE2: TE2 expands to coincide its bidding deadline with that of the TB agent. It may eventually lose γ-domination over TB agent, depending on the task deadline distribution. When this happens, it switches back to original D_{TE}. Thus the TE2 agent switches its low-end bid deadline between D_{TE} and D_I.

A TE2 agent can, depending on the task distribution, "starve" the TB agent out of the market. If the TB agent then leaves the market, we have the scenario of a TE agent against G agents, which will eventually lead to the TE agent monopolizing the market. In other situations, a TE2 agent may outperform a TE1 agent when facing the same TB and G competitors, but not each other.

4. EXPERIMENTAL FRAMEWORK

We use a simple two-level supply chain consisting of a single contractor at the top level, two to three contractees in the second level, and two suppliers under each such contractee at the third level. A representative instance is presented in Figure 2. We ran our experiments in two discrete phases. In the first phase we compared the behavior of two agents bidding, one TB and one TE agent when they play separately against the same G agent, for a contract offered by the customer, the contractor. In the next phase we allow all three agents to operate simultaneously and compared their performances. The deadlines for the contracts are generated randomly from a discretized triangular distribution. We used triangular distribution as an approximation of likely real-life distributions because of easy calculations of the cumulative probabilities. There are two suppliers under each of these two contractees, and the contractee agents must procure supplies from these downstream suppliers before it can process an assigned task.

Since the total time taken by the contractee is the sum of the time taken by its supplier to produce the necessary supplies plus the time it takes to use these supplies to process the tasks, the corresponding task distributions are chosen so that the resultant performance distribution range matches that of its task distribution. We have used identical triangular distributions for the contractee and the suppliers

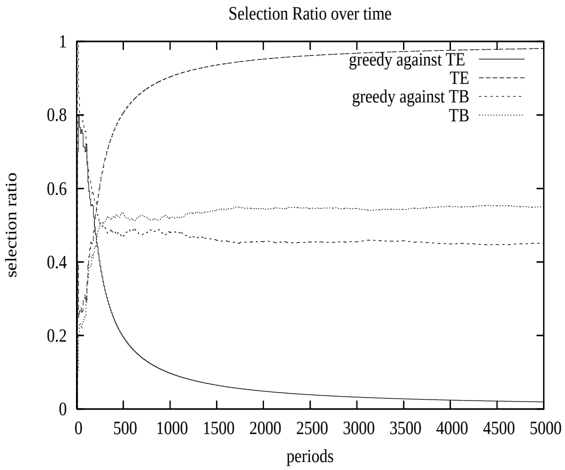

Figure 3: Selection Ratio ($T_p = 8$, $\gamma = 1.29$).

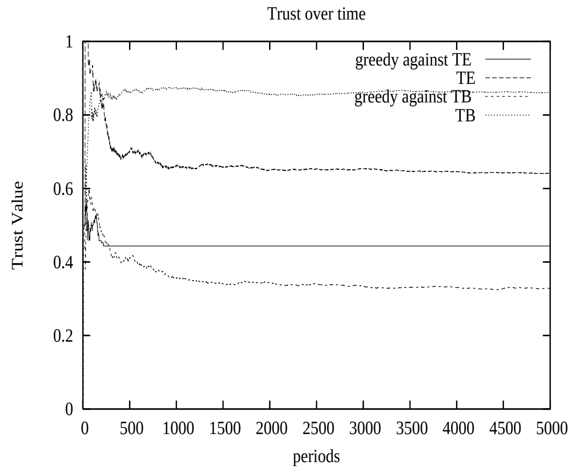

Figure 4: Trust over time ($T_p = 8$, $\gamma = 1.29$).

defined over the range [0,6] with the highest point of the distribution at the midpoint of the range. The resultant performance distribution of the contractee then ranges over [0,12].

In our experiments, we iteratively generated new task deadlines from the task distribution. In a particular iteration, the TB agent decides to bid if the deadline is greater than D_I. The G agent bids on all tasks and TE agent bids if the deadline is greater than its current D_I value.

In the first 100 iterations the customer randomly selects between the agents with bids to estimate the trustworthiness of the contractees. Thereafter a contractee for a task is chosen using the selection criteria specified in Section 2.2. We also assume that the customer selects any new bidder entering the marketplace with a small probability, ϵ, to estimate its performance. After the new agent is selected 50 times using this ϵ-exploration, the customer reverts back to its previous selection strategy using the computed trust values based on these exploration phase. This exploration phase is used only when the TE agent enters the market after TB agent is γ-dominating the G agents.

When a contractee is awarded a task, it generates a task announcement to procure necessary supplies from the suppliers. The time taken by the suppliers followed by that of the selected contractee are then generated from their performance distributions using the standard inverse transform method and are added to calculate the contract fulfillment or delivery time. The agent succeeds if the delivery time is less than or equal to the contract deadline, and fails otherwise.

5. EXPERIMENTAL RESULTS

In our first experiment, we compare the selection ratio and the trust value of the greedy and trust-based contractees over 5000 iterations, where the selection ratio is the proportion of time the agent is selected. The experiment is repeated on three different task distributions: T_p= 4, 6 and 8 and for $\gamma = 1.29$. The results from the $T_p = 8$ scenario is shown in Figure 3 (plots for other values used for T_p are similar and are omitted due to space constraints).

From Figure 3 we observe that though TB agents are selected more often than the greedy agents, it can never get

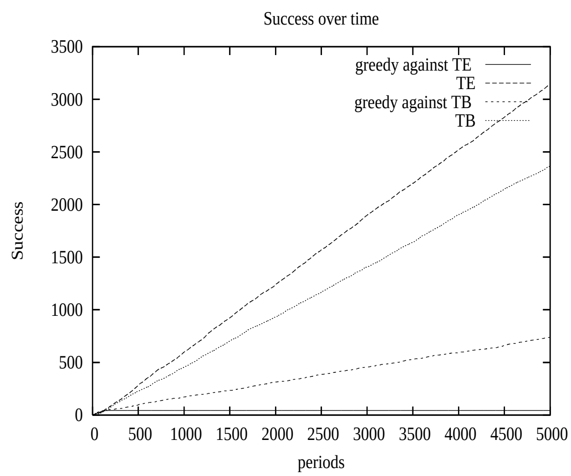

Figure 5: Success over time ($T_p = 8$, $\gamma = 1.29$).

more than 60% of the task as it never bids on a task whose deadline is less than D_I. On the other hand, the selection ratio of the TE agent asymptotically reaches 1 as it starts taking over the market once it has gained sufficient trust. To ensure that, it observes the market continuously and after winning successive 20 bids it sets its new D_I at 2 forcing the G agent to operate only on the high-risk tasks having deadline from 0 to 1. This causes the greedy agent to fail more often and thus loose its trust even more. The trust value of TE agent also decreases and is less than that of TB as shown in Figure 4. But the rate of decrease of trust for the TE agent is much slower than that of the G agent which enables the former to maintain γ-domination. As a result it continues winning all the tasks it bids, and after 20 wins it shifts D_I to 0 and takes over the entire market. From this point onward the G agent does not win any bid, and as a result its trust value remains constant (Figure 4).

We observe from Figure 5 that the TE agent receives more success compared to the TB agent when they play separately against the greedy agent. Also note that, the G agent wins no contracts once TE takes over the market.Due to sampling biases, a sequence of tasks with short deadlines may be generated which the TE agent may fail to meet, thereby loosing γ-domination over G agents. To counter this problem, the monopolist TE agent monitors the situation and whenever it looses enough bids to loose γ-domination over the greedy agent, it reverts back to its initial D_I value to rebuild the lost trust. Once it gains back the trust, it uses Algorithm 1 to re-monopolize the market.

In our next set of experiments, only the TB and G contractee agents are in the system for the first 400 iterations. Once the TB contractee gain γ-domination over G, a TE contractee is introduced. Here again we assume that the customer uses ϵ-exploration to evaluate the new entrant. We used two different versions of TE agent: TE_1 and TE_2. Both of these agents choose a higher task deadline threshold D_{TE} to achieve γ-dominance over not only the G agent, but also the TB agent. Once it gains the required trust, TE_1 tries to monopolize the market by setting its D_I to 0 and oscillates back and forth to maintain its trust domination. From the results (see Figure 6) we observe that the success of TE_1 depends on the nature of the task distribution; it does well for the distribution with $T_p = 4$, but for task distributions with $T_p = 6$ and higher, TB outperforms TE_1. This is because, in order to γ-dominate the TB agent, the TE_1 bids on a very small range of the task spectrum for such distributions when tasks with longer deadline is common. In such situations, it is not able to win sufficient number of tasks to outperform the TB agent.

TE_2 behaves less greedily than TE_1: after it gains γ-domination over the TB agent, instead of monopolizing the market it sets its D_I at the same position as the TB agent. The TB agent, therefore, loses some of its market share to the TE_2 agent. From Figures 6 and 7 we observe that the success rates of TE_2 are better than that of TE_1. The TE_2 agent, however, loses out to the TB agent for task distributions with $T_p = 8$ and above. In such situations, eventually TE_2 loses γ-domination over TB and moves its bidding deadline to D_{TE}, and after regaining trust switches it back to D_I of TB agent, and the cycle repeats. TE_2 also retains its γ-domination over the greedy agent through this entire process. But this oscillation does not happen for all task distributions, e.g., for $T_p = 4$ and 6, TE_2 never needs loses γ-domination over TB even when it continues to play in the range $[D_I, h]$. This is because, when TE_2 becomes γ-dominant over TB while playing in the range $[D_{TE}, h]$, the TB agent is only awarded tasks with deadline in the range $[D_I, D_{TE})$. For $T_p = 4$ and 6, not only is $\bar{P}(D_{TE}, h) > \gamma \bar{P}(D_I, D_{TE})$ required for establishing γ-domination, but $\bar{P}(D_I, h) > \gamma \bar{P}(D_I, D_{TE})$, i.e., the TE_2 agent continues to maintain γ-domination even after bidding over the range on which the TB agent was bidding when the TE agent enters the market. As a result we observe that in these situations (see the middle and left plots in Figure 7), the TB agent receives no contracts after some time, i.e., it is effectively "starved". Note that the TE_2 agent cannot monopolize all the contracts, and hence the G agent continues to receive the more riskier contracts on which none of the trust-based agents are bidding. Ironically, the presence of the TB agent benefits the G agent: TB starves, but prevents TE from starving G as well!

6. RELATED WORK

More recently the problem of supply chain management has drawn the attention of multiagent systems (MAS) researchers. In MAS, a supply chain is conceptualized as a group of collaborative autonomous software agents [11]. It is argued that managers can better coordinate and schedule processes by distributing the organization-wide business management system into autonomous problem solving agents. This approach is more appropriate when the participating enterprises are geographically distributed and manually controlling $B2B$ trading is not possible. Tackling coordination in supply chains using partial constraint satisfaction problems by mediating agents is investigated by [2].

Swaminathan *et al.* has provided a framework for efficient supply chain formation [21]. MAS researchers have paid attention to the emergence of the optimal supply chain configuration. Walsh *et al.* has shown the optimal dynamic task allocation in a supply chain using combinatorial auctions [23]. Given a task, composed of a group of subtasks, they provide a mechanism for dynamic formation of a supply chain to produce maximum profit. Collins and Gini have provided a testbed, MAGNET, for multiagent contracting for supply chain formation using multiple criteria [7]. A few years back, the Trading Agent Competition (TAC) [16] introduced a realistic supply chain framework where the challenges include the design of strategies for effectively handling factory schedules, efficiently contracting with the suppliers, competitively and profitably bidding to the customers, reducing the inventory cost and penalty for late delivery or order cancellation.

In recent research, trust is established as a key factor to build profitable and long term B2B and B2C collaborations in the Industry [12]. Trust plays an even more critical and important role in the domain of Electronic commerce [3]. Several trust management systems have been proposed to handle the development of trust and its impact in such systems [9].

Use of trust as a basis for interaction strategies has been widely used in multiagent systems. Marsh was one of the first to attempt a computational model of trust [13]. Sen *et al.* have developed a comprehensive reciprocity-based reputation framework for learning to trust like-minded agents and outperform exploitative agents in fully distributed environments like P2P networks [17, 19, 1]. Castelfranchi and colleagues have argued for the necessity of trust in social interactions between agents with complex mental attitudes [5]. Brainov and Sandholhm [4] have shown that trust based contracting can significantly increase market efficiency measured by social welfare, trade volume and agent utilities. Yu and Singh [26] have proposed a mechanism for combining reputation from multiple sources to obtain trust ratings. Glass and Grosz present the SPIRE framework [8], where agents who fulfill social obligations are treated favorably by other members of the society. Sen *et al.* have developed a comprehensive reciprocity-based reputation framework for using trust in

7. CONCLUSIONS AND FUTURE WORK

In this paper, we have argued for the use of trust models to award contracts in the context of supply chains. We work with a particular trust-based contracting framework where contractees with significantly higher historical success rates in meeting contract deadlines are preferentially selected over less "punctual" agents. Our proposed trust-building agents bid less often for tasks, but win and suc-

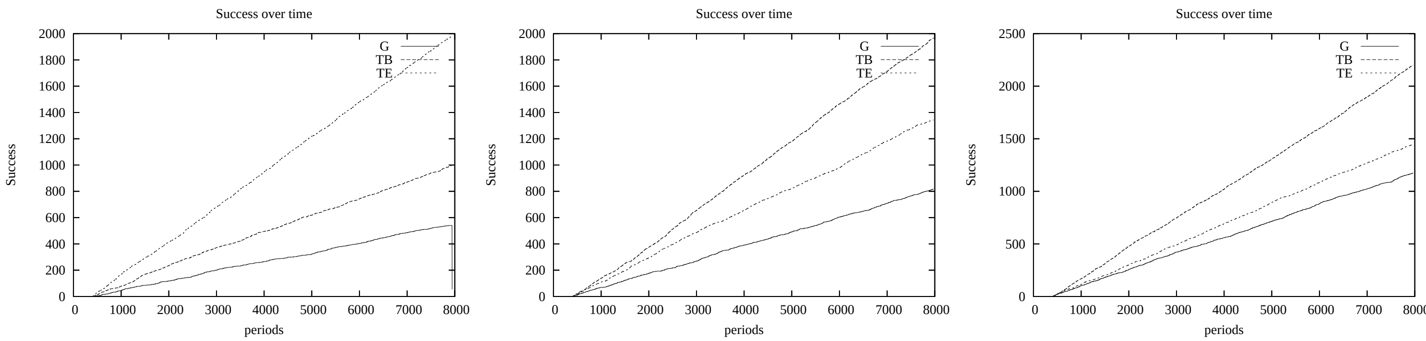

Figure 6: Success of G, TB, and TE_1 agents over time for different task distributions: $T_p = 4$ (left), 6 (middle), 8 (right); $\gamma = 1.15$.

cessfully meet the deadline of tasks they bid for more often than "greedy" agents who bid for all tasks.

Given task deadline and performance, i.e., task processing distributions, we use a probabilistic analysis to analytically derive the task deadline threshold below which a trust-building contractee will not bid for a task. We expand on this analysis to develop exploitative agents who incrementally take over the entire market from greedy agents. We also propose variations of exploitative agents when competing against both greedy and non-exploiting trust-building agents. We provide experimental verification on a small supply chain to demonstrate the competitive advantage of these trust-exploiting strategies.

While the basic motivations and outlines for trust-exploitation strategies presented in this paper are intuitive, the contribution of our paper lie in formally specifying the trust-exploitation problem and deriving precise strategies for exploiting trust. We present implementable strategies with predictable performance given task deadline and performance distributions and the trust-based contract allocation criterion. We first derived a trust-based strategy that trades off short-term loss (not bidding on certain tasks) for long-term profits (higher success rates after achieving γ-domination). Following this we introduced variants of aggressive trust-exploitation schemes that not only establishes trust as the previous scheme, but then leverages it strategically to monopolize the market. We show that though such market monopolization is not possible if non-exploitative trust-based agents are present, the latter can be effectively dominated under certain task deadline distributions.

In this paper, the trust constant, γ, and the contract allocation procedure are assumed to be common knowledge. The task distribution is assumed to be constant and known *a priori*. Similarly, the performance distribution is fixed and is assumed to be identical for all contractees. Some or all of those assumptions may be violated in a real environment. For example, the contractor may not be consistent in either maintaining or enforcing the trust-based contract award scheme, task distribution may vary over time, agents may be heterogeneous in their capabilities, etc. We plan to relax some of this approaches in future work. In particular, we plan to explore various adaptive mechanisms to cope with dynamic, uncertain environments.

Acknowledgments This work has been supported in part by an NSF award IIS-0209208.

8. REFERENCES

[1] D. Banerjee, S. Saha, P. Dasgupta, and S. Sen. Reciprocal resource sharing in p2p environments. In *Proceedings of the Fourth Intenational Joint Conference on Autonomous Agents and Multiagent Systems*, pages 853–859, New York, NY, 2005. ACM Pres.

[2] J. C. Beck and M. S. Fox. Supply chain coordination via mediated constraint relaxation. In *Proc. of the 1st Canadian Workshop on Distributed Artificial Intelligence*, 1994.

[3] A. Bhattacherjee. Individual trust in online firms: Scale development and initial test. *Journal of Management Information Systems*, 19(1):211–241, 2002.

[4] S. Brainov and T. Sandholm. Contracting with uncertain level of trus. *Computational Intelligence*, 18(4):501–514, November 2002.

[5] C. Castelfranchi and R. Falcone. Principles of trust for MAS: Cognitive autonomy, social importance, and quantification. In *Proceedings of the Third International Conference on Multiagent Systems*, pages 72–79, Los Alamitos, CA, 1998. IEEE Computer Society.

[6] C. Castelfranchi and Y.-H. Tan. The role of trust and deception in virtual societies. In *Proceedings of the 34th Hawaii International Conference on System Sciences*, Los Alamitos, CA, 2001. IEEE Computer

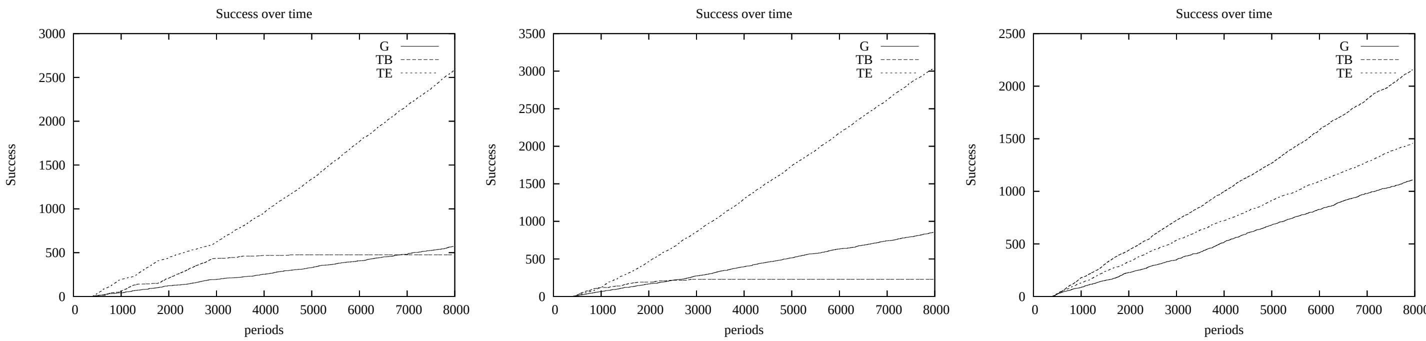

Figure 7: Success of G, TB, and TE_2 agents over time for different task distributions: $T_p = 4$ (left), 6 (middle), 8 (right); $\gamma = 1.15$.

Society.

[7] J. Collins, W. Ketter, and M. Gini. A multi-agent negotiation testbed for contracting tasks with temporal and precedence constraints. *International Journal of Electronic Commerce*, 7(1):35–57, 2002.

[8] A. Glass and B. Grosz. Socially conscious decision-making. *Autonomous Agents and Multi-Agent Systems*, 6(3):317–339, May 2003.

[9] T. Grandison and M. Sloman. A survey of trust in internet applications. *IEEE Communications Surveys and Tutorials*, 3(4), 2000.

[10] N. Griffiths and M. Luck. Coalition formation through motivation and trust. In *Proceedings of the Second Intenational Joint Conference on Autonomous Agents and Multiagent Systems*, pages 17–24, New York, NY, 2003. ACM Pres.

[11] N. R. Jennings, P. Faratin, T. J. Norman, P. O'Brien, and B. Odgers. Autonomous agents for business process management. *International Journal of Applied Artificial Intelligence*, 14(2):145–189, 2000.

[12] Y. L. Lederer-Antonucci, P. S. Greenberg, M. zur Muehlen, and G. Ralph. Establishing trust in a business-to-business collaboration: Results from an international simulation. In *Proc. of the IRMA 2003 Conference*, pages 922–924, 2003.

[13] S. Marsh. *Formalising Trust as a Computational Concept*. PhD thesis, University of Stirling, April 1994.

[14] R. Miles and C. Snow. Organizations: New concepts for new firms. *California Management Review*, 28(3):62–73, 1986.

[15] M. Schillo, P. Funk, and M. Rovatsos. Using trust for detecting deceiptful agents in artificial societies. *Applied Artificial Intelligence*, 14:825–848, 2000.

[16] Tac, supply chain management. url: http://www.sics.se/tac/.

[17] S. Sen. Believing others: Pros and cons. *Artificial Intelligence*, 142(2):179–203, 2002.

[18] S. Sen, S. Saha, and D. Banerjee. Trust-based contracting in supply chains. In *Proceedings of the Fourth Intenational Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1251–1252, New York, NY, 2005. ACM Pres.

[19] S. Sen, S. Saha, and P. S. Dutta. Helping based on future expectations. In *Proceedings of the Second Intenational Joint Conference on Autonomous Agents and Multiagent Systems*, pages 289–296, New York, NY, 2003. ACM Pres.

[20] S. Sen and N. Sajja. Robustness of reputation-based trust: Boolean case. In *Proceedings of the First Intenational Joint Conference on Autonomous Agents and Multiagent Systems*, pages 288–293, New York, NY, 2002. ACM Pres.

[21] J. M. Swaminathan, S. F. Smith, and N. M. Sadeh. Modeling supply chain dynamics: A multiagent approach. *Decision Sciences*, 29(3):607–632, 1998.

[22] Y. Tan and W. Thoen. An outline of a trust model for electronic commerce. *Applied Artificial Intelligence*, 114(8):849–862, 2000.

[23] W. E. Walsh and M. P. Wellman. Decentralized supply chain formation: A market protocol and competitive equilibrium analysis. *Journal of Artificial Intelligence Research*, 19:513–567, 2003.

[24] B. Welty and I. Becerra-Fernandez. Managing trust and commitment in collaborative supply chain relationships. *Communications of the ACM*, 44(6):67–73, 2001.

[25] P. Yolum and M. P. Singh. Engineering self-organizing referral networks for trustworthy service selection. *IEEE Transactions on System, Man, and Cybernetics*, 2005.

[26] B. Yu and M. P. Singh. Distributed reputation management for electronic commerce. *Compuational Intelligence*, 18(4):535–549, 2002.