# Continuous Double Auctions with Execution Uncertainty

**Gert van Valkenhoef***
valkenhoef@ai.rug.nl
Artificial Intelligence
University of Groningen
The Netherlands

**Sarvapali D. Ramchurn**
sdr@ecs.soton.ac.uk
Electronics and Computer Science
University of Southampton
United Kingdom

**Perukrishnen Vytelingum**
pv@ecs.soton.ac.uk
Electronics and Computer Science
University of Southampton
United Kingdom

**Nicholas R. Jennings**
nrj@ecs.soton.ac.uk
Electronics and Computer Science
University of Southampton
United Kingdom

**Rineke Verbrugge**
rineke@ai.rug.nl
Artificial Intelligence
University of Groningen
The Netherlands

**Abstract**

We propose a novel variant of the *Continuous Double Auction* (CDA), the Trust-based CDA (T-CDA), which we demonstrate to be robust to *execution uncertainty*. This is desirable in a setting where traders may fail to deliver the goods, services or payments they have promised. Specifically, the T-CDA provides a mechanism that allows agents to commit to trades they believe will maximize their expected utility. In this paper, we consider agents that use their *trust* in other agents to estimate the expected utility of a transaction. We empirically evaluate the mechanism, both against the optimal solution given perfect and complete information and against the standard CDA. We show that the T-CDA consistently outperforms the traditional CDA as execution uncertainty increases in the system. Furthermore, we investigate the robustness of the mechanism to unreliable trust information and find that performance degrades gracefully as information quality decreases.

## Introduction

Resource allocation is an important problem in computer science. Traditionally, it has been studied in settings where computational entities are cooperative and the allocation is determined by a central authority (e.g., the operating system kernel allocating available CPU time to different processes). However, with the advent of Grid computing, peer-to-peer systems, and ad-hoc networks, distributed systems are now being populated by an increasingly large number of computational entities. In such circumstances, a fully centralized approach to resource allocation may not be feasible, as the central resource broker will become a bottleneck for system performance (Wolski et al. 2003). Furthermore, such settings are not necessarily cooperative: stakeholders may have conflicting interests and may be motivated by their individual profit. Therefore, an approach that acknowledges the autonomy of the different actors within a multi-agent system is required.

In more detail, if we consider a truly open infrastructure, there may be a very large number of agents providing a certain resource and a large number of agents that need such a resource. For a number of reasons, some agents may be more reliable (i.e., more likely to provide the goods as promised, or to settle the payment) than others. For example, a desktop computer providing its idle CPU time will typically be less reliable than a dedicated machine in a data center with equivalent CPU power. The situation is complicated further by the fact that agents may enter or leave the system at any time.

In these systems, there may be a limited demand or supply of specific resources. Now, in a setting where agents compete for a limited demand and supply of a resource, market-based resource allocation mechanisms are a natural choice (Clearwater 1995), since they are designed such that desirable overall system behavior emerges despite the agents' selfish, profit-motivated behaviors. Specifically, in adopting a decentralized approach, we also want to ensure a good system-wide allocation of resources, i.e., we aim to maximize *social welfare*. Given that we need to deal with large numbers of agents (buyers and sellers in the market-based paradigm), who may enter or leave the market at any time and who want to be able to trade at any time, the Continuous Double Auction (CDA) is an appropriate choice (Dash et al. 2007). Indeed, the CDA has already been shown to be a highly efficient mechanism for such cases (Gode and Sunder 1993).

However, the CDA does not inherently deal with the varying degrees of reliability that may be exhibited by the different agents. This means that a buyer will always choose a low priced offer that is almost certainly faulty, over a reliable offer that is priced slightly higher. Indeed, one of the key roles of CDA-based exchanges is to organize trading so that contract defaults (i.e., agents opting out of the contract due to non-execution or non-payment) are avoided (Hull 2006). They do so by having investors deposit funds in an account that is adjusted at the end of each trading day to cover their demand and supply in the market in case of defaults. However, this can often be a *barrier to entry* in the market for new investors that lack the funds to join the market. Moreover, this requires a trusted third party to be established to provide these guarantees.

Thus, to address these shortcomings, we propose a novel mechanism in which agents themselves manage the risk of

*This paper presents work done while the author was at the University of Southampton.

defaulting. This implies that the market can be set up in an ad-hoc manner and does not require traders to trust a specific agent (typically the auctioneer) to manage risks for them. To this end, the Trust-based CDA (T-CDA), allows agents to use a *trust model* in their decision making process to assess whether to accept or reject offers based on cost and the *perceived* reliability of the proposer. Specifically, each agent has a private trust function that represents its best estimate of the reliability of each potential trading partner. The T-CDA decouples the commitment that is inherent in making an offer in the CDA. Thus, trade is separated into a bidding and a commitment phase, further decentralizing the decision making. Specifically, this circumvents the need for a trusted third party and agents *do not* need to reveal their trust information. Note that this is different from a multi-attribute auction where the trader's reputation is part of the bid, as in that case the traders would have to agree on a trusted source of reputation information, or trust each other to report this value truthfully. The mechanism proposed in this paper may be viewed as the decentralized version of a (centralized) combinatorial auction where agents submit their valuation for each potential transaction partner to the auctioneer (however, we *do not* require agents to reveal this information).

We empirically demonstrate our mechanism to be at least as efficient, and *usually more efficient*, in terms of maximizing social welfare as the CDA. Moreover, agents in the T-CDA never trade at negative expected utility, as can happen to their CDA counterparts. Finally, we empirically evaluate the T-CDA's robustness against unreliable trust models and observe that performance decreases linearly with the error introduced in the trust model.

In this paper, we advance the state of the art in the following ways:

1. We develop a decentralized market-based mechanism that is robust to execution uncertainty.
2. Although this mechanism allows agents to use a trust model in their decision making, it does not require agents to reveal their trust function.
3. We separate the offer and commitment phases inherent in traditional CDA trade, providing a new way of managing risk that can replace or complement the guarantee funds that are conventionally used.

The paper is structured as follows: first we review related work, then the problem is formalized. Againts this background, we give some desiderata of a mechanism for such a setting. Subsequently, the new mechanism is detailed and the agent behavior is defined. We then turn to an empirical evaluation and, finally, we conclude and state directions for further work. Table 2 summarizes the notations introduced in this paper.

## Related Work

Early investigation of the CDA was based on the market protocol discussed by Smith (1962). For example, Gode and Sunder showed that most of the efficiency achieved by trade in the CDA is due to the mechanism itself and not due to the intelligent behaviors of the traders, by showing that the efficiency with a *Zero Intelligence* (ZI) strategy (where agents bid randomly) was close to that of human traders (Gode and Sunder 1993). In particular, a ZI-C (ZI Constrained) agent randomly shouts any price from the range of prices that will not result in a loss for the agent. Later, Cliff and Bruten (1997) showed that there are significant differences between the behavior of ZI-C traders and humans. Specifically, they point out that the high efficiency shown by Gode and Sunder is due to the specific demand and supply curves used and that ZI-C traders are much more erratic than human traders. To remedy this, a more intelligent strategy was developed, the *Zero Intelligence Plus* strategy (Cliff and Bruten 1997): a minimally intelligent strategy with human-like trader behavior.

A number of other lines of work have extended the CDA in some domain-dependent way. In particular, Dash *et al.* adapt the CDA for a scenario where sellers have a limited capacity and a complex cost structure, defined by a fixed overhead cost and a constant marginal cost (Dash et al. 2007). In that work, the extension of the CDA is empirically compared to a centralized mechanism that is known to find the optimal allocation, an approach we also adopt in this paper. Finally, other work has proposed similar market-based approaches to computational resource allocation (Buyya, Abramson, and Venugopal 2005; Wolski et al. 2001; 2003) and architectures to put this in practice have been designed (Buyya, Abramson, and Giddy 2000). Indeed, specific approaches based on the CDA have also been proposed, e.g., (Tan and Gurd 2007).

As can be seen, there is a precedent for adaptation of the CDA to solve new problems. In addition to this, there is also past work on the integration of trust in mechanism design (Dash, Ramchurn, and Jennings 2004) and, specifically, auctions (Porter et al. 2008). However, to date, there is no work that attempts to extend the CDA to a domain where the expected utility of a transaction is dependent on the reliability of the other party and where the reliability of traders may vary greatly. This is because originally the CDA was intended for commodity markets, where these issues do not arise. Moreover, the application of the CDA to computational resource allocation has focused on cases where the services being traded can reasonably be treated as commodities, which is clearly not the case in an open distributed system, where failure to deliver what was promised is a real possibility. Therefore, a variant of the CDA is required in our context. Thus, we propose a mechanism that allows traders to differentiate between potential transaction partners based on their expected reliability (i.e., trust).

## Modeling the Trading Environment

We denote the set of buyers as $b_1, b_2, \ldots, b_n \in B$ and the set of sellers as $s_{n+1}, s_{n+2}, \ldots, s_{n+m} \in S$. Then, the set of agents is denoted as $A = B \cup S$. As a convention, we generally refer to a generic buyer as $b_i$, a seller as $s_j$ and an agent that can be of either type as $a_k$.

Every agent participating in the market is given an *endowment*. For a buyer, an endowment is an order to buy a single unit of resource for at most the specified *limit price*, $\ell_i^b$. For a seller, an endowment is an order to sell a single unit of resource for at least the specified *cost price*, $\ell_j^s$.

Given their endowments, buyers place *bids* (offers to buy) and sellers place *asks* (offers to sell) in the market.[1] Based on the submitted bids and asks, the market mechanism determines when a *transaction* takes place between a buyer and a seller. We will denote a transaction at price $q$ between a buyer $b_i \in B$ and seller $s_j \in S$ as $t_{i,j}(q)$. After agreeing on a transaction $t_{i,j}(q)$, the buyer pays the seller and the seller transfers some goods to the buyer. The way the shouts are managed in the market can be regimented by different market rules.

The setting described above is the one traditionally considered in market-based mechanisms. Moreover, in this work, we do not assume that successful *execution* of a transaction is guaranteed. Instead, we assume that the execution of a transaction is binary, that is, either failure or success.[2] We denote the outcome for the buyer as $e_i^b \in \{0,1\}$ and for the seller as $e_j^s \in \{0,1\}$. The probability that a buyer is successful (i.e. $P(e_i^b = 1)$) is denoted as $p(b_i)$ and that the seller is successful (i.e. $P(e_j^s = 1)$) as $p(s_j)$. For example, after $t_{i,j}(q)$, if $e_i^b = 1$ and $e_j^s = 0$, buyer $b_i$ has paid for a service, but $s_j$ did not provide that service. In general, every agent $a_i$ is assigned a certain *probability of success* (POS) $p(a_i) \in [0,1]$, which indicates the likelihood that an agent will honour its agreement.

Given the execution of a transaction, the agents derive utility as follows:

$$\begin{aligned} u_i^b(t_{i,j}(q), e_j^s) &= \begin{cases} \ell_i^b - q & , e_j^s = 1 \\ -q & , e_j^s = 0 \end{cases} , \\ u_j^s(t_{i,j}(q), e_i^b) &= \begin{cases} q - \ell_j^s & , e_i^b = 1 \\ -\ell_j^s & , e_i^b = 0 \end{cases} , \end{aligned} \tag{1}$$

where $\ell_i^b$ is the limit price of $b_i$ (i.e., the maximum $b_i$ is willing to pay) and $\ell_j^s$ is the cost price of $s_j$ (i.e., the minimum price at which $s_j$ is willing to sell). These functions follow naturally if we assume that agents are not malicious; i.e., regardless of their own success and regardless of their partner's success, they will incur the cost associated with the action they agreed to perform.[3] Although this definition of utility is not necessarily appropriate in every setting, it was chosen to represent a worst case scenario: if agents derive non-negative utility in this scenario, they will certainly do so in a more forgiving scenario. Given this, the expected utility of a transaction $t = t_{i,j}(q)$ is given by:

$$\begin{aligned} \bar{u}_i^b(t) &= u_i^b(t,1)p(s_j) + u_i^b(t,0)(1 - p(s_j)) \\ &= \ell_i^b p(s_j) - q \\ \bar{u}_j^s(t) &= u_j^s(t,1)p(b_i) + u_j^s(t,0)(1 - p(b_i)) \\ &= q p(b_i) - \ell_j^s \ . \end{aligned} \tag{2}$$

In order to make informed decisions, an agent needs to evaluate the utility it expects to derive from each of the possible transactions. Since, in general, we cannot assume that agents have perfect and complete knowledge of each other's POS, agents hold an estimate of the POS of the other agents. Thus, each agent $a_i$ has a *trust function* ($\mathrm{trust}_i : A \to [0,1]$), which represents its best estimate of the probability of success for each other agent. So ideally, $\mathrm{trust}_i(a_j) = p(a_j)$. This allows $a_i$ to estimate the expected utility $\bar{u}$ (Equation 2) of a transaction:

$$\begin{aligned} \tilde{u}_i^b(t_{i,j}(q)) = \ & u_i^b(t_{i,j}(q),1)\mathrm{trust}_i(a_j) + \\ & u_i^b(t_{i,j}(q),0)(1 - \mathrm{trust}_i(a_j)) \\ \tilde{u}_j^s(t_{i,j}(q)) = \ & u_j^s(t_{i,j}(q),1)\mathrm{trust}_j(a_i) + \\ & u_j^s(t_{i,j}(q),0)(1 - \mathrm{trust}_j(a_i)) \ . \end{aligned} \tag{3}$$

Now, it is rational to agree to a transaction only if the estimated expected utility $\tilde{u}_i(t) \geq 0$. In this paper we remain agnostic to the origin of this trust function; agents might learn the reliability of others through the observation of market interactions, or they could employ some outside source of information. Rather than implementing one of these approaches, we simulate the trust model by endowing agents with trust information that has certain properties (see 'Empirical Evaluation').

Note that our model is equivalent to the setting in which the CDA is normally evaluated, when

$$\mathrm{trust}_i(a_j) = p(a_j) = 1;\ \forall a_i, a_j \in A \ .$$

## Mechanism Desiderata

Given our problem setting, we define a number of *desiderata* that we believe our mechanism should exhibit. In particular, the market mechanism should be *efficient*: it should maximize the sum of the expected utilities of the individual agents, since we want to maximize social welfare. It should also be *individually rational*, i.e., individual agents will not participate in loss-making transactions. This ensures that we do not disincentivize agents from participating in our market. Furthermore, an equal and, thus, fair distribution of profits between buyers and sellers is desirable (again to ensure we have approximately equal numbers of each). Additionally, since our model incorporates the notion of POS, we desire the mechanism to be robust against agents having an inaccurate representation of each others' POS, since in the real world, it is unrealistic to assume that agents have perfect and complete information about the reliability of other agents.

In order to evaluate the *efficiency* of the mechanism, we need to define and find the optimal solution, given complete and perfect information of all agents. This provides an upper bound on the efficiency we can expect from our mechanism. Given our model, we aim to find the allocation that maximizes the sum of the expected utilities of the individual agents, subject to certain constraints.

First, let us consider how to choose the transaction price given that two agents interact. In order to optimize efficiency, we should maximize the sum of the agents' individual utilities:

$$\begin{aligned} U_{i,j}(q) &= \bar{u}_i^b(t_{i,j}(q)) + \bar{u}_j^s(t_{i,j}(q)) \\ &= \ell_i^b p(s_j) - \ell_j^s + q(p(b_i) - 1) \ . \end{aligned} \tag{4}$$

[1] Collectively, bids and asks are referred to as *shouts*.

[2] Failure is binary to simplify our analysis, but this work can easily be generalised to be continuous, to reflect partial success or failure if that is appropriate in a given setting.

[3] For example, when a buyer pays $q$ and if he receives the goods (or service), which are worth $\ell_i^b$ to him, he will derive a utility of $\ell_i^b - q$. Otherwise, his utility is $-q$.

From the above formula, we see that when the probability of success of the buyer $p(b_i) = 1$, the transaction price $q$ has no influence on the total expected utility of the transaction. However, when $p(b_i) < 1$, a higher transaction price leads to a lower expected utility. Therefore, if we choose $q$ to optimize $U_{i,j}$, sellers will derive negative expected utility. Hence, participation is not individually rational.

To remedy this, we could demand that $\bar{u}^s_j(t_{i,j}(q)) \geq 0$, however when $p(b_i) < 1$, the result will be that sellers will always break even and thus have no incentive to take part in the market. Instead, we demand that the expected utilities of both parties are equal, to achieve a fair distribution of utility between buyers and sellers:

$$\bar{u}^b_i(t_{i,j}(q)) = \bar{u}^s_j(t_{i,j}(q)) \ . \tag{5}$$

This constraint determines a unique solution for the transaction price $q$. Then, given that in our current model, each agent can take part in only one transaction, we can find the set of pairs $T = \{(b_i, s_j) \mid b_i \in B \wedge s_j \in S\}$ that maximizes $U = \sum_{(b_i,s_j)\in T} U_{i,j}$.

Here $U$ gives an upper bound on the performance of the T-CDA, under the constraint that utility is equally distributed between buyers and sellers. However, the solutions the T-CDA finds do not necessarily obey this constraint, because it cannot enforce it since the solution is not centrally determined. Therefore, in evaluating the mechanism, we must seperately compare both buyer and seller utilities to $0.5U$.

## Designing the Trust-Based CDA

As we pointed out earlier, traditional market mechanisms ignore the execution phase present in every interaction. In traditional settings, this is justified because the implications of execution uncertainty can be dealt with outside the scope of the market. However, the ad-hoc nature of the markets considered here makes this unacceptable. Given this, here we first describe the CDA and then propose our extension, the Trust-Based CDA, that does allow agents to factor the execution phase into their decision making.

In more detail, the market protocol that defines the CDA consists of a number of simple rules. In order to keep track of the offers that have been made, bids and asks are queued into *order books*, which are sorted lists of orders. Bids are sorted from highest to lowest, asks from lowest to highest. The following rules define the CDA protocol in detail:

**Shout Accepting Rule** Determines which bids and asks are allowed in the market. Primarily, the price must be within the interval $[0, \mathrm{maxprice}]$[4]. Furthermore, the commonly implemented NYSE shout accepting rule imposes that a new shout must improve upon the current best shout by that agent. When a trader submits a new shout, provided that it improves upon the current shout by that trader, the current shout is simply replaced by the new one.

**Information Revelation Rule** Determines what information is published to buyers and sellers. Here, we assume this to be the current bid and ask prices.

[4] $\mathrm{maxprice}$ is an arbitrary price limit set by the market.

**Clearing Rule** The market *clears* continuously, whenever the highest bid price is at least as high as the lowest ask. Then a transaction takes place, at a *transaction price*, determined according to the pricing rule. The matched shouts are removed from the order books.

**Pricing Rule** Determines the transaction price. The average of the matched bid and matched ask prices is often used in the CDA, and will be used here.

The CDA may be seen as consisting of two components. First, the *bidding component* manages the agents' interaction with the order books, through the shout accepting rule. Second, the *clearing component* determines how transactions arise, through the clearing and pricing rules.

Now, in our setting, the CDA is modified to additionally let agents accept or reject transactions based on the identity of the other agent. To this end, agents not only submit their bids or asks to the market, but also have to explicitly indicate their willingness to interact with a specific agent before a transaction takes place. We call this declaration of willingness a *commitment*. This allows us to leave most of the rules and structure of the CDA intact and also maintains the decentralized nature of the CDA, by leaving the management of trust information and the decision making up to the agents themselves. Indeed, our mechanism does not require agents to reveal this information. As in the CDA, the T-CDA merely provides the necessary means for the agents to communicate their desires effectively. Conversely, this means that agent strategies will be more complex and play an important role in determining individual agent utilities as well as system efficiency, as is the case for the CDA.

We may think of the mechanism as consisting of three components: the bidding and clearing components identified earlier and a new one, the *commitment component*, which manages the interaction with the commitment book, through the commitment accepting rule. This is shown in Figure 1.

In more detail, if $b_i \in B$ has placed a bid $o^b_i$ and $s_j \in S$ has placed an ask $o^s_j$, we denote the *commitment* of $b_i$ to a transaction based on $o^b_i$ and $o^s_j$ as $c_i(o^b_i, o^s_j)$. A commitment by $s_j$ would be $c_j(o^b_i, o^s_j)$. Two matching commitments result in a transaction. We do not allow more than one commitment by an agent on its own shout, since there can be only one transaction based on a particular shout. However, we do allow agents to withdraw a commitment, for example because the other agent is not responding. Agents may reject a commitment made by others on their shout.

In addition to the *order books*, the T-CDA has a *commitment book*, in which a list of all current commitments is maintained. We define an additional rule and adapt the Clearing Rule to deal with commitments:

**Commitment Accepting Rule** A commitment $c_k(o^b_i, o^s_j)$ is accepted when the prices of the shouts concerned match (i.e. $o^b_i \geq o^s_j$) and one of the shouts was made by the agent committing (i.e. $k = i \vee k = j$). Furthermore, any agent may have only one commitment for a specific shout in the commitment book at any one time. Commitments can be withdrawn by the agent that made them, or rejected by the agent that is being committed to. In either case, the commitment is removed from the commitment book.

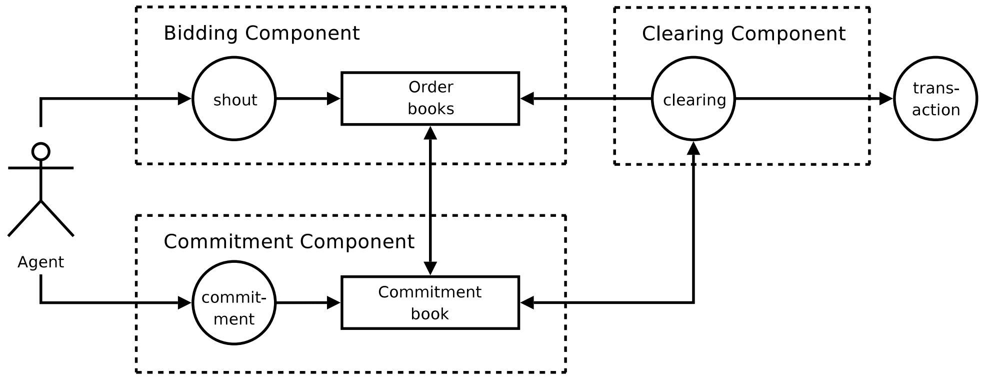


Figure 1: Information in the Trust-based CDA flows through three different components. The Commitment component distinguishes the T-CDA from the traditional CDA.

**Clearing Rule** Two commitments match when both the buyer and the seller commit. So, commitments $c_i(o_i^b, o_j^s)$ and $c_j(o_i^b, o_j^s)$ match and would result in a transaction $t_{i,j}(q)$, where $q$ is a transaction price determined by the *Pricing Rule*. After the matching, both the commitments and the shouts concerned are removed from the books.

To illustrate the trading process, consider a scenario with one buyer, $b_0$, with $p(b_0) = 1$ and $\ell_0^b = 8$ and one seller, $s_1$, with $p(s_1) = 0.85$ and $\ell_1^s = 5$. For simplicity, assume both agents have perfect knowledge of $p(\cdot)$. After some bidding, we have the offers $o_0^b = 7$ and $o_1^s = 6.8$ in the order books. Now, in the traditional CDA, the market would immediately clear and a transaction would take place at price $q = 6.9$. However, in the T-CDA, agents consider their expected utility (Equation 2) in order to decide whether to commit. It happens that $\bar{u}_1^s(t_{0,1}(6.9)) \geq 0$, so $s_1$ will commit to $c_1(o_0^b, o_1^s)$. However, $\bar{u}_0^b(t_{0,1}(6.9)) < 0$, so $b_0$ will reject the commitment, removing it from the commitment book. If $s_1$ were to improve its ask to $o_1^s = 6.4$, both agents have positive expected utility (at price $q = 6.7$) and they will both commit, resulting in a transaction $t_{0,1}(6.7)$.

## Behavior

In the traditional CDA, an agent's strategy is specified through its bidding behavior, which dictates the offers an agent submits in the market. In additon to this, a commitment behavior is also required when trading in the market, to determine when an agent commits.

A basic bidding strategy used in the CDA is the ZI-C behavior, which randomly picks a shout price from the range of acceptable prices (i.e., from $[0, \ell_i^b]$ for buyers and from $[\ell_j^s, \mathrm{maxprice}]$ for sellers). We augment this bidding strategy with a commitment behavior. This allows us to evaluate the structure of the mechanism, rather than agents' behavior, as per Gode and Sunder for the traditional CDA. The ZI-C strategy is easily extended to work in the T-CDA. For more advanced strategies, however, this is non-trivial. Thus, although evaluation with intelligent strategies is desirable, we leave this for further work.

The commitment strategy is based on a single heuristic: if the expected utility is non-negative, an agent $a_i$ is keen on transacting.[5] Hence, the following actions are tried in the order given:

1. Given commitments to its own shout, $a_i$ picks the best and commits if $\tilde{u}_i(t) \geq 0$;
2. If $a_i$ is already committed, it does nothing more;
3. Given compatible shouts, $a_i$ picks the best and commits if $\tilde{u}_i(t) \geq 0$;
4. $a_i$ submits an offer based on the ZI-C strategy.

If necessary, the agent $a_i$ will withdraw a previous commitment, while any unaccepted commitments on its own shout will be rejected.

## Empirical Evaluation

In this section we detail the empirical evaluation of the T-CDA. In particular we aim to see how it performs with respect to the desiderata specified above, under 'Mechanism Desiderata.' Specifically, we investigate:

- The efficiency of the mechanism;
- The distribution of utility between buyers and sellers;
- The robustness of the mechanism to errors in the trust information.

In doing so, we assume that:

- Agents are risk neutral (i.e. utility is a linear function of profits);
- Utility functions have no discount factor (i.e. agents value future reward just as much as current reward);
- Traders act according to the ZI-C strategy;
- The set of buyers and sellers is fixed;
- No new demand or supply appears during a run.

None of these assumptions are required by the mechanism, but they provide a simple scenario in which to evaluate the mechanism. In what follows, we first detail the experimental settings, the independent variables and the metrics used. Then, the experiments and results are discussed.

[5] This heuristic was chosen for its simplicity, but others may be equally appropriate.

**Experiment Settings.** For some variables, although they may impact on the performance of the mechanism in some way, the results in this paper are not sensitive to their specific values. Therefore, for these variables, reasonable default values were chosen. These values were identified by trial runs and they represent informative scenarios and reasonable performance (i.e., runs can be completed in acceptable time). More specifically, there are 50 buyer and 50 seller agents. The agents' endowments, which determine the orders the agents have to complete, are generated from a uniform distribution with the range $[6, 8]$ for sellers and $[10, 12]$ for buyers. Although it appears that all traders should transact (i.e., it seems that all traders are intra-marginal), this may not be the case, because not all traders may be matched with positive expected utility due to execution uncertainty. Thus, in most scenarios (defined below), there are extra-marginal traders in the market. The maximum price is set to 15. As agents do not learn over trading days (see 'Behavior'), a run will consist of a single trading day. Experiments consist of 300 runs per condition. The buyer POS is fixed at 1, because this allows for more insightful analysis, though similar results occur if failure is two-sided.

**Independent Variables.** There are three independent variables. The first two are the expected value $\mathrm{E(pos)}$ and variance $\mathrm{Var(pos)}$ of the probability of success of sellers. In total, 65 combinations of these variables are run. If $\mathrm{Var(pos)} = 0$, every seller has POS $\mathrm{E(pos)}$. Otherwise, POS values are drawn from a Beta distribution[6] with appropriately chosen parameters. The third variable determines the way in which trust (in sellers) is initialized for the buyers. If trust is CDA-LIKE, a trust of 1 is placed in every seller. This condition thus exhibits the same behavior as the traditional CDA. With RANDOM trust, trust values are drawn from a uniform distribution. Trust can also be initialized as the MEAN seller POS, or as a PERFECT copy of the POS value of each seller.

**Metrics.** Performance is measured as the sum of the actual (derived) utilities of all buyers, $V_B$, and the sum of the actual utilities of all sellers, $V_S$. When the optimal allocation has an expected utility $U \neq 0$, we may express these measures relative to the optimum, as $2V_B U^{-1}$ and $2V_S U^{-1}$, respectively.

Now, we analyze the performance of the mechanism, given that agents have a correct perception of their counterparts' probabilities of success. The analysis serves three main goals. First, it confirms that the emergent behavior of the system is as we expect. Second, we evaluate the behaviour of the mechanism, with respect to the optimal performance (as derived under 'Mechanism Desiderata') and to the traditional CDA. Finally, we evaluate the robustness of the mechanism to errors in the trust information.

[6]The Beta distribution was chosen because it generates values in $[0, 1]$ and allows us to choose the desired expected value $\mu$ and variance $\sigma^2$, by setting the parameters $\alpha$ and $\beta$. We may find the parameters $\alpha$ and $\beta$ from the desired $\mu$ and $\sigma^2$ as done in (Teacy 2006):

$$\alpha = \frac{\mu^2 - \mu^3}{\sigma^2} - \mu, \quad \beta = \frac{\alpha}{\mu} - \alpha \ .$$

The constraint that $\alpha > 0$ bounds $\sigma^2$ for specific $\mu$: $\sigma^2 < \mu - \mu^2$.

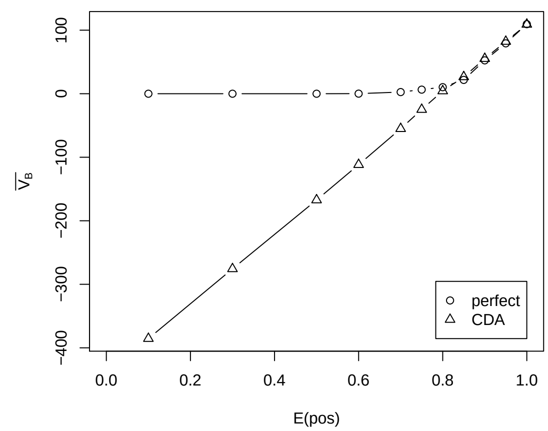


Figure 2: Buyer utility for the Trust-based CDA given PERFECT trust and the normal CDA. PERFECT trust avoids making a loss, where the CDA does make a loss. Confidence intervals are small.

**Positive Payoff** First of all, calculating the optimal allocation tells us when a positive payoff is possible. We expect that given perfect information, on average, the mechanism will derive a positive utility if that is at all possible.

**Hypothesis 1** *If for a certain setting of* $\mathrm{E(pos)}$ *and* $\mathrm{Var(pos)}$*, the optimal buyer utility is positive, so is the expected performance for the* PERFECT *trust setting.*

For the 60 out of 65 combinations of $\mathrm{E(pos)}$ and $\mathrm{Var(pos)}$ where the optimal expected utility is greater than zero, we do a t-test with the null hypothesis that the mean buyer utility is equal to zero. The alternative hypothesis is that the mean is greater than zero. At the $\alpha = 0.05$ level, we reject the null hypothesis in 56 of the 60 cases.[7]

In the cases where the null hypothesis is not rejected (and the mean buyer utility is thus roughly equal to zero), the estimated mean is greater than zero, so we need not consider the alternative that the actual mean is smaller than zero. Furthermore, these cases all have a very small optimal expected utility. Hence, in general, the mechanism does derive a positive expected utility if this is possible.

**Comparison to CDA** In this experiment we want to show that not only do we avoid making a loss and that we turn a profit whenever possible, we also do better than we would if we would ignore trust information altogether, as in the CDA-LIKE and the RANDOM conditions.

[7]If we protect the null hypothesis against spurious results by setting $\alpha' = 1 - 0.95^{1/65}$, the null hypothesis is rejected in 54 cases.

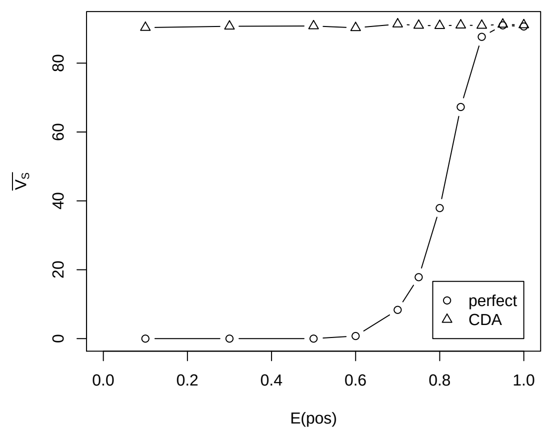


Figure 3: Seller utility for the Trust-based CDA given PERFECT trust and the normal CDA. The normal CDA allows unreliable sellers to exploit buyers. Confidence intervals are small.

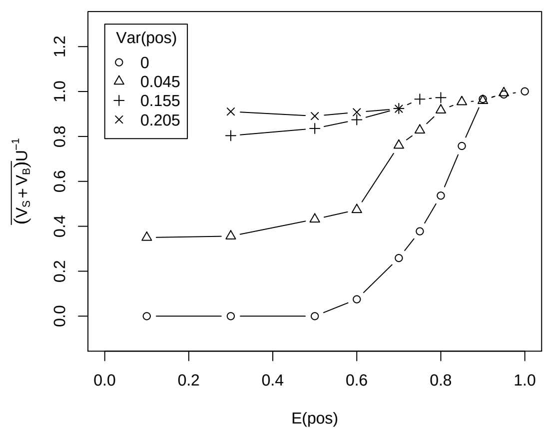


Figure 4: The normalized utility, or *efficiency* derived by the mechanism. Confidence intervals are small.

To this end, Figure 2 shows a typical outcome when $\mathrm{Var(pos)} = 0$, for different levels of $\mathrm{E(pos)}$. The PERFECT condition does not trade for low values of $\mathrm{E(pos)}$, where a profit is not possible. For higher values of $\mathrm{E(pos)}$, the utility for the PERFECT condition increases more or less linearly. For the CDA-LIKE condition, the relationship between $\mathrm{E(pos)}$ and buyer utility is linear, which is what we expect, since it will ignore the probability of success of sellers altogether. Hence, it derives a (very large) negative utility for low $\mathrm{E(pos)}$. Beyond a certain threshold, there is very little distinction between the PERFECT and CDA-LIKE conditions. This is to be expected, since then transactions are usually desirable and given random (ZI) bidding both conditions will lead to approximately identical results.

In Figure 3 we see that the influence of $\mathrm{E(pos)}$ on seller utility is quite different. Clearly, accurate trust information prevents the buyers from being exploited by sellers. We return to this point later, when we discuss Figure 5.

The above conceptions are formalized as follows:

**Hypothesis 2** *Under any setting of* $\mathrm{E(pos)}$ *and* $\mathrm{Var(pos)}$, PERFECT *trust will do at least as well (in terms of buyer utility) as the* RANDOM, CDA-LIKE *and* MEAN *conditions.*

To test this hypothesis, for all combinations of $\mathrm{E(pos)}$ and $\mathrm{Var(pos)}$, pairwise comparisons of the PERFECT condition were done against the other conditions. Two t tests were performed for each pair, in both cases the null hypothesis is that the means are equal. In the first test, the alternative is that the mean in the PERFECT condition is greater, in the second, that the mean is less. The resulting p-values were inspected at $\alpha' = 1 - 0.95^{1/65}$, protecting the null hypothesis (no difference) against spurious results.

Note that given the experiment settings, when $\mathrm{Var(pos)} = 0$ and $\mathrm{E(pos)} \geq 0.82$, the decisions made by CDA-LIKE trust are, on average, rational.[8] Hence, we cannot expect much advantage from good trust information in that case. In comparison to MEAN, we expect no difference when $\mathrm{Var(pos)} = 0$. Also see Figure 6, Table 1 and the corresponding discussion, which show that for $\mathrm{E(pos)} \geq 0.80$ and $\mathrm{Var(pos)} = 0$, errors in the trust information have very little impact on the overall system performance.

Looking at the 'PERFECT $>$ other' alternative hypothesis, at $\alpha'$, PERFECT is significantly better than RANDOM in 64 of the 65 cases, better than CDA-LIKE in 57 of the 65 cases and better than MEAN in 43 of the 65 cases. The cases of no difference correspond to the expectations mentioned above. For the 'PERFECT $<$ other' alternative, there are no significant differences at $\alpha'$.

Thus, it is safe to say that the PERFECT condition improves upon the control conditions RANDOM, CDA-LIKE and MEAN. Moreover, it is clear that the T-CDA does better than the CDA when faced with uncertainty about the result of transactions.

**Benchmark** In this experiment we benchmark the T-CDA's performance against the optimal performance and make some overall qualitative observations about its behaviour.

To this end, Figure 4 shows the total utility achieved by the system, normalized by the maximum expected utility from the optimal allocation. The mechanism does well when either $\mathrm{Var(pos)}$ is high, or $\mathrm{E(pos)}$ is high, or both. This is because, in both cases, the part of the population from which profit can be derived have $\mathrm{E(pos)} \approx 1$. Hence, when buyers bid randomly from $[0, \ell]$, they are submitting profitable bids. If, however, a large group from which profit may potentially be derived has a low POS, the bidding strategy does poorly. This is because it submits bids that are too high (overbid-

[8] Assume the transaction price is, on average, the equilibrium price $\bar{q} = 9$. Then, given the average limit price for buyers, $\bar{\ell}^b = 11$ and that all sellers have the same POS $p$, we can find $p$ such that expected buyer utility (Equation 2), on average, is non-negative: $\bar{u^b} = \bar{\ell}^b p - \bar{q} \geq 0 \Rightarrow p \geq \frac{9}{11} \approx 0.82$.

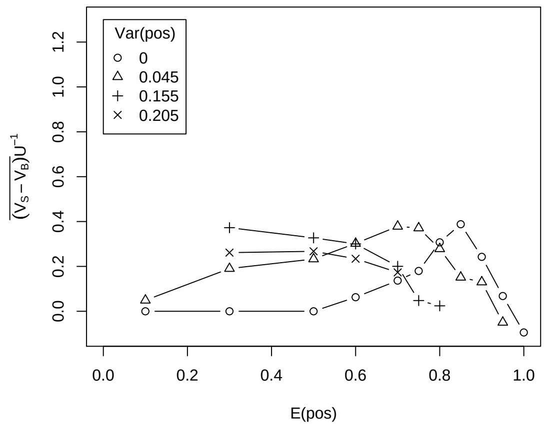


Figure 5: Disparity between seller and buyer utilities in the T-CDA. Confidence intervals are small.

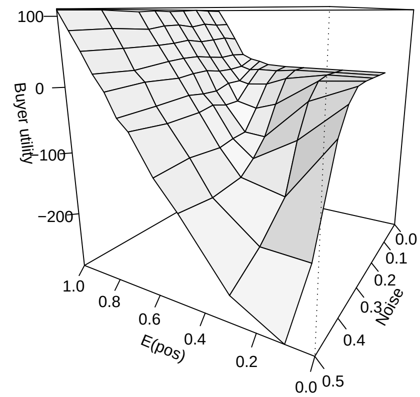


Figure 6: Performance degrades when $\mathrm{E(pos)}$ is lowered and when more noise is added to trust values. The noise level represents the variance of the Gaussian noise distribution that is applied to individual agents' trust function.

ding). Hence, the agent itself is not willing to transact at that price, given the execution uncertainty. Hence, the figure reveals the need for a bidding strategy to be informed by a trust model.

Another relevant aspect of the behaviour is the balance of utility between buyers and sellers. This is shown in Figure 5. In the $\mathrm{Var(pos)} = 0$ case, it appears that sellers are the first to profit from an increase in $\mathrm{E(pos)}$, with the balance being restored only for the highest values of $\mathrm{E(pos)}$. Specifically, for $\mathrm{E(pos)} = 0.60$, observe that the difference of seller and buyer utility is almost identical to the total utility in the system, i.e., only the sellers turn a significant profit. The higher $\mathrm{Var(pos)}$ levels show an imbalance that decreases when $\mathrm{E(pos)}$ increases. Once again, the imbalance is caused by the bidding strategy, which is uninformed about the actual worth of the sellers' offers.

**Effect of Unreliable Information** Now, because we cannot assume agents to have perfect and complete information of each others' POS, we analyse the effect of the degradation of trust information on the mechanism. To simulate unreliable trust information, each buyer's trust function is initialized to the actual POS values with some arbitrary level of Gaussian noise applied to it. Figure 6 provides an overview of the results.

The figure provides a number of interesting insights. First, if the noise level is high, performance degrades almost linearly with $\mathrm{E(pos)}$. This is to be expected, since interaction partners are chosen almost completely at random, and this randomness leads to a linear relationship between buyer utility and $\mathrm{E(pos)}$. Second, if $\mathrm{E(pos)}$ is very low, performance increases linearly with a decreasing noise level, until a 'plateau' is reached where utility is zero. A linear regression (Table 1) shows that a linear relation can indeed account for a large proportion of the variance in these cases. Adding noise means that agents will overestimate POS in some cases and hence that they may transact even if it is not in their best interest, leading to losses. The 'plateau' where utility is zero exists because even with some overestimation of the POS, agents do not see transactions as desirable.

| $\mathrm{E(pos)}$ | *noise* | $r^2$ | $F$ | $p$ |
|---|---|---|---|---|
| 0.10 | $[0.15, 0.50]$ | 0.74 | 280 | $\ll 0.01$ |
| 0.30 | $[0.10, 0.50]$ | 0.70 | 226 | $\ll 0.01$ |
| 0.50 | $[0.05, 0.25]$ | 0.61 | 155 | $\ll 0.01$ |
| 0.80 | $[0.00, 0.50]$ | 0.00 | 0.15 | $> 0.50$ |
| 0.85 | $[0.00, 0.50]$ | 0.00 | 0.04 | $> 0.50$ |
| 0.90 | $[0.00, 0.50]$ | 0.00 | 0.01 | $> 0.50$ |
| 0.95 | $[0.00, 0.50]$ | 0.00 | 0.01 | $> 0.50$ |
| 1.00 | $[0.00, 0.50]$ | 0.00 | 0.00 | $> 0.50$ |

Table 1: Linear regression of buyer utility on noise, for $\mathrm{Var(pos)} = 0$, significance tested against $F$ distribution. $r^2$ is the proportion of the total variance accounted for by the regression line. $F$ is the value of the F-test statistic for a linear regression and $p$ is the significance of the regression line given by an $F$ distribution.

Last, when $\mathrm{E(pos)} \geq 0.80$, the noise level seems to have very little impact on the total utility derived by buyers, rather increasing linearly with an increasing $\mathrm{E(pos)}$. Linear regression of buyer utility on noise (Table 1) confirms this. This may be explained by the fact that in Figure 6, $\mathrm{Var(pos)} = 0$ and hence there is no benefit in distinguishing between sellers. The intuition behind this is that the application of noise introduces an arbitrary preference for certain sellers, which is different for each buyer, and transactions are usually desirable. Thus, the effects of noise on the individual cancel out over the entire population.

## Conclusions

In this paper, we propose a novel resource allocation mechanism based on the CDA, that allows the varying degrees of reliability of trading agents to be taken into account in the decision making process. We empirically demonstrate the efficiency of our mechanism and, specifically, its robustness against increasing execution uncertainty. The standard CDA mechanism, on the other hand, is shown to break down, with agents ending up with considerable losses in such a setting. Moreover, we show that our mechanism is robust to errors in the trust information employed by agents. Specifically, performance degrades linearly with the information error. We believe that our approach is a significant step for more realistic and uncertain environments in which execution cannot be guaranteed.

For future work, we first will extend our model to cover settings where agents can enter and leave the system at any time, where there are more dynamic market shocks (i.e. drastic changes in demand and supply) and where there is changing execution uncertainty. Second, we intend to develop more intelligent strategies for traders in the T-CDA, that are capable of learning from market observations and interactions and improving their behaviours over trading days. The first step to developing such strategies is to extend the ZI Plus strategy (Cliff and Bruten 1997) to factor in estimates of execution uncertainties of the competition in its bidding behaviour. Finally, we intend to develop a trust model and analyze how its different properties influence the individual trader's efficiency and the global social welfare.

| Symbol | Meaning |
|---|---|
| $B$, $S$ | The set of buyers, sellers |
| $A$ | The set of agents $A = B \cup S$ |
| $b_i$ | Buyer $i$ |
| $s_j$ | Seller $j$ |
| $a_k$ | Agent $k$ |
| $\ell_i^b$ | Limit price of $b_i$ |
| $\ell_j^s$ | Limit (cost) price of $s_j$ |
| $t_{i,j}(q)$ | Transaction between $b_i$ and $s_j$ at price $q$ |
| $e_i^b$, $e_j^s$ | Execution outcome for $b_i$, $s_j$ |
| $p(b_i)$, $p(s_j)$ | Probability of success of $b_i$, $s_j$ |
| $u_k^a(t,e)$ | Utility of transaction $t$ with outcome $e$ for $a_k$ |
| $\bar{u}_k^a(t)$ | Expected utility of transaction $t$ for $a_k$ |
| $\tilde{u}_k^a(t)$ | Estimated expected utility of transaction $t$ for $a_k$ |
| $\text{trust}_i(a_j)$ | Trust of $a_i$ in $a_j$ |
| $o_i^b$ | Bid by $b_i$ |
| $o_j^s$ | Ask by $s_j$ |
| $c_i(o_i^b, o_j^s)$ | Commitment by $b_i$ on transacting based on $o_i^b$ and $o_j^s$ |

Table 2: Table of notations.

## References


Buyya, R.; Abramson, D.; and Giddy, J. 2000. An economy driven resource management architecture for global computational power Grids. Las Vegas, USA: CSREA Press.

Buyya, R.; Abramson, D.; and Venugopal, S. 2005. The Grid economy. *Proc of the IEEE* 93(3):698–714.

Clearwater, S. H., ed. 1995. *Market-Based Control: A Paradigm for Distributed Resource Allocation*. New Jersey, USA: World Scientific.

Cliff, D., and Bruten, J. 1997. Zero is not enough: On the lower limit of agent intelligence for continuous double auction markets. Technical Report HPL-97-141, HP Laboratories Bristol.

Dash, R. K.; Vytelingum, P.; Rogers, A.; David, E.; and Jennings, N. R. 2007. Market-based task allocation mechanisms for limited capacity suppliers. *IEEE Transactions on Systems, Man and Cybernetics (Part A)* 37(3):391–405.

Dash, R. K.; Ramchurn, S. D.; and Jennings, N. R. 2004. Trust-based mechanism design. 748–755.

Gode, D. K., and Sunder, S. 1993. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *The Journal of Political Economy* 101(1):119–137.

Hull, J. 2006. *Options, futures, and other derivatives*. New Jersey, USA: Prentice-Hall.

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

Smith, V. L. 1962. An experimental study of competitive market behavior. *Journal of Political Economy* 70(2):111.

Tan, Z., and Gurd, J. R. 2007. Market-based grid resource allocation using a stable continuous double auction. In *Proc 8th IEEE/ACM Int Conf on Grid Computing*, 283–290.

Teacy, W. T. L. 2006. *Agent-Based Trust and Reputation in the Context of Inaccurate Information Sources*. Ph.D. Dissertation, University of Southampton.

Wolski, R.; Plank, J. S.; Brevik, J.; and Bryan, T. 2001. Analyzing market-based resource allocation strategies for the computational Grid. *International Journal of High Performance Computing Applications* 15:258–281.

Wolski, R.; Brevik, J.; Plank, J. S.; and Bryan, T. 2003. Grid resource allocation and control using computational economies. In Berman, F.; Fox, G.; and Hey, T., eds., *Grid Computing: Making The Global Infrastructure a Reality*, Wiley Series in Communications Networking & Distributed Systems. John Wiley & Sons, Ltd. 747–771.