Cooperation based Dynamic Team Formation in Multi-Agent Auctions

Charles E. Pippin[a*] and Henrik Christensen[b]

[a]Georgia Tech Research Institute, Georgia Institute of Technology, Atlanta, GA, USA;
[b]Center for Robotics and Intelligent Machines, Georgia Institute of Technology, Atlanta, GA

ABSTRACT

Auction based methods are often used to perform distributed task allocation on multi-agent teams. Many existing approaches to auctions assume fully cooperative team members. On in-situ and dynamically formed teams, reciprocal collaboration may not always be a valid assumption.

This paper presents an approach for dynamically selecting auction partners based on observed team member performance and shared reputation. In addition, we present the use of a shared reputation authority mechanism. Finally, experiments are performed in simulation on multiple UAV platforms to highlight situations in which it is better to enforce cooperation in auctions using this approach.

Keywords: multi agent systems, market-economy, decision theory

1. INTRODUCTION

Auction based methods are often used to perform distributed task allocation on multi-agent teams. Many existing approaches to auctions assume fully cooperative team members, and team members may have cooperation explicitly built in. However, on in-situ and dynamically formed teams, reciprocal collaboration may not always be a valid assumption.

The basic auction approaches to the task allocation problem assume that team members can be trusted and have the goal of the team in mind (to reduce the overall cost).[1] These algorithms serve as a mechanism for distributed task allocation and generally do not need to consider team members' cooperation levels or performance characteristics. As such, these methods do not explicitly account for trust between team members, but assume that *a)* team members will participate in auctions that are presented to them and *b)* team members will attempt to perform tasks that are assigned to them. However, there are situations in which teams may be formed dynamically. While the team members may have the same common goal, the individuals may have different levels of interest in the cooperation. That is, some of the team members may place a higher utility on successful completion of tasks, while others are obligated to participate, but wish to conserve resources.

This paper presents an approach for dynamically forming auction partners based on observed team member performance and shared reputation information. In addition, we present the use of a trust and reputation mechanism in a practical setting. Each team member models the other individuals and these models are updated through repeated interactions. Agents can use the model to detect team members that are not contributing, and those team members can be removed from future collaboration, thereby losing the benefits of cooperation. Finally, experiments are performed in simulation on a UAV platform using this approach.

The rest of this paper is organized as follows. In Section 2, we present the background and related work for trust in multi-agent auctions. In Section 3, we discuss the use of a trust model applied to multiple dimensions of trust in an auction framework. In Section 4, we present results of simulated experiments in which agents learn trust models for non-cooperative team members and exclude those members from future auctions. Finally, in Section 5, we conclude and present future work.

* pippin@gatech.edu, phone: 404-407-7904

Ground/Air Multisensor Interoperability, Integration, and Networking for Persistent ISR III,
edited by Tien Pham, Proc. of SPIE Vol. 8389, 838919 · © 2012 SPIE
CCC code: 0277-786X/12/$18 · doi: 10.1117/12.919551

2. MOTIVATION AND RELATED WORK

In traditional multi-agent systems approaches, each team member explicitly operates as part of a team and has the team's goals either explicitly or implicitly encoded. Future robotic teams may have different internal goals as well as operational capabilities. Not only will it be important for statically organized teams of heterogenous vehicles to work together in dynamic and changing environments, but also for dynamically formed teams of heterogenous vehicles to work together. Such teams may need to learn which team members are trustworthy and capable, and dynamically form their team composition accordingly.

Pippin and Christensen investigated the use of incentives to enforce cooperation on multi-agent teams.[2] The teams members performed auctions and observed which team members participated. Team members that did not participate were isolated from future cooperation with the rest of the team and this resulted in worse performance.

Jones et al. present the problem of 'forming pickup teams' of heterogeneous, cooperative robots to perform tasks using an auction framework. The ability to form dynamic teams has several advantages: robots may be expensive or scarce, and it makes sense to share them across organizational boundaries; robots may need to be organized quickly into ad-hoc teams (such as at disaster locations), and robots should be easily replaced when they fail.[3] The auction framework presented allows for the specification of roles that are needed to perform a task. Robots that are added to the team are labeled with the roles that they can fulfill. Each task that is given to the system is announced to all team members, and they negotiate task assignments based on local bid estimates and the roles that they can perform. This assumes, of course, that the robots all negotiate using the same auction protocol, that they accurately define roles and calculate utility using the same basis and that they correctly perform tasks that they bid on.

In other works, robots learn to form altruistic models of trust for determining bidding rules.[4,5] Altruism is defined as the amount of cost (in terms of time) that a robot is willing to spend to perform a task for another. This approach relies on a control law to drive the level of altruism that $robot_i$ will allow for $robot_j$ to be the observed level of altruism displayed by $robot_j$. This level of altruism is used to determine whether a robot will bid on another's task, if the task cost is less than that amount.

An approach for learning trust strategies is described by Fullam and Barber.[6] That work enumerates the types of decisions and strategy profiles that an agent can learn and compares reputation based strategy learning (based on indirect observations) with experience based learning (based on direct observations.) Other work investigates reputation with the concept of multi-dimensional trust.[7] Trust can be described by different characteristics, such as quality, reliability and availability. They show that modeling trust with multiple dimensions can lead to greater agent rewards.

Matei, Baras and Jiang investigated several different approaches to a trust model representation, including the use of continuous and discrete numerical models, binary models and probabilistic models.[8] In that work, trust is applied to the information fusion problem, by incorporating it into a Kalman filter process. Trust is describe as being multi-dimensional, based on the domain. For instance, in computer networks, trust can refer to the trustworthiness of a sensor (whether it has been compromised), the quality of data from the sensor, or the security of the link between sensors. That work defines trust as a robot performance metric, however, the use of multiple trust dimensions is instructive.

3. APPROACH

3.1 Dimensions of Trust

This paper investigates the use of observation based trust for determining when to remove a non-cooperative team member from an auction team by ignoring its auction requests. If an agent is no longer on the team, it loses opportunities for others to assist it with tasks when those tasks could be done more efficiently as part of a team than alone. In the auction context, agents that do not bid on each other's tasks, or complete them successfully can be viewed as uncooperative and removed from a team. From an agent's viewpoint, it is better to have team members that cooperate and participate in the auction algorithm as this leads to more efficient outcomes. From a global viewpoint, it is desirable to have an efficient team that is composed of cooperative members; each uncooperative member decreases the overall team performance. Finally, in this work, we assume that a currency exchange mechanism is not available for enforcing cooperation. In teams that are dynamically formed or consist of temporary alliances, it is reasonable to assume that such an exchange and accounting mechanism may not be present.

In a dynamically formed auction team, agents may encounter other agents for which they have no prior experience. The use of a trust model would allow for an agent to reason about other agent's trustworthiness using observation histories and reputation information. In these settings, there are multiple dimensions that could be used to define trust, such as whether an agent participates in the auctions of others and whether an agent successfully completes tasks that are assigned to it. Agents that regularly violate the trust dimensions are considered to be *defectors*, while agents that cooperate fully are labeled as *cooperators*. These dimensions of trust can be considered separately or in combination. Each agent can build models of other team members behaviors from observation histories and use those models to determine levels of trust.

3.1.1 Bid Participation

The *Bid Participation* dimension considers whether an agent is participating by submitting bids as part of the auction process, but does not evaluate the bid. In fact, there are legitimate situations in which an agent might not wish to submit a bid, if the calculation is costly.[9] However, in domains in which the bid calculation is easily computed and communicated, this can be a useful gauge of auction participation. When an agent announces an auction, it keeps track of the agents that received the auction announcement and compares this to the list of agents that submitted bids. This approach assumes that agents are uniquely identifiable and that a protocol exists for acknowledging the receipt of an announcement. When an agent does not bid on received auctions announcements, this negatively updates the trust model, while the submission of a bid positively updates the model (described further in section 3.2.) Using the model, if an agent determines that a team member is not trusted, the agent refuses to bid on the untrusted team member's future auctions, effectively isolating it from the auction team. In this paper, we will primarily discuss the bid participation dimension; however, the trust model presented in the next section could be used to combine additional dimensions into a single trust valuation.

3.2 Trust Model

This work relies on the use of a a probability based trust model, using the Beta distribution.[10,11] We rely especially on the trust mechanism from Teacy, et al.[10] for incorporating direct trust and reputation into a probabilistic formulation. This mechanism provides not only a trust belief about an agent, but also a confidence value. The approach can incorporate positive and negative histories (direct observations) to calculate the belief and confidence.

Each agent maintains a set of α and β vectors that represent the histories of interactions with each team member. For a given team member, if the calculated trust value is less than the trust threshold, τ, and with confidence greater than γ, it is not trusted. However, a succession of positive observations (direct or indirect) can move an untrusted agent back to being trusted again. Furthermore, this approach is tolerant of noise as it can take multiple observations to move the value above or below the trust threshold. To better explain this model, the equations from Teacy, et al.[10] for calculating the trust value τ and confidence, γ, are included below.

When an agent receives new α and β updates for a dimension of trust, it can calculate the Expected Value for trust using the trust model as follows.

$$E_{trust_{i,j}} = \frac{\alpha}{\alpha + \beta} \tag{1}$$

The value, $E_{trust_{i,j}}$, is the expected trust that $robot_i$ has toward $robot_j$, given a set of observations. Therefore, the trust value, τ, is

$$\tau = [E_{trust_{i,j}} | O^{1:t}] \tag{2}$$

The confidence factor, γ, is calculated as the proportion of the beta distribution that is within ϵ of τ.

$$\gamma = \frac{\int_{\tau-\epsilon}^{\tau+\epsilon} X^{\alpha-1}(1-X)^{\beta-1}\,\mathrm{d}X}{\int_0^1 U^{\alpha-1}(1-U)^{\beta-1}\,\mathrm{d}U} \tag{3}$$

3.3 Shared Reputation

In addition to the observations from direct interactions with other agents, this approach allows for the agents to incorporate indirect observations from other trusted team members, known as shared reputation information. The intent for sharing reputation information among team members is to quickly spread information about trusted or untrusted agents to the rest of the team. If an agent can rely on reputation information from other agents, it might be spared from negative direct interactions with uncooperative agents. However, the shared reputation information must be combined with the locally observed trust vectors. In the auction framework, each agent can regularly post their trust model's α and β vectors to all other team members that are within range, and agents only incorporate those updates from other currently trusted team members.

3.4 Auction Approach using the Trust Model

In the basic multi-agent auction algorithm, the problem is to assign tasks to agents. The tasks in this case are to visit a target location and perform an observation. In the auction framework, each robot is a *bidder* and the items to be auctioned are the 'visits'. Each of the agents in the system also participates as an *auctioneer* and periodically auctions new task requests (it is assumed that the task requests are provided to the agent by an external process, such as a human operator or other event). This approach can easily be used on teams with different robot characteristics: each robot knows their own location and cost function and submits cost based bids to the auctioneer. While costs and rewards use the same basis for calculation, no revenue is actually exchanged. Rather, an agent awards itself a utility value when one of its own tasks is completed.

The *auctioneer* first handles any auctions that have already been announced and are ready to close. This step is shown in detail in Procedure 1. In lines 1-3, the *auctioneer* selects the minimal cost bid from all bids received by the agents within communications range (including their own) as the winner of that auction and performs the task assignment by announcing the winning bidder. In lines 5 and 7, the *auctioneer* updates the trust model (described in Section 3.2) for each possible *bidder* that was sent the auction announcement. The trust model is referenced by the *bidder* in Procedure 2, when an auction announcement is received. If the originator of the auction announcement is not trusted, then the auction announcement is ignored, effectively isolating the untrusted agent from the benefits of cooperation.

Procedure 1 *Auctioneer* :: $HandleAuctionBids$

Input: An auction, a.
Input: The set of posted bids, B_a.
Input: The set of announcement recipients, $Recipients_a$.

1: $winner \leftarrow Min(B_a)$
2: $AnnounceWinner(winner, a)$
3: **for all** $a : Recipients_a$ **do**
4:   **if** $a \in B_a$ **then**
5:     $UpdateParticipation(TRUST_o, 1)$
6:   **else**
7:     $UpdateParticipation(TRUST_o, 0)$
8:   **end if**
9: **end for**

4. EXPERIMENTAL RESULTS

4.1 Multi-Agent Experiments

A set of experiments were performed to test the trust strategies in simulated auctions using the Mason simulated multi-agent environment.[12] The Mason environment was used to a run large number of low-fidelity simulations to demonstrate the trust model. In these experiments, the agents are represented as Unmanned Aerial Vehicles (UAVs), modeled as points in a 2d plane, and each UAV is assigned tasks to perform by an external process. Each UAV has an auctioneer and can auction their tasks to other agents, assigning the task to the agent that submits

Procedure 2 *Bidder* :: *HandleAnnouncements*(A)

Input: An set of announced auction tasks, A.
Input: The auction originator trust model, $TRUST_o$.

```
for all a : A do
  if CanTrust(TRUST_o) then
    bid ← CalculateBid(a)
    if bid > 0 then
      PostBid(bid)
    end if
  end if
end for
```

the minimal cost bid. The UAVs in the simulation have a limited communications range and can therefore only perform auctions or exchange reputation information with a subset of the other team members at a given time.

In addition, each UAV periodically re-auctions the last n tasks to other agents in range. This allows tasks to be more optimally assigned by giving other agents a chance to bid on them if they were not in range during the initial auction. Each experiment was performed using 10 UAVs, with results averaged over 100 iterations. Each UAV has 50 tasks that arrive at regular intervals and are sequentially auctioned. The initial locations of the UAVs and the tasks are randomly chosen for each iteration. The results show the average score for each of the *cooperator* agents as the auctions are completed and rewards are assigned.

In this experiment, a fraction of the agents on the team defect by not participating in auctions (not bidding on others' tasks). Each *defector* agent only participates in auctions 10% of the time. At this level they are occasionally participating but do not contribute effectively. As a result, *Naive* agents that trust unconditionally (using no trust mechanism) end up doing additional work for the *defector* agents and receive little in return. The objective for the *cooperator* agents is to detect those team members that regularly fail to participate in auctions and to isolate them from future cooperation by not bidding on the *defectors'* tasks.

For each auction, the trust strategies update the trust model for each agent that was sent an auction announcement. If the agent submitted a bid, the trust model is updated, ($Trust_{a,t} = 1$), and ($Trust_{a,t} = 0$) otherwise. Once an agent is no longer trusted, with high confidence, they are removed from future auction participation as the *cooperators* isolate them by refusing to bid on their tasks. The results of this experiment, shown in Figure 1, reflect that the agents running the *Beta Trust* and *Reputation* strategies receive better scores than those that apply the *Naive* strategy or the *No Cooperation* strategy, once the agents are able to observe which team members are participating in the auctions. In addition, the *Reputation* strategy which shares trust information across team members performs better than the *Beta Trust* strategy which relies on direct observations alone.

4.2 UAV Platform Experiments

The UAV platform leverages off-the-shelf, readily available components, and is based on a quarter-scale Piper Cub airframe with a base model Piccolo avionics and autopilot system from Cloud Cap Technology.[13] The airframe has a wingspan of 104 inches, and carries a mission computer and sensor payloads, see Figure 2(a). Over 60 field tests of this platform have been performed, including multi-UAV cooperative autonomy and UAV-UGV teaming demonstrations.

The platform can also be tested in high-fidelity simulations, as shown in the simulation architecture diagram in Figure 2(b). The flight dynamics of each UAV are simulated using the *software in the loop* (SIL) capabilities of the Piccolo autopilot. In addition the auction algorithms that run on the mission computer are executed within a separate Linux virtual machine (VM) for each aircraft to be simulated. Messages are sent to the FalconView™map display using simulated radio messages. Vehicle positions and assigned waypoints are displayed over the FalconView™map as shown in Figure 3(a). In addition, the UAVs in simulated flight are displayed in the MetaVR™visualization as shown in Figure 3(b).

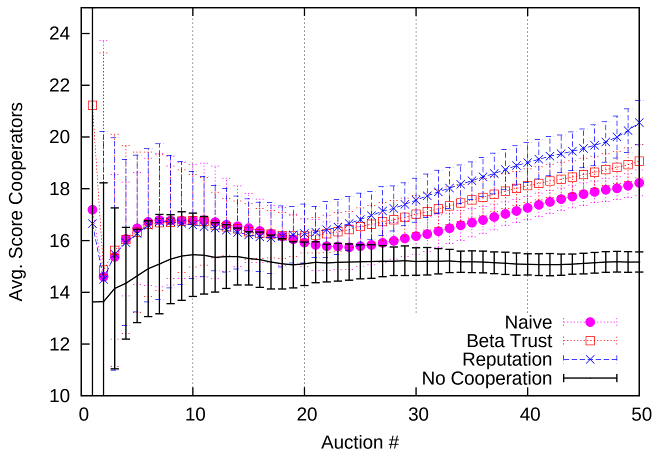

Figure 1: Bid Participation: Agents that defect by not bidding can be detected and isolated using observation based trust mechanisms. The average unit score of the cooperative agents is plotted against the number of auctions completed for the different trust strategies. The error bars reflect one standard deviation.

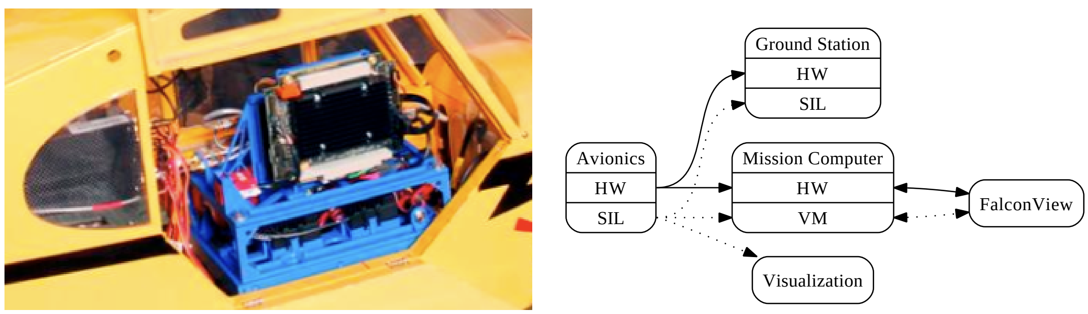

(a) Autonomy Payload (b) Architecture

Figure 2: The UAV Simulation Architecture. (a) The UAV platform carries an autonomy payload, consisting of the autopilot and mission computer. (b) The architecture can support the autonomous behaviors in real flight or simulation.

Again in this set of experiments, a fraction of the UAVs do not participate in group auctions, but exploit the others on the team by allowing them to perform tasks on behalf of the *defectors*. In this case, we simulated 4 UAVs flying in different sectors of the environment and awaiting tasks from their operators. To simulate operators, a separate process regularly assigns a new task to a UAV that is picked at random. The location of the task varies in the environment, over an area of 15x10 km. The simulation ends after 100 tasks have been assigned and completed. When a UAV is assigned a task by their operator, it has the responsibility of completing it. However, a UAV has the option of auctioning the task to another member on the team. No currency is exchanged in this domain, but rather the vehicles consist of a loosely formed team that can benefit from cooperation because the tasks are distributed throughout the environment.

In one set of experiments, the cooperative team members perform a basic or *naive* auction strategy, and do not consider whether other UAVs have reciprocated cooperation. In the second set, the cooperative team members apply the direct observation based trust mechanism that was described in Section 3.2 to detect team members that do not participate in auctions by bidding on other's tasks. The results are shown in Table 1. The average task cost represents the amount of time (in seconds) that was taken to complete the task by the UAV that won the auction.* The average time spent represents the amount of time spent performing assigned and won tasks. When the *naive* auction strategy is used, the cooperators are exploited by the defectors and spend a

*A task could be completed by a different UAV than the one it was initially assigned to, if it was re-assigned as part of an auction.

Table 1: Auction Participation with the Trust Model

		Global	Cooperators	Defectors
Naive Auction	Avg. Task Cost	1282	1172	1383
	Avg. Time spent	1282	1306	79
Beta Trust Model	Avg. Task Cost	1138	850	1415
	Avg. Time spent	1138	850	1415

much greater amount of time performing tasks. The tasks are performed less efficiently as a result, because the cooperators are doing most of the work. In contrast, when the trust model is employed, the cooperators quickly learn to not trust the defectors, and isolate them from further participation by refusing to bid on the untrusted team members' auctions. This results in a lower global cost, and the cooperators have a much lower cost and time spent as well, because they are not being exploited in this case.

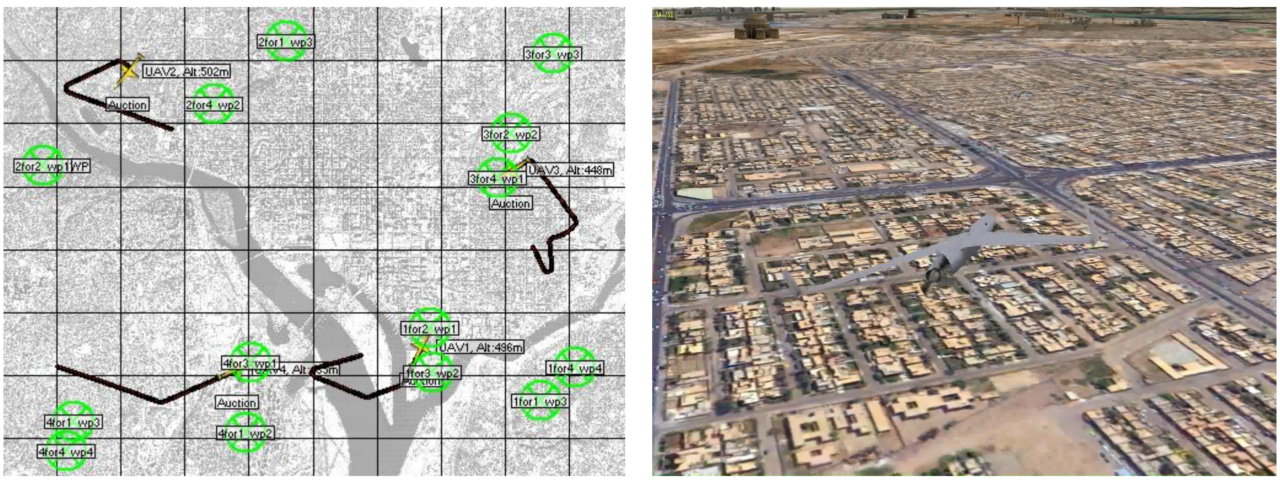

(a) Simulating Multiple UAVs

(b) UAV Visualization

Figure 3: Simulating multi-UAV auctions: Multiple UAVs are simulated using the autopilot and autonomous auction behaviors. (a) The UAVs and assigned waypoints are shown in FalconView.$^{\text{TM}}$(b) A simulated UAV is shown rendered in a visualization using the MetaVR$^{\text{TM}}$scene generation tool.

5. CONCLUSIONS AND FUTURE WORK

The above experiments showed that trust and reputation mechanisms can be effective for detecting and isolating uncooperative team members in auctions, when compared to the naive approach. This may prove useful in situations in which auction based teams are dynamically formed and not all team members are likely to participate equally.

Traditional auction algorithms for performing the robot task assignment problem assume that robots are equally incentivized to participate in auctions. However, there are situations in which agents may assign tasks to others on the team, without taking on a fair number of additional tasks in return. This paper presents an approach for using observation based trust and a shared reputation mechanism in determining which agents to include in multi-agent auctions. The results show that by incorporating the use of trust strategies into the basic auction mechanism, agents can perform better than agents that trust unconditionally. There are legitimate situations in which a team member may not be able to participate in auctions, such as when the agent is not capable of performing the task or is otherwise preoccupied. However, the trust model presented is able to tolerate noise in observations and also can incorporate forgiveness when an agent is able to participate again. Rather, the focus is on detecting those team members that exploit the team and isolate them from the benefits of future cooperation.

Future work will consider additional learning mechanisms relevant to task performance. This is related to the problem of determining how to recognize when tasks that were assigned to another agent were not only completed according to the initial cost estimate, but completed within stated quality parameters. In addition, we would like to perform additional UAV experiments to explore how performance data can be used to affect either the task assignment function or to again perform UAV team formation.

REFERENCES

[1] Koenig, S., Keskinocak, P., and Tovey, C., "Progress on agent coordination with cooperative auctions," in [*AAAI Conference on Artificial Intelligence*], (2010).

[2] Pippin, C. and Christensen, H., "Incentive based cooperation in multi-agent auctions," in [*AAAI 2012 Spring Symposia - Game Theory for Security, Sustainability and Health*], (March 2012).

[3] Jones, E., Browning, B., Dias, M. B., Argall, B., Veloso, M., and Stentz, A. T., "Dynamically formed heterogeneous robot teams performing tightly-coordinated tasks," in [*International Conference on Robotics and Automation*], 570 – 575 (May 2006).

[4] Morton, R. D., Bekey, G. A., and Clark, C. M., "Altruistic task allocation despite unbalanced relationships within multi-robot communities," in [*Proceedings of the 2009 IEEE/RSJ international conference on Intelligent robots and systems*], *IROS'09*, 5849–5854, IEEE Press, Piscataway, NJ, USA (2009).

[5] Clark, C. M., Morton, R., and Bekey, G. A., "Altruistic relationships for optimizing task fulfillment in robot communities," in [*Distributed Autonomous Robotic Systems 8*], Asama, H., Kurokawa, H., Ota, J., and Sekiyama, K., eds., 261–270, Springer Berlin Heidelberg (2009).

[6] Fullam, K. K. and Barber, K. S., "Learning trust strategies in reputation exchange networks," in [*Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*], *AAMAS '06*, 1241–1248, ACM, New York, NY, USA (2006).

[7] Ahn, J., DeAngelis, D., and Barber, S., "Attitude driven team formation using multi-dimensional trust," in [*Intelligent Agent Technology, 2007. IAT '07. IEEE/WIC/ACM International Conference on*], 229 –235 (Nov. 2007).

[8] Matei, I., Baras, J., and Jiang, T., "A composite trust model and its application to collaborative distributed information fusion," in [*Information Fusion, 2009. FUSION '09. 12th International Conference on*], 1950 –1957 (july 2009).

[9] Busquets, D. and Simmons, R., [*Learning when to Auction and when to Bid*], vol. Distributed Autonomous Robotic Systems 7., 21–30, Springer (2006).

[10] Teacy, W. T. L., Patel, J., Jennings, N. R., and Luck, M., "Travos: Trust and reputation in the context of inaccurate information sources," *Journal of Autonomous Agents and Multi-Agent Systems* **12**, 2006 (2006).

[11] Commerce, B. E., Jøsang, A., and Ismail, R., "The beta reputation system," in [*In Proceedings of the 15th Bled Electronic Commerce Conference*], (2002).

[12] Luke, S., Cioffi-Revilla, C., Panait, L., Sullivan, K., and Balan, G., "MASON: A multi-agent simulation environment," *Simulation: Transactions of the society for Modeling and Simulation International* **82**((7)), 517–527 (2005).

[13] Vaglienti, B., Hoag, R., Niculescu, M., Becker, J., and Miley, D., *Piccolo User's Guide (v2.1.0).* Cloud Cap Technology, www.cloudcaptech.com (October 14, 2009).