A new perspective of trust through multi-attribute auctions

Ferran Torrent-Fontbona, Albert Pla, Beatriz López
Univeristy of Girona (Spain)
Control Engineering and Intelligent Systems Group (eXiT)
{ferran.torrent,albert.pla,beatriz.lopez}@udg.edu

Abstract

The use of trust in auctions is a well-studied problem; however, most of the works in the literature focus on how to model trust rather on how trust is used in the mechanism. In this paper we propose a method to manage trust in multi-attribute auctions. Complementary, we propose a new trust model and we compare it with other models in the literature. The proposed methodology is tested using a real data-based simulator, showing the benefits of incorporating trust in multi-attribute auctions.

Introduction

Auction mechanisms are very well known methods to allocate tasks when several agents are involved. Tasks' prices are established as a result of the auctions. Some mechanism such as Vickrey-Clark-Grove (VCG) (MacKie-Mason and Varian 1994) incentivize bidders to be honest providing their price (bid truthfully). On the other hand, multi-attribute auctions are a particular mechanism that enables the consideration of task attributes other than prices, as delivery time or energy consumptions. Our research is concerned with such kind of multi-attribute scenarios, where industries have to focus on the delivery times of their providers, in addition to price. Moreover, social concern about environment and global warming conveys that industries have to deal with their environmental impact and footprint of their products. Thus, a task allocation problem is not only about cost and delivery time, but also about other attributes like energy consumption or the environmental footprint.

Incentive compatibility regarding multi-attribute auctions based only on prices do not work if other attributes, like task delivered time or energy consumed, committed during the bidding process should also be fulfilled. In (Pla et al. 2014), the authors, facing the impossibility of guaranteeing this desirable property, propose a mechanism to minimize the auctioneer's loss of utility when bidders do not keep their word. However, the breach of contract by some bidder could be not only about agents who intentionally provide false bids in order to dishonestly try to win auctions but also about the accuracy of bidders' estimations about their performances. Incentive compatible mechanisms encourage agents to reveal the attributes which agents estimate truthful, however, thus mechanisms by themselves cannot know if such estimations are reliable or not due to the uncertainty caused by the available data to bidders (Jurca and Faltings 2003; Pla et al. 2014). Under such circumstances, trust could complement incentive compatibility reducing the risk of losses by the auctioneer.

This work presents a new perspective of trust in a multi-attribute framework. The trust model is multi-facet, so the auctioneer keeps track of each verifiable (e.g. traceable, checkable) attribute provided by bidders (that is, all attributes but price). Using separated trusts provides a higher flexibility (Pinyol and Sabater-Mir 2013). For instance, in a moment with a high work load with a tight schedule an auctioneer might be more concerned about delivery times than to energy consumptions and therefore it could give more importance to being reliable on delivering a task at the agreed time. Using a global trust, the agent would not be able to distinguish between which agents are reliable in terms of time and which are reliable in terms of energy consumption. While keeping separated trusts, the auctioneer knows which bidders are more reliable regarding an attribute.

The main goal of this paper is to provide an easy and systematic way to include a multi-faced model of trust into multi-attribute auctions. Conversely to other previous works where trust is only used in the winner determination problem, in our approach, trust is used both in deciding the winner of the auction and the payment to the corresponding bidder. Taking into account trust also in the payment reduces the losses of the auctioneer, defining positive synergies between truthful biding and trust. In addition, we also present a simple trust learning methodology to test our approach. Experiments show promising results compared with other state of the art models.

This paper is organized as follows: first we present some works regarding trust in multi-agent systems and particularly in auctions. We continue by presenting the methodology we propose in this paper. Next, we show the experimentation we performed to analyze our methodology and we discuss the results. Finally, we provide the conclusions of this work, and future research directions.

Related Work

Trust and uncertainty have been widely studied in the auction literature. Many authors have proposed trust mechanisms based on the experiences of other buyers like the methods presented in (Regan and Cohen 2005; Schillo, Funk, and Rovatsos 1999; Schillo, Funk, and Rovatsos 2000; Hu et al. 2013) and the Dirichlet reputation systems explained in (Jø sang, Ismail, and Boyd 2007). In these works, buyers evaluate sellers and share this information with other buyers. These mechanisms are useful when there are a lot of buyers that can share information regarding the sellers to help each other. However, the scenario we consider in this paper consists of one buyer (auctioneer) that wants to outsource some tasks. In an scenario like that, Porter et al. (Porter et al. 2008) presented an auction mechanism which deals with uncertainty by the inclusion of a confidence attribute defining the probability of a bidder to accomplish its task (POS). This attribute is provided by the bidders and influences both the winners of the auction and their pay-off. In (Ramchurn and Mezzetti 2009), the authors proposed an extension of the previous mechanism including the sense of reputation: the perception that bidders have about a bidder of being reliable or not. Conversely to these approaches, we base trust in the auctioneer's perception of the bidders. Moreover, all these approaches are designed for uni-attribute auctions while we intend to tackle a multi-attribute auction problem.

Other authors have also presented trust mechanisms based on the direct experience of the buyer with a particular seller. Some examples are the methods presented in (Marsh 1994) and (Regan and Cohen 2005). These methods, as those explained in (Schillo, Funk, and Rovatsos 1999; Schillo, Funk, and Rovatsos 2000; Jø sang, Ismail, and Boyd 2007), calculate trust on a particular seller and then the buyer decides to either commerce with it or not depending on such trust. As well as these approaches, we also use trust mechanisms based on direct experiences of the buyer, nevertheless, we do this in a multi-attribute auction setting.

Besides, for a deeper analysis, the reader can find exhaustive surveys of trust in multi-agent systems in (Chiu, Huang, and Yen 2010) and (Pinyol and Sabater-Mir 2013).

Beyond, in the literature exist other mechanisms than trust to incentivize agents to behave honestly. For example Vickrey-Clark-Grove (VCG) (MacKie-Mason and Varian 1994) based systems incentivize bidders to bid truthfully. In this paper we follow the VCG based method called VMA2 (Pla et al. 2014) which is a multi-attribute auction system. Despite VMA2 has been proved to incentivize bidders to send bids with attribute values that they can fulfill, it is vulnerable when bidders are not able to accurately estimate the values of the attributes of the tasks they are bidding for. Thus, in this paper we include trust as a mechanism to protect the auctioneer from involuntary errors of the bidders during the bidding process.

Methodology

We propose a multi-attribute mechanism to deal with a multi-faced model of trust. Our multi-attribute mechanism distinguish three kinds of attributes involved in an auction (Pla et al. 2014):

- Unverifiable attributes: as the prices.
- Verifiable attributes: as the delivery time, or energy consumed. They could be checked by the auctioneer upon the reception of the tasks
- Auctioneer provided attributes: managed by the auctioneer and related to bidders information

Both, verifiable and unverifiable attributes are provided by bidders in response to the call for proposals requested by the auctioneer. On the other hand, auctioneer provided attributes imply that auctions are repeated over time, so the auctioneer keeps track of an history of auctions outcomes and bidders behaviors. Therefore, a multi-faced trust model can be set as a collection of auctioneer provided attributes, one per each of the verifiable attributes. In consequence the use of trust in a multi-attribute auction has an advantage: a mechanism which fails in guaranteeing incentive compatibility, become less vulnerable to cheaters agents.

Following, the different steps of our method are described according to the auction protocol. As we are using auctions for allocating or outsource tasks to third party companies, we follow a reverse auction schema: an auctioneer needs a task to be done and offers to pay an external provider for carrying it out (becoming the buyer who aims to buy a service at the cheapest price) while bidders offer their working capacity at a given price (becoming the sellers who compete to offer the best working conditions at the cheapest price).

Call for Proposals

The first step consists of the auctioneer offering a particular task to do under some constraints of time and energy (attributes) to the agents available in the market. During the call for proposals, the auctioneer makes public the valuation function of the bids it will use to combine the attributes together with the price.

Bidding

We consider that bidders are greedy agents which are willing to perform activities in order to increase their own utility. Once a bidder receives the auctioneer's proposal, if it is interested in the auctioned task, it offers a bid with the conditions (price, time, energy) under which the bidder can perform the task. We define the bid of bidder i as the set of attributes that describe the conditions of the task execution

$$B_i = \langle b_i, t_i, e_i \rangle \tag{1}$$

where b_i is the price of the bid, t_i is the delivery time and e_i is the energy consumption. The price is a unverifiable attribute since the auctioneer will never know the true value of the bid; conversely, the delivery time and the energy consumed are verifiable attributes, because if the bidder wins the auction, the auctioneer can check them. As the auctioneer is buying tasks, the lowest the price the best; regarding delivery time and energy consumed, the lowest the best as well.

Despite agents can perform the task with different configurations of attributes' values (e.g. faster but also with a higher energy consumption or price, or slower with a lower energy consumption),we assume that they only need to send the best option according to the valuation function made public by the auctioneer.

Winner Determination Problem

Once the period of receiving bids is closed, the auctioneer must decide who the winner of the auction is: the bidder who offered the bid that maximizes the auctioneer expected utility (Ramchurn and Mezzetti 2009). Given a set of attributes $a_1, \ldots, a_n$, the utility u of the auctioneer can be defined as follows:

$$u\left(T_0, a_1, \ldots, a_n\right) = v\left(T_0\right) - f\left(a_1, \ldots, a_n\right) \tag{2}$$

where T_0 is the auctioned task, $v\left(T_0\right)$ is the value the auctioneer gives for having the task completed and $f\left(a_1, \ldots, a_n\right)$ is the valuation function which evaluates the bid attributes. Note that maximizing $u\left(T_0, a_1, \ldots, a_n\right)$ is equivalent to minimize $f\left(a_1, \ldots, a_n\right)$.

Among the set of available attributes, there are the attributes provided in the bidders' bids (verifiable and unverifiable attributes) as well as attributes regarding bidders information (auctioneer provided attributes).

According to the needs of the problem we are facing (where delivery time and energy consumptions are involved), we propose to introduce two trust parameters $\tau^t_{i,r}$ and $\tau^e_{i,r}$ as auctioneer provided attribute. $\tau^t_{i,r}$ and $\tau^e_{i,r}$ define the confidence the auctioneer has in bidder i at round r regarding time and energy attributes according to its past experience. Of course, more trust parameters could be added if more verifiable attributes were available. Both trust attributes, $\tau^t_{i,r}$ and $\tau^e_{i,r}$, are defined in $(0, 1]$, and the higher the trustee.

Therefore, we propose to maximize the expected utility with the chances the bidder has to fail delivering the task in the agreed conditions according to the following expression:

$$\overline{u}\left(T_0, b_i, t_i, e_i, \tau^t_{i,r}, \tau^e_{i,r}\right) = v\left(T_0\right) - V\left(b_i, \frac{t_i}{\tau^t_{i,r}}, \frac{e_i}{\tau^e_{i,r}}\right) \tag{3}$$

where $V\left(b_i, \frac{t_i}{\tau^t_{i,r}}, \frac{e_i}{\tau^e_{i,r}}\right)$ is the expected valuation that the auctioneer gives to the bid B_i. As stated above, the lowest delivery time and energy consumed the better outcome. Therefore, dividing the delivery time and energy consumed values provided by bidders by the corresponding trust value, results in a augmented value for untrusted agents, and so, a lower chance to become the winners. The new value can be seen as the value of the attribute plus a security margin for the auctioneer.

According to Equation 3, solving the winner determination problem means to minimize the value of V:

$$\min_i \left\{ V\left(b_i, \frac{t_i}{\tau^t_{i,r}}, \frac{e_i}{\tau^e_{i,r}}\right) \right\} \tag{4}$$

In this paper we use the weighted sum as evaluation function (see (Pla et al. 2014) for alternative evaluation functions):

$$V\left(b_i, \frac{t_i}{\tau^t_{i,r}}, \frac{e_i}{\tau^e_{i,r}}\right) = w_0 b_i + w_1 \frac{t_i}{\tau^t_{i,r}} + w_2 \frac{e_i}{\tau^e_{i,r}} \tag{5}$$

subject to $\sum_k w_k = 1$.

Payment rule

The payment rule is used to establish the economic amount p_i that the auctioneer must pay to the auction winner (agent a_1) after performing a task. It is a key aspect for ensuring the incentive compatibility of an auction mechanism. Due to the multidimensional nature of the allocation problem we are dealing with, the payment is not only conditioned by the price of the bid but also the value of the rest of attributes.

In such situations, there are no mechanisms guaranteeing incentive compatibility, but (Pla et al. 2014) propose a two case method depending on whether the bidder delivers the task as agreed or not, that minimize auctioneer losses in case of cheater agents participation. In the case the task is successfully delivered, the payment will be carried out following a classical Vickrey (or second price) schema, meaning that the winner of the auction will receive the economic amount it should have offered in order to obtain the same evaluation as the second best bid[1], as follows:

$$V\left(p_1, \frac{t_1}{\tau^t_{1,r}}, \frac{e_1}{\tau^e_{1,r}}\right) = V\left(b_2, \frac{t_2}{\tau^t_{2,n}}, \frac{e_2}{\tau^e_{2,n}}\right) \tag{6}$$

where the sub-index j indicates the *second best* bid sent in the auction. Thus, b_2, t_2, e_2 are the attributes of the second best bid and $\tau^t_{2,n}$ and $\tau^e_{2,n}$ are the trust attributes the auctioneer gives to the bidder of this second best bid. Given that in the winner determination problem, all of the attributes including trust are evaluated together, the payment rule need to use all those parameters in order to asses the payment corresponding to the auction winner. To this end, we are assuming that the auctioneer is not able to change (intentionally or not) the trust values assigned to each bid.

In the case the bidder delivers the task in worse conditions, the bidder receives a smaller payment in such a way that the valuation of the initially presented bid matches with the valuation of the actual delivered task, as follows:

$$V\left(p_1, \frac{t'_1}{\tau^t_{1,r}}, \frac{e'_1}{\tau^e_{1,r}}\right) = V\left(b_1, \frac{t_1}{\tau^t_{1,r}}, \frac{e_1}{\tau^e_{1,r}}\right) \tag{7}$$

where t'_1 and e'_1 are the real delivery time and energy consumption respectively. This payment will avoid the auctioneer to be harmed in case of receiving a task in worse conditions than its valuation during the winner determination problem.

[1] In case of a draw, a tie-breaking rule should be used. In such circumstance both, the best and the second best bidders, will obtain 0 payoff (Maskin and Riley 2003).

Therefore, if we define a $V^{-1}(v,\ldots,a_n)=a_1$ as the reverse function of $V(a_1,\ldots,a_n)=v$ which given the evaluation of a bid v and the non-economic attributes of such bid, returns the economic attribute of the bid, the payment is defined as follows:

$$p_1=\begin{cases}V^{-1}\left(V\left(b_2,\frac{t_2}{\tau^t_{2,n}},\frac{e_2}{\tau^e_{2,n}}\right),\frac{t_1}{\tau^t_{1,r}},\frac{e_1}{\tau^e_{1,r}}\right) & \text{if } t_1^{'}\le t_1,\, e_1^{\prime}\le e_1\\ V^{-1}\left(V\left(b_1,\frac{t_1}{\tau^t_{1,r}},\frac{e_1}{\tau^e_{1,r}}\right),\frac{t_1^{'}}{\tau^t_{1,r}},\frac{e_1^{'}}{\tau^e_{1,r}}\right) & \text{otherwise}\end{cases} \tag{8}$$

Trust learning

After any bidder delivers a completed task, the auctioneer can collect information regarding bidder's performance (e.g. delivery on time and appropriate energy consumption) and update its trust on the bidder. If the delivered task has been successful, the auctioneer increases trust on the corresponding bidder, but if the bidder delivers a task in bad conditions, the auctioneer reduces trust on the bidder. It is important to note that the sense of success or failure will be different in every domain (e.g. in certain domains a successful task will be delivered just at a certain moment while in others a task will be considered successful if it is delivered before the deadline).

The updating function we propose for each trust attribute is given by Equations (9) and (10).

$$\tau^t_{1,r+1}=\begin{cases}\tau^t_{1,r}+\alpha_t\left(1-\tau^t_{1,r}\right) & \text{if } t_1^{'}\le t_1\\ \tau^t_{1,r}-\beta_t\tau^t_{1,r} & \text{otherwise}\end{cases} \tag{9}$$

$$\tau^e_{1,r+1}=\begin{cases}\tau^e_{1,r}+\alpha_e\left(1-\tau^e_{1,r}\right) & \text{if } e_1^{'}\le e_1\\ \tau^e_{1,r}-\beta_e\tau^e_{1,r} & \text{otherwise}\end{cases} \tag{10}$$

where α_t, β_t, α_e and β_e are coefficients in $[0,1]$ which determine the rate of reinforcement.

Results and Discussion

In this section we test this new perspective of trust as attributes in a multi-attribute auctions (in the WDP and in the payment) by means of a multi-agent system simulation where an agent tries to allocate tasks to resources (Pla et al. 2012). We use our trust learning method as well as trust learning methods of the state of the art.

Experimental set up

Simulations are based on a real business process[2] where one auctioneer outsources tasks to external agents.

For the experimentation we have modeled 6 couples of competitive and greedy bidders whose time and energy values (for executing tasks) are randomly distributed according to real probability distributions[2]. Each couple of bidders consists of two equal bidders regarding time and energy distributions, but one of them is able to exactly estimate the values of time and energy it needs to perform the tasks whilst the other one is only able to estimate the values according to the mean of the distributions. Thus, there are 6 *accurate* bidders and their *inaccurate brothers*.

Regarding incentive compatibility, bidders follow an adaptive strategy: they adapt their offers (increase or decrease their economic pretensions) according to the resulting allocations in order to increase their chances of winning the auction and maximizing their benefits (Lee and Szymanski 2005).

Finally, to study the behavior of trust we have tested the following methods:

- *No trust*: no trust model is used.
- *T-Trust model*: this is the trust learning method proposed in this paper with the learning algorithm of the previous section.
- *Schillo model*: the trust learning method is taken from (Schillo, Funk, and Rovatsos 2000) and consists of calculating the honesty of a bidder checking what it claimed and what it finally did. The estimated probability of a bidder of being honest is then $\frac{h}{n}$ where h is the times it has been honest (regarding time or energy) in the past and n is the number of task delivered.
- *Dirichlet models*: the trust learning method is described in (Jø sang, Ismail, and Boyd 2007) and consists of rating the task delivered by the bidders according to a discrete and finite set (e.g. $\{$*very bad, bad, average, good, very good*$\}$). Then the auctioneer calculates a probability distribution according to this set, which represents the probability that the bidder has to act as stated in each one of the categories.

We have set up two experiments:

- *Trust versus no trust*: we compare the outcomes of a multi-attribute auction without trust with the method presented in this paper. Our hypotheses is that exploitation of a multi-attribute mechanism for using trust results in a reduction of the auctioneer losses.
- *Trust models*, where we compare the different trust learning methods.

Results: Trust versus no trust

The performed experiments are evaluated using the percentage of bad delivered tasks and the percentage of winner bids from unreliable bidders. The first metric evaluates the reliability of the resulting allocations, where a high percentage of bad delivered task implies poor reliability on the resulting allocations (the auctioneer cannot rely that its tasks will be successfully performed). The second metric is useful to evaluate if for a bidder it is important or not to be reliable, indicating if a bidder wins more auctions when it is or when it is not reliable.

Figure 1 shows the percentage of bad delivered tasks (tasks delivered without the agreed conditions) and the percentage of winner bids from inaccurate bidders using different trust methods and not using trust. The initial trust value used in all the models have been 0.5 while α_t and β_t values of our approach have been set to 0.1. As expected, the results tell us that the use of trust reduces the number of winner bids

[2]Data available at `http://eia.udg.es/~apla/fac_data/`

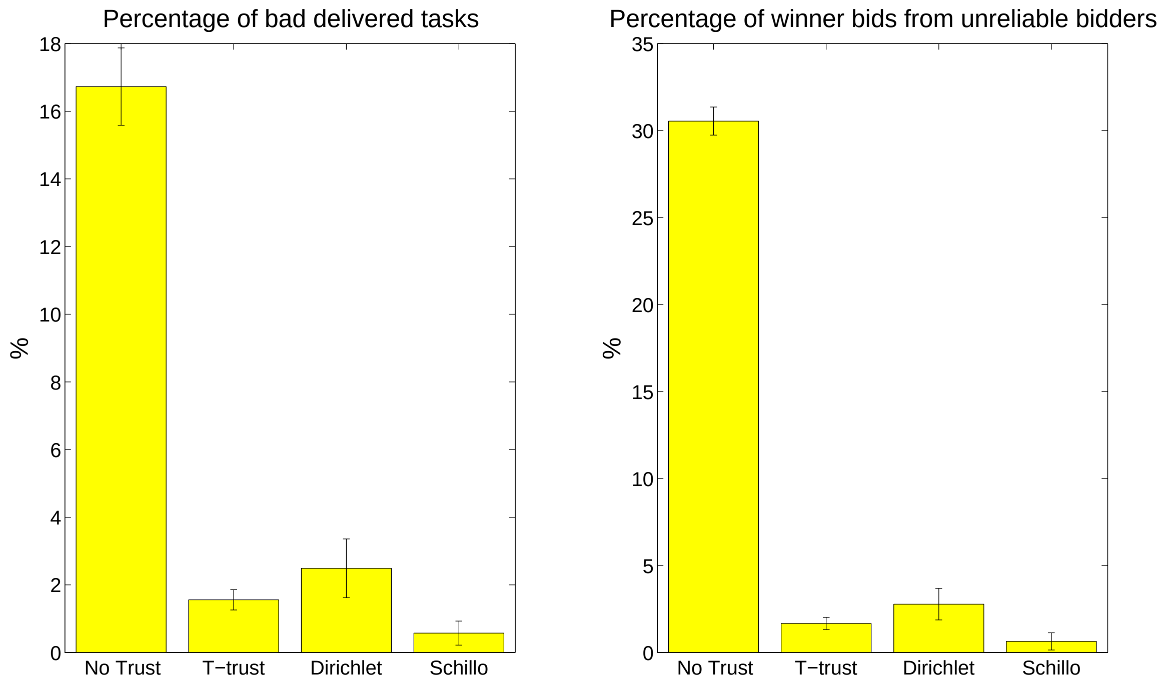

Figure 1: On the left, percentage of tasks delivered in worse conditions than the agreed using different trust models and not using trust (20 repetitions). On the right, percentage of winner bids from unreliable bidders using different trust models and not using trust (20 repetitions). All trust values have been initialized to 0.5.

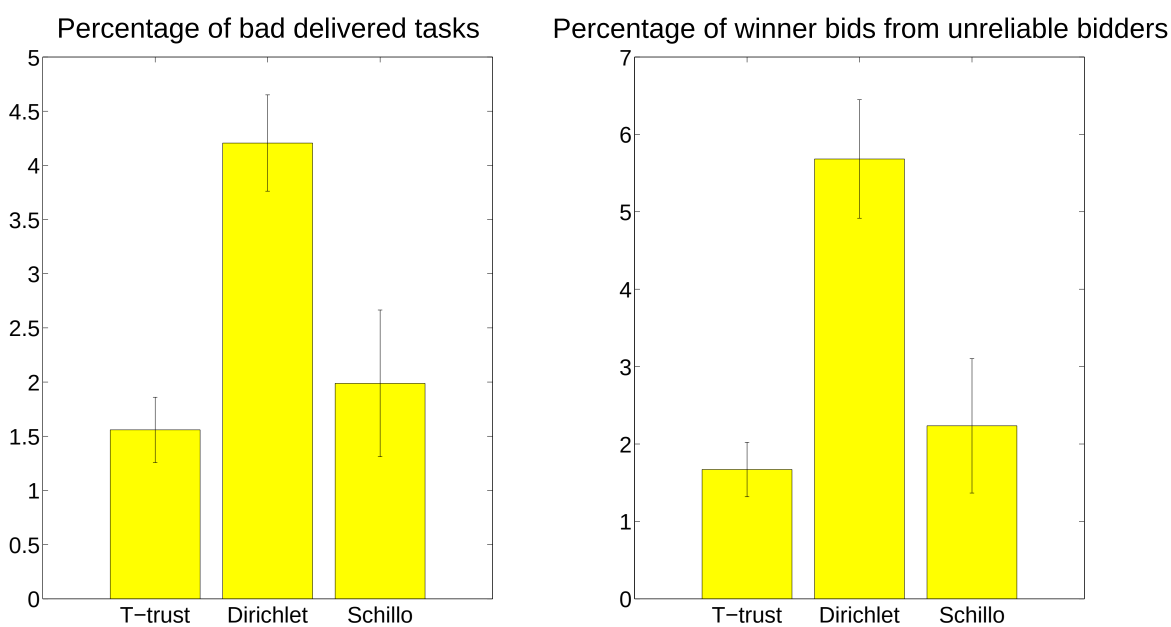

Figure 2: On the left, percentage of tasks delivered in worse conditions than the agreed using different trust models (20 repetitions). On the right, percentage of winner bids from unreliable bidders using different trust models (20 repetitions). All trust values have been initialized to 0.5 but the Schillo and Dirichlet mechanism have an initial memory of 10 values (half of them good) for each bidder.

from unreliable bidders (inaccurate bidders), and therefore, the number of tasks badly delivered. These results show the improvement respect the previous work of (Pla et al. 2014). However, the improvement depends on the trust model.

Results: Trust models comparison

According to Figure 1, the best results are obtained using the Schillo model. For example, Schillo model obtains the best results according to Figure 1 because its simple model is able to quickly discriminate between reliable and unreliable bidders. On the other hand, T-trust and Dirichlet model obtain worse results because they are slower. Dirichlet model needs more information to make up the probability distribution function because it considers more states (bad, neutral and good). T-trust answer time depends on α_t and β_t and the values used make it slower than the Schillo model without memory.

An important issue of the Dirichlet and Schillo models is that they use all the past information without emphasizing the most recent. This conveys a problem of rigidity when agents change their behavior. To tackle this problem Schillo and Dirichlet approaches can use a memory parameter that will determine how many of the last auctions should be considered to compute the trust. Conversely, the trust model we propose does not need such parameter as it automatically gives more relevance to the most recent auctions. Figure 2 shows the same information of Figure 1 regarding the different trust models analyzed in this paper, but Schillo and Dirichlet models have been initialized with a memory of 10 values from each bidder, which five were *good* delivered tasks and the other were *bad* (*very bad* for Dirichlet) delivered tasks. Note that this models a change of behavior of the agents respect their last 10 actions. Our model has been also initialized with a value of 0.5 for each bidder and 0.1 for the (α_t and β_t values). Regarding Schillo and Dirichlet models, the results obtained with these initializations are worse than the results of Figure 1. This proves the drawback these models have respect the approach presented in this paper. Comparing the Schillo model with the Dirichlet model, we can say that Schillo model outperforms again Dirichlet model because it needs less instances to re-shape the probability distribution function of each agent.

We have also repeated the experiment with all trust values initialized to 1.0 and Schillo and Dirichlet models with an initial memory of 10 good delivered tasks for each bidder. The best results are obtained by our approach confirming it is more robust to bad initializations and changes in agents' reliability. Comparing Figures 1 and 2 we see that the three models are sensible to the initialization values. That fact was expected because they are based on past experience and, therefore, if the initialization values does not correspond to the behavior of the agents, the performance will be worse. However, the important point is the flexibility of the models.

Summing up we can say that including trust in the bid valuation according to the methodology presented in this paper highly reduces the percentage of bad delivered items because it reduces the chances any unreliable has bidder of winning an auction. Nevertheless, the results are strongly linked to the model of trust and its initialization. In this sense, we see that Schillo model is the one that obtains better results when agents have a constant reliability. However, when the reliability of the agents in not constant, the performance of this model, as well as Dirichlet model, drops compared to the trust model we propose. This problem can be solved adding to the model a sliding window or weighting the past values according to the time, but this adds a complexity to these simple and easy models. On the other hand the trust model presented in this paper becomes a simple and robust solution against changes in the reliability of the agents.

Conclusion

This paper presents a new perspective of trust from a multi-attribute mechanism, to deal with task allocation problems that require several attributes to be fulfilled in addition to task prices (task delivery time, energy consumption). Particularly, the approach consists of defining trust value for any attribute provided by the bidder. Trust, as any other attribute, is included in the valuation function of the bids, and thus, it implicitly includes trust in the winner determination problem and in the payment rule. According to the experiments explained in this paper, the use of trust using this methodology helps to reduce the number of winner bids from unreliable bidders and, therefore, the number of tasks executed in worse conditions than the agreed.

In addition to the methodology to include trust in multi-attribute auctions, this paper proposes a new trust adaptation method which consists of increasing or decreasing the trust value (depending on whether the task is executed properly or not) according to a simple mathematical function with asymptotes on 0 and 1. This model does not present the rigidity problem present in other models of the literature when it comes to agents that have inconstant performances (e.g. agents starts acting oppositely to how they were acting before).

As a future work, it should be studied how multi-faced trust can handle the cold start problem in multi-attribute auction configurations. In addition, it could be interesting to transfer the approach presented in this paper towards a combinatorial multi-attribute auction configuration; nevertheless the complexity of combinatorial auctions makes it a challenging problem.

Acknowledgment

This research project has been partially funded through BR-UdG Scholarship of the University of Girona granted to Ferran Torrent-Fontbona. Work developed with the support of the research group SITES awarded with distinction by the *Generalitat de Catalunya* (SGR 2014-2016) and the MESC project funded by the Spanish MINECO (Ref. DPI2013-47450-C2-1-R).

References

[Chiu, Huang, and Yen 2010] Chiu, C.; Huang, H.; and Yen, C. 2010. Antecedents of trust in online auctions. *Electronic Commerce Research and Applications* 9(2):148–159.

[Hu et al. 2013] Hu, J.; Ning, Z.; Lu, X.; and Hao, S. 2013. A Trust Model for Effective Selection of Sellers Based on Social Network. In *Proceedings of the 2nd International Workshop on Incentives and Trust in E-Commerce, IJCAI WIT-EC*.

[Jø sang, Ismail, and Boyd 2007] Jø sang, A.; Ismail, R.; and Boyd, C. 2007. A survey of trust and reputation systems for online service provision. *Decision support systems* 43(2):618–644.

[Jurca and Faltings 2003] Jurca, R., and Faltings, B. 2003. Towards incentive-compatible reputation management. In *In Proceedings of the AAMAS 2002 Workshop on Deception, Fraud and Trust in Agent Societies*, 92–100.

[Lee and Szymanski 2005] Lee, J.-S., and Szymanski, B. K. 2005. A novel auction mechanism for selling time-sensitive e-services. In *E-Commerce Technology, 2005. CEC 2005. Seventh IEEE International Conference on*, 75–82.

[MacKie-Mason and Varian 1994] MacKie-Mason, J. K., and Varian, H. R. 1994. Generalized vickrey auctions.

[Marsh 1994] Marsh, S. 1994. *Formalising trust as a computational concept*. Ph.D. Dissertation, University of Stirling.

[Maskin and Riley 2003] Maskin, E., and Riley, J. 2003. Uniqueness of equilibrium in sealed high-bid auctions. *Games and Economic Behavior* 45(2):395–409.

[Pinyol and Sabater-Mir 2013] Pinyol, I., and Sabater-Mir, J. 2013. Computational trust and reputation models for open multi-agent systems: a review. *Artificial Intelligence Review* 40(1):1–25.

[Pla et al. 2012] Pla, A.; Gay, P.; Meléndez, J.; and López, B. 2012. Petri net-based process monitoring: a workflow management system for process modelling and monitoring. *Journal of Intelligent Manufacturing* 25(3):539–554.

[Pla et al. 2014] Pla, A.; López, B.; Murillo, J.; and Maudet, N. 2014. Multi-attribute auctions with different types of attributes: Enacting properties in multi-attribute auctions. *Expert Systems with Applications* 41(10):4829–4843.

[Porter et al. 2008] Porter, R.; Ronen, A.; Shoham, Y.; and Tennenholtz, M. 2008. Fault tolerant mechanism design. *Artificial Intelligence* 172(15):1783–1799.

[Ramchurn and Mezzetti 2009] Ramchurn, S., and Mezzetti, C. 2009. Trust-based mechanisms for robust and efficient task allocation in the presence of execution uncertainty. *Journal of Artificial Intelligence Research* 35(1):119.

[Regan and Cohen 2005] Regan, K., and Cohen, R. 2005. A model of indirect reputation assessment for adaptive buying agents in electronic markets. *Proceedings of the Business Agents and Semantic Web (BASeWEB05)*.

[Schillo, Funk, and Rovatsos 1999] Schillo, M.; Funk, P.; and Rovatsos, M. 1999. Who can you trust: Dealing with deception. *Proceedings of the workshop Deception, Fraud and trust in Agent Societies at the Autonomous Agents Conference*.

[Schillo, Funk, and Rovatsos 2000] Schillo, M.; Funk, P.; and Rovatsos, M. 2000. Using trust for detecting deceitful agents in artificial societies. *Applied Artificial Intelligence* 14(8):825–848.