

Market-Based Task Allocation Mechanisms for Limited Capacity Suppliers

Rajdeep K. Dash, Perukrishnen Vytelingum, Alex Rogers, Esther David and Nicholas R. Jennings

Abstract—This paper reports on the design and comparison of two economically-inspired mechanisms for task allocation in environments where sellers have finite production capacities and a cost structure composed of a fixed overhead cost and a constant marginal cost. Such mechanisms are required when a system consists of multiple self-interested stakeholders that each possess private information that is relevant to solving a system-wide problem.

Against this background, we first develop a computationally tractable centralised mechanism that finds the set of producers that have the lowest total cost in providing a certain demand (i.e. it is efficient). We achieve this by extending the standard Vickrey-Clarke-Groves mechanism to allow for multi-attribute bids and by introducing a novel penalty scheme such that producers are incentivised to truthfully report their capacities and their costs. Furthermore our extended mechanism is able to handle sellers' uncertainty about their production capacity and ensures that individual agents find it profitable to participate in the mechanism. However, since this first mechanism is centralised, we also develop a complementary decentralised mechanism based around the continuous double auction. Again because of the characteristics of our domain, we need to extend the standard form of this protocol by introducing a novel clearing rule based around an order book. With this modified protocol, we empirically demonstrate (with simple trading strategies) that the mechanism achieves high efficiency. In particular, despite this simplicity, the traders can still derive a profit from the market which makes our mechanism attractive since these results are a likely lower bound on their expected returns.

Index Terms—distributed decision making, decision theory, multi-agent systems, market-based control.

I. INTRODUCTION

Task allocation is an important and challenging problem for computer science (see [36], [5] for an overview). To this end, in this paper we specifically consider it in the context of assigning tasks to a set of autonomous software agents. Now, when a designer has complete control over both the agents and the way in which they interact, an inherently cooperative approach can be developed whereby the agents work together for the common good according to algorithms specified by the designer in order to find a system-wide solution [20], [7]. Moreover when the agents participating in the system are only concerned about the effectiveness of the overall system, planning, distributed constraint optimisation and scheduling algorithms have also been proposed [23], [10], [1], [31], [46].

However, such methods fail in systems where the agents represent distinct stakeholders whose aim is to maximise their own profit in the system (e.g. in Grid computing where the

agents represent end users and e-commerce scenarios such as e-Bay and electronic markets where agents represent the buyers and sellers). They would fail because they present the opportunity for the agent to gain an advantage by misreporting their position (either their needs or their resources). For example, an agent might over-report its need for memory capacity on a computational grid so that when the distributed constraint optimisation process is carried out, it gets allocated more memory than its share in an efficient allocation. Another example is in peer-to-peer systems where the case of free-riding (i.e. where agents under-state their available resource so as not to be asked to contribute to the system) has been well documented [2]. Moreover, it can only be assumed that if they could obtain such benefit they would do so since the agents are self-interested, rational problem solvers. Against this background, market-based techniques are attractive because their point of departure is developing protocols that achieve good system wide properties despite the fact that the agents act selfishly [5]. Specifically, *market-based control* (MBC) is concerned with using technologies and metaphors from economics to develop effective task and resource allocation systems that operate in a robust and decentralised manner [5]. To date, such techniques have been shown to have great potential in domains such as grid computing [44], peer-to-peer systems [39], multirobot coordination [15] and mobile computing [4] but have typically looked at standard cost functions and do not provide a comparison between the two main strands of MBC: namely, centralised and decentralised auctions.

In more detail, the aim of this paper is to consider the use of these market mechanisms in cases where the agents that are trying to sell the goods/resources (or provide services/tasks) have a particular form of *cost structure* (consisting of a fixed overhead cost and a constant marginal cost) and *finite production capacities* (which are both privately known to them). We believe that these traits are typical of many real world applications such as electricity markets, job-shop scheduling and grid computing applications. For example, a power plant will typically have a fixed startup cost and a constant marginal cost of running the plant upto its maximum capacity. The classic job shop scheduling problem consists of running periods composed of an initial machine set-up time (overhead cost) plus a cost per unit time (the marginal cost) and a finite capacity which these machines can run upto. Finally, agents providing computational resources on the grid incur an overhead cost (computational cost of setting up the agent managing the resource on the machine) and marginal costs as they accept tasks upto the limit that their machines can support.

In general, there are two broad classes of market mechanisms that can be considered when dealing with such problems. The first class, the reverse auction, involves a centralised mechanism in which sellers report their values to a centre (that has already aggregated the demand from the buyers) which then decides on the optimal allocation and the payments. The most popular of the centralised mechanisms is the *Vickrey-Clarke-Groves* (VCG) mechanism¹. Its popularity arises from two attractive economic properties: it is allocatively efficient (i.e. it guarantees an efficient solution in terms of finding the cheapest set of sellers that satisfies the demand) and it is individually rational (i.e. it ensures that the agents are better off joining the mechanism rather than opting out of it) [25], [8]. Unfortunately, in our case, the finite capacities of the sellers and the particular cost structure of our problem mean that the VCG no longer preserves these desirable economic properties. Thus, we need to extend the VCG mechanism in order to restore them. Such modification is important because we wish to guarantee that we find the cheapest providers and we want to ensure that participants willingly join the system. Here, we achieve these dual objectives by allowing agents to report on the triples (fixed cost, unit cost, and capacity) that characterise their types and via the use of a novel penalty scheme (detailed in section IV). We prove that the ensuing mechanism is strategyproof (i.e. every agent finds it in its best interest to truthfully reveal its private information) and robust to sellers being uncertain about their production capacity². Furthermore, we show that the mechanism is computationally tractable since the optimal allocation can be computed in pseudo-polynomial time via the use of a dynamic programming solution.

However, a potential drawback of our modified VCG mechanism (indeed of all the mechanisms in this class) is that it is inherently centralised. That is, the task allocation is computed by a single entity, the auctioneer, who does so by collecting all the private information about the costs and capacities from the various agents. Now, in some cases, this is not a problem and the optimality of the mechanism is the over-riding concern. However, in other cases, issues such as robustness to a single point of failure and scalability are more important and this gives rise to the desire for decentralised mechanisms [8]. Thus to cope with this situation, we also consider the second broad class of MBC mechanism that can be employed for the task/resource allocation problem: the *continuous double auction* (CDA) [40], [13]. In this protocol, buyers and sellers continuously submit bids (an offer to buy at price p_b) and asks (an offer to sell at price p_a) respectively (which are listed on a billboard) and the market clears (i.e. a transaction occurs) whenever the bid of a buyer matches the ask of a seller (i.e. when $p_b \geq p_a$). Such an auction is decentralised in that the

allocation of the tasks is not computed by any single agent, but rather emerges out of the interactions of the agents in the protocol³. Nevertheless, despite this decentralisation, CDAs still produce solutions that are very close to the optimal, even when the participants adopt very simple strategies⁴.

However, most work on CDAs assumes a cost structure that consists of a fixed marginal cost for each unit supplied and no startup cost. This choice of cost structure is quite natural in macro-economic models and it results both in an equilibrium market price (a unique price at which buyers and sellers agree to trade) for the commodity and in efficient allocations [25]. Unfortunately, the particular cost structure of our domain implies that no such equilibrium exists. This is due to the average unit cost of producing lower quantities is greater than that when producing larger quantities as a result of the startup cost (this is akin to models where there are economies of scale in which the startup cost is shared over a greater product run [25]). The presence of a capacity constraint further complicates matters since, in general, a single seller will not be able to fully satisfy the total demand. Furthermore, since we are developing a protocol for task allocation, we consider buyers with inelastic demand (i.e. buyers do not vary their demand according to price) which, in turn, means that the CDA is focused on finding the cheapest set of seller(s) given an exact demand from the buyers⁵. Given these points, we need to modify the standard CDA mechanism by designing suitable clearing rules and constraining the type of offers allowed in the market in order to deal with the aforementioned issues. We then assess the allocative efficiency of our market mechanism using the same methodology as was employed by Gode and Sunder in their seminal study of the standard CDA mechanism⁶ [18]. This assessment shows that the allocative efficiency of our CDA protocol is fairly high (with an average value of 83% in the scenario we consider) and that our ZI2 agents are always profitable (this condition is broadly equivalent to the individual rationality condition of the centralised mechanism).

These two mechanisms have been developed because they represent complementary task allocation mechanisms for the same domain (i.e. where the sellers have finite production capacity and the cost structure we outline). Thus, while the extended VCG mechanism *guarantees* that the cheapest set of seller(s) is always found, it is centralised. In contrast, the

³Even the seemingly centralised billboard in the CDA can be implemented using a broadcast communication protocol that mimics the typical "shouts" in the original trading pit [13].

⁴In this context, a strategy is simply a method of generating a bid or an ask given the observed current market conditions. In CDAs, it has been shown that a strategy that randomly generates bids/asks between a set lower and upper bound can be extremely efficient (both for the individual participant and in terms of the effectiveness of the overall market). Such strategies are known as zero-intelligence (ZI) strategies [18].

⁵Inelastic demand also ensures a fair comparison with the centralised case. This is because allowing for elastic demand will result in an allocation which satisfies a demand defined by the demand and supply curves rather than a prior demand that has been made by the buyers (which would occur with inelastic demand). It also allows us to characterise the cost of decentralising the market-based mechanism in terms of its efficiency loss.

⁶While their study employed ZI agents that operate purely on price, in our case, the sellers have to provide both a price and quantity vector. Thus we modify the ZI strategy to a ZI2 strategy that applies the same basic idea to both price and quantity.

¹It should be noted that before the reverse auction is conducted all information about the agents' costs are held privately by the individual agents themselves (and thus it is informationally decentralised) [25]. However, the operation of these mechanisms involves an agent that collects all information and performs a centralised computation.

²In certain scenarios, sellers may be uncertain about their capacity and would only have a best estimate of that capacity (e.g. in power generation scenarios a wind farm's capacity will depend on the strength of the wind and in a job-shop scheduling context the capacity of a machine might degrade stochastically over time).

mechanism derived from the CDA is decentralised, but it does not guarantee to find the cheapest set of sellers. Thus, in some cases, the centralised mechanism is more appropriate because efficiency cannot be compromised (e.g. when the costs involved are high or the set of agents participating in the market is low, thereby abating the disadvantages of centralisation). However, when decentralisation is a more desirable aspect (such as in cases where there are large numbers of agents or when robustness to failure is important), the CDA-based solution is more appropriate. Furthermore, our experimental results quantify the loss in efficiency that occurs when the decentralised system is implemented instead of its centralised counterpart (an average of 17% in the case we study). It is important to note that under both mechanisms, the sellers, though competitive, are profitable and they are hence always incentivised to participate in our systems.

The remainder of this paper is organised as follows: section II presents the main related work in this area. In section III, we then present our basic model of the cost structure we are considering. We explain in section III why a simple extension of the VCG mechanism would not work via the use of an example. In section IV we detail our centralised mechanism and prove its economic and computational properties. In section V, we present our decentralised protocol and empirically evaluate its properties. We conclude and suggest areas of future work in section VI.

II. RELATED WORK

The VCG mechanism and its various extensions have been used in a variety of computer systems for task allocation situations. The two broad issues that have been investigated are the economic and computational properties of these mechanisms under various scenarios (e.g. [29], [38], [41], [32], [22], [33]). Most solutions in this area consider standard demand functions (not our cost structure) in order to derive approximate solutions to the problem or to find instances where these can be solved exactly in polynomial time [35].

However recently, there has been increasing interest on the economic and computational properties of mechanisms using non-standard cost functions. In particular, a decreasing marginal cost structure has been considered in [21] and a polynomially solvable, approximately strategyproof and approximately efficient (i.e. solutions which are within a bound of the optimal) auction mechanism has been devised. In addition, more general piece-wise linear continuous curves have been considered in [11], but the incentives for truthful bidding were not taken into account. Furthermore, in [37], [16] more realistic cost curves such as those related to volume-quantity discounts are considered and expressed using particular bidding languages (which express variations on XOR and AND bids) have been investigated. However, none of these approaches would work for the cost structure of our domain since they do not consider both the economic and computational properties of problems with overhead cost, constant marginal cost and limited capacity simultaneously. Furthermore, unlike our work, they do not derive an efficient, strategyproof and individually-rational solution or compare it

with a decentralised auction. Also, they do not consider the problem of suppliers not fulfilling their commitment. This latter problem is studied in [9] and [30]. However, the mechanism in [30] considers success and failure as a binary variable and thus does not try to incentivise agents to produce up to their maximum if ever they cannot fulfil their commitment. In [9], both the producers and consumers report over the success of a transaction and thus their mechanism is more appropriate in an iterated market place where the consumers can form an opinion about the success rate of each producers. As a result, in their case, the consumers bear the risk of correctly evaluating the success rate of a producer, unlike in our mechanism where it is up to the producers to correctly estimate their capacities.

The double auction class of market mechanism consists fundamentally of two categories: the clearing-house and the CDA. The former involves all bids and asks being submitted to an auctioneer and the market being cleared periodically by that auctioneer (who calculates the allocation). In contrast, the latter clears continuously, with the competition in the market deciding the allocation rather than an auctioneer. In this context, one particularly relevant application of the double auction is by Nicolaisen et al. in a wholesale electricity market [27]. Specifically, they use a clearing-house double auction with discriminatory pricing. Now, while they do not look at the complexity involved with a cost structure, they do describe a market mechanism for resource allocation. In particular, the agents populating their markets adopt a sophisticated bidding behaviour (a modified Roth-Erev Reinforcement Learning algorithm [34]), and they evaluate the efficiency of their mechanism using such strategies. Other relevant works on the double auction include that by McCabe et al [26] on the design of a clearing-house, and Xia et al.[45] on solving combinatorial double auction mechanisms. However, these mechanisms are not decentralised like the CDA since they involve an auctioneer who computes the allocation and prices.

Speaking more generally, most research on the CDA has been on the structure and behaviour of the mechanism. Indeed the initial stimulation for this work comes from the field of experimental economics where experiments with human volunteers showed that small groups of traders could quickly find the equilibrium price in simulated single commodity markets [40], [18]. In line with this seminal work, many researchers then extended these simple trading strategies to generate sophisticated software agents that are capable of observing the trading behaviour of other agents in order to learn the market equilibrium price of a commodity, and thus trade more efficiently [42], [17], [19], [43]. However, in all of this work the existence of the market equilibrium at which both buyers and sellers seek to trade is a consequence of the assumption of a cost structure with an increasing marginal cost and no startup cost. Unfortunately the cost structure of our domain destroys this market equilibrium and thus the close to optimal efficiency usually obtained by CDAs cannot be guaranteed. Specifically, this is because the different startup costs and the inelastic demand mean that a single price on which buyers and sellers agree to trade cannot be reached. To remedy this, we develop a variant of the CDA that is still

reasonably efficient, but that can deal with the specific cost structure and capacity constraint in our domain.

III. THE ALLOCATION PROBLEM

We now discuss in more detail the problem structure that we consider in the remainder of this paper. The system which we wish to control consists of a set $\mathcal{I} = \{1, \dots, n\}$ of n suppliers of a resource and a number of consumers with total demand D . Each supplier, $i \in \mathcal{I}$, is characterised by a maximum capacity that it can provide, c_i , and a cost function, C_i . The cost function is defined as a combination of a fixed price, f_i , payable for any amount of production and a separate per unit price, u_i :

$$C_i = \begin{cases} 0 & \text{if } x = 0, \\ f_i + x_i u_i & \text{if } 0 < x_i \leq c_i \end{cases} \quad (1)$$

where x_i is the quantity of production allocated to seller S_i . Thus, an allocation vector $\mathbf{x} \in \mathcal{X}$ is one in which each agent S_i is asked to supply a quantity x_i . We assume that both the demand and the details of the cost function are private information of the producers (also referred to as suppliers or sellers) since they represent distinct self-interested stakeholders. Given this, the overall aim of the system is to satisfy the total demand by allocating production between the different producers. Here, we assume that the resource is bought and sold in small indivisible units (as is common in most billing systems) and thus $x_i \in \mathbb{N}$.

As the designer of the whole system, we are interested in ensuring that the overall allocation, \mathbf{x}^* , of the resource under consideration is optimum in the sense that it minimises the total cost of production. In this case, it is an optimisation problem where we minimise the sum of the individual production costs, whilst satisfying the total demand, $\sum_i x_i = D$, and the capacity constraints of each individual producer:

$$\mathbf{x}^* = \arg \min_{\mathbf{x}} \sum_i (\alpha_i f_i + u_i x_i) \quad (2)$$

such that $0 \leq x_i \leq c_i$ and where:

$$\alpha_i = \begin{cases} 0 & \text{if } x_i = 0 \\ 1 & \text{otherwise.} \end{cases}$$

The problem as described here is similar to two standard problems from the literature of operational research and scheduling; specifically the knapsack problem [24] and the capacitated lot-size problem [3]. Comparing this problem to the knapsack problem, we note that we can consider each supplier to be an item to be fitted into a knapsack. The size and value of each of these items is represented by the number of units of production allocated to this supplier and the cost of producing this allocation. Unlike the standard knapsack problem, where we seek to maximise the value of items without exceeding the size of the knapsack, our goal here is to exactly fill the knapsack (i.e. satisfy demand) whilst minimising the value of items placed inside (i.e. minimise the production costs). Although we can place fractional items within the knapsack, the size of these items is restricted to

integer units of production and the corresponding value of the item is given by the cost structure shown in equation 1.

Comparing to the standard capacitated lot-size problem, which attempts to schedule the production of a single producer over a number of days to meet a specific daily demand, we are attempting to schedule production over a number of different producers to satisfy an aggregate demand. Despite this difference, both problems share a similar cost structure, most specifically the combination of a fixed and per unit cost, and most importantly, both models share the concept of producers who have a constrained production capacity. We could thus adapt algorithms developed for the capacitated lot-size problem to our problem. However, in this paper the goal is to show that the problem can be solved in a computationally efficient manner rather than solve the problem in the most computationally efficient manner.

Now, both the knapsack and the capacitated lot-size problems have been shown to be \mathcal{NP} -hard [12], [14]. However both can be solved in pseudo-polynomial time using a dynamic programming approach [14] and we present a suitable implementation of this technique for our specific problem in section IV-C.

Given this problem description, in the following sections we describe our two task allocation mechanisms, starting with the centralised one.

IV. A CENTRALISED MECHANISM

Our centralised mechanism builds upon the standard VCG mechanism since this mechanism has a number of desirable economic properties with respect to task allocation (as outlined in section I). Specifically, it is *efficient*, incentivises the agents to reveal their costs truthfully to the auctioneer in *dominant strategy* (i.e. an agent finds no better option than to reveal its costs truthfully) and guarantees a non-negative utility to the participating agents.

The standard VCG mechanism for task allocation represents the producers as agents participating in a reverse auction to satisfy the demand of the auctioneer. The agents submit their respective private information about their costs, known as their types, θ_i , in sealed bids to the auctioneer. After this stage, the auctioneer finds the efficient allocation and then calculates the transfers (i.e. the amount of money that is to be paid to each agent). It is this transfer scheme that results in the agents having truthful reporting as a dominant strategy.

However, there are two key differences between our setting and that of a standard VCG mechanism. First, each agent's type has three dimensions that characterise its cost function instead of the usual one. Specifically, these dimensions are the fixed price or setup cost, f_i , the unit cost, u_i , and the capacity, c_i . Second, the capacity of the agent does not directly impact on the cost of supplying an allocated quantity of a resource, but rather puts a limit on the amount that it can supply. This differs from the standard setting of a VCG where an agent's type directly impacts on its cost. Thus, an agent overstating its capacity does not change its payment in the traditional VCG mechanism (as we show in section IV-A), but does change the efficient set of suppliers calculated by the centre.

To deal with these differences, the standard VCG needs to be extended in three ways. The first change is to have agents report the attributes that define their cost functions rather than a single cost price. The second change is to have a separate allocation and payment phase (as opposed to the traditional VCG mechanism where this is amalgamated into a single phase) since it is the very reports of the agents (i.e. that of their capacities) which define the space of *feasible* allocations. The third change is the introduction of a penalty scheme that incentivises the agents to report truthfully on their capacities⁷.

Given this, we present the payment as a two-part scheme: a transfer scheme and a penalty scheme (presented in sections IV-A and IV-B). This two-part mechanism is presented for explanatory purposes only and the overall combined mechanism is presented in section IV-C. In section IV-D, we prove the economic and computational properties of our mechanism.

A. The Transfer Scheme

The allocation problem is the same as that introduced in section III. If the agents are incentivised to report truthfully, then the auctioneer can just take their reports and solve the optimization problem introduced in section III. More generally, however, agents might not report their types truthfully if they believe that they will derive a higher profit by lying. Thus, if agents report $\hat{\theta}_i = (\hat{f}_i, \hat{u}_i, \hat{c}_i)$, the auctioneer then solves:

$$\mathbf{x}^* = \arg \min_{\mathbf{x}} \sum_i (\alpha_i \hat{f}_i + \hat{u}_i x_i) \quad (3)$$

such that $0 \leq x_i \leq \hat{c}_i$ and where:

$$\alpha_i = \begin{cases} 0 & \text{if } x_i = 0 \\ 1 & \text{otherwise.} \end{cases}$$

Hence, comparing equations 2 and 3, in order to achieve an efficient allocation we are left with the problem of incentivising the agents to report truthfully. If we assume rational self-interested agents, then this implies that they should maximise their own utility when reporting truthfully (otherwise they will lie!). As with most other work in this area, we consider the case that the agents have a quasi-linear utility function⁸ [25].

Definition 1: Quasi-linear Utility Function. A quasi-linear utility function is one that can be expressed as:

$$U_i(\mathbf{x}, t_i, \theta_i) = t_i - v_i(\mathbf{x}, \theta_i) \quad (4)$$

where $v_i(\mathbf{x}, \theta_i)$ is the cost of the allocation \mathbf{x} to agent i given its type θ_i and t_i represents the transfer of money from the centre to agent i .

⁷We should note here that the second difference does not result in inter-dependent valuations (i.e. valuations which depend on other agents' observed signals). While the capacity of each agent does change the allocation of other agents (the cheapest agent will determine how much the remaining agents will obtain via its capacity), it only does so in an indirect way. Therefore, we can still aim to achieve an efficient mechanism despite the multi-dimensionality of the types since we are firmly in the realm of private values [22].

⁸The quasi-linear utility function is a characteristic of a standard VCG mechanism and is required so as to circumvent the Gibbard-Satterthwaite impossibility result about achieving efficiency in a setting that considers general utility functions (i.e. non quasi-linear ones) [25].

Sellers			
	S_1	S_2	S_3
Capacity	100	150	175
Fixed Price	100	200	120
Unit Price	1.5	1	2

TABLE I
A SET OF THREE PRODUCERS BIDDING TO SATISFY A DEMAND OF 200 UNITS.

The standard VCG mechanism achieves truth-telling by aligning the goal of each agent with that of the mechanism designer. It does so by imposing a transfer on the agent which is equivalent to its marginal contribution to the society. Now, applying this insight to our multi-dimensional type domain, we advocate the following transfer scheme in which the agents report on all three dimensions of their types (i.e. on $\theta_i = (f_i, u_i, c_i)$):

$$t_i = \left[\min_{\mathbf{x}} \sum_{x_j \leq c_j, j \in \mathcal{I} \setminus i} (\alpha \hat{f}_j + \hat{u}_j x_j) \right] - \left[\sum_{j \in \mathcal{I} \setminus i} (\alpha^* \hat{f}_j + \hat{u}_j \hat{x}_j^*) \right]. \quad (5)$$

where \hat{x}_j^* is the allocation to agent j in the optimal allocation, $\hat{\mathbf{x}}^*$, calculated with the reports of all the agents. This scheme is the only one that completely captures an agent's marginal contribution to the system and it is therefore the only possible scheme that can be used in this context.

The transfer scheme of equation 5 consists of two parts. The first calculates the total cost of the optimal allocation if agent i were not included in the set of suppliers. In the second part, first the optimal allocation with agent i is found and then the total cost of this allocation is calculated minus the cost of this allocation to agent i . Thus, the payment that i receives is its marginal contribution to reducing the total cost of the optimal allocation. It can be observed that i will always receive a non-negative payment since the addition of a seller will only decrease the cost of the optimal allocation.

However, this is not the only change that is required to incentivise the agents to report truthfully. We now present an example which illustrates the need for an additional penalty scheme. Consider a set of producers $\{1, \dots, n\}$ with different types who are participating in a reverse auction to fulfill a demand of 200 units (i.e. $D = 200$). The producers' types, (i.e. $\theta_i = (c_i, f_i, u_i)$), are depicted in table I. They report their types to the auctioneer which then calculates the transfers according to equations 3 and 5.

Let us suppose for now that the capacity c_i of the agents are known by the auctioneer. Then, implementing our mechanism with the transfer described by equation 5, the auctioneer first chooses the optimal allocation. In this case, it would be S_2 producing 150 units and S_1 producing 50 units (i.e. $\mathbf{x} = \{50, 150, 0\}$) thereby giving a total cost of 525 to the system. The transfers would then be 220 to S_1 , 395 to S_2 and 0 to S_3 (i.e. $\mathbf{t} = \{220, 395, 0\}$). However, given this scheme, S_3 has an incentive to lie about its capacity and give a capacity greater than 200 (i.e. $\hat{c}_3 \geq 200$). It would then be allocated to produce the whole demand and would be paid 525 to do so.

However, as its true capacity is only 175 units, the demand will not be satisfied.

Thus from the above example, we can observe that an agent has an incentive to report a higher capacity than it actually has. However, an agent has no incentive to report a lower capacity. This is because the utility derived by an agent is equal to its marginal contribution to the society. Now, if an agent reports a capacity lower than its actual one and this misreport has an effect on the optimal allocation (i.e the capacity it reports is lower than the allocation it would have got under an optimal allocation), then it increases the total cost to the society since the minimisation in equation 3 would have tighter constraints. This would mean that the marginal contribution of the agent to decreasing the total cost in the society is less and hence the agent would derive a lower utility. We thus only need to worry about agents reporting a higher capacity than they actually have. We therefore impose a penalty scheme that incentivises agents to report truthfully about their capacity. In a standard VCG, such a penalty scheme does not exist since it is assumed that the producers have unlimited capacity. Furthermore, a penalty scheme imposed after the agents have supplied their allocations is the only way in which we can incentivise agents to report truthfully about their capacity. This is because the auctioneer will only know whether an agent has overstated its capacity if ever that agent has been allocated to produce over its true capacity (but under its declared one) after the agent has supplied its allocation.

B. The Penalty Scheme

We wish to penalise agents that report a higher capacity than they actually have. However, we are not concerned with untruthful reporting if this does not change the resulting efficient allocation. This is because such agents will not derive a higher utility if their untruthful reporting has not changed the efficient allocation. Thus, we will call agents whose reported capacity changes the optimal outcome *active agents*.

For example in the allocation problem given in table I, if there was a supplier S_4 with $(c_4, f_4, u_4) = (150, 200, 2)$, then even if it lied and reported $\hat{c}_4 = 400$, it would not make a difference in our optimal allocation (since its cost of supplying 200 units is 550 and this is still greater than the efficient outcome calculated previously).

In order to know whether the active agents have truthfully reported their capacity, we require a post-production stage that checks how much they actually produced. We shall assume that if an agent is asked to supply a certain amount \hat{x}_i^* , and actually produces only \bar{x}_i , ($\bar{x}_i < \hat{x}_i^*$), then the capacity of that agent is \bar{x}_i . We shall see that given the penalty we design, this assumption is satisfied with rational agents. It is only in the case of malicious agents who want to increase the cost to the system with no consideration to their own utility for which the following penalty scheme would not work.

In more detail, we impose the following penalty, p_i , if the agent does not supply the amount that it was required to supply under the optimal allocation (i.e. if $\bar{x}_i < x_i^*$):

$$p_i = t_i(x_i \leq \hat{c}_i) - t_i(x_i \leq \bar{x}_i) + \delta \quad (6)$$

where $t_i(x_i \leq \hat{c}_i)$ is the transfer in equation 5 computed with the constraint $x_i \leq \hat{c}_i$, $t_i(x_i \leq \bar{x}_i)$ is the one computed with the constraint $x_i \leq \bar{x}_i$ and $\delta > 0$. Intuitively, the penalty scheme ensures that an agent overstating its capacity would derive strictly less utility than when it provided a truthful report by making such an agent derive an overall transfer of $t_i(x_i \leq \bar{x}_i) - \delta$. Note that in the event that such an agent has misreported so as to be in the active set, $t_i(x_i \leq \bar{x}_i)$ would then be zero and thus that agent would derive a negative utility equal to $-\delta$ since $t_i(x_i \leq \hat{c}_i)$ would also be removed from its utility.

This penalty scheme, which is a transfer of money from the agent to the auctioneer, consists of three parts. The first is the transfer that occurs with the reported capacity \hat{c}_i . The second part is the transfer that would have resulted if the agent had reported its capacity as the amount that it has successfully supplied. This penalty scheme thus only penalises agents in the case where their misreported capacity has changed the allocation of supply. The third part is the one that ensures that the utility an agent derives from misreporting its capacity is strictly lower than when it tells the truth (i.e it is then a strongly dominant strategy for the agent to report its truthful capacity).

It should also be noted that although this penalty scheme has been developed for the case of agents misreporting their capacity, it would also penalise agents that have not produced the specified amount due to other reasons. This penalty scheme thus puts the onus on the agents to provide an accurate report of the amount they can produce. The value of δ can thus be set by the mechanism designer depending on how critical it is to meet demand. The more critical the requirement, the higher δ should be set. Evidently, this sacrifices efficiency (the agents report a lower capacity than their most likely capacity) for robustness. Another attractive aspect of this penalty scheme is that if ever an agent realises after the allocation that it cannot produce the amount assigned to it, it would still produce till its limit so as to reduce the ultimate penalty. This penalty scheme can also be potentially coupled with a reputation mechanism such that instead of having a δ which is uniform over all the agents, each agent i could have a specific δ_i which the centre tunes according to past interactions with the agent (e.g. one could use the simple trust model in [9] in order to model past interactions and condition the penalties accordingly). Further details about the actual calibration of δ when the system is designed for the case in which agents are unsure about their capacity are given in the appendix.

Thus, in our example in table I, if agent S_3 reported $\hat{c}_i = 200$, it would be penalised $525 + \delta$ (from equation 6). As a result, the agent does not profit by lying. In the case of the two other agents, S_1 and S_2 , misreporting their types, they incur a loss in utility equal to δ .

C. The Mechanism

We can amalgamate the two-part payment-scheme presented in sections IV-A and IV-B into an equivalent one-stage payment, thus yielding the following centralised mechanism:

- 1) First the seller agents, S_i , provide reports of their types $\hat{\theta}_i = (\hat{f}_i, \hat{u}_i, \hat{c}_i)$ to the centre.

- 2) The centre, having gathered total demand from the buyer agents, solves equation 3 and assigns production to the agents according to the optimal allocation vector $\hat{\mathbf{x}}^*$.
- 3) The centre then provides the overall payment m_i to the agents once they have produced their allocation:

$$\begin{aligned} m_i &= t_i - p_i \\ &= t_i(x_i \leq \bar{x}_i) - \delta\beta_i \end{aligned} \quad (7)$$

where β_i is a binary variable indicating when a supplier has over-reported its capacity and equals 1 when $\bar{x}_i < \hat{x}_i^*$.

D. Properties of the Mechanism

We now prove the properties of our mechanism. To this end:

Proposition 1: The mechanism is strategyproof.

Proof: A mechanism is *strategyproof* if it is a dominant strategy for the players to reveal their types truthfully (i.e. stating their types truthfully is a best strategy for an agent no matter what the type of the agent and no matter what strategy the other agents follow). Here, we need to prove that truthful reporting is a dominant strategy for the agents given the transfer and penalty schemes in our mechanism. We first consider the case that the agent has not over-reported its capacity. Then its strategy is to report $\hat{\theta}$ so as to maximise its utility:

$$\begin{aligned} \hat{\theta}_i &= (\hat{u}_i, \hat{f}_i, \hat{c}_i) = \arg \max_{\theta_i \in \Theta_i} (U_i(\theta_i), \mathbf{x}) \\ &= \arg \max_{\theta_i \in \Theta_i} \left[(\hat{\alpha}_i^*(\hat{f}_i - f_i) + (\hat{u}_i - u_i)\hat{x}_i^*) \right. \\ &\quad \left. - \min_{\substack{\mathbf{x} \\ x_j \leq \hat{c}_j, j \in \mathcal{I}}} \sum_{j \in \mathcal{I}} (\alpha_j \hat{f}_j + \hat{u}_j x_j) \right] + \\ &\quad \left. \min_{\substack{\mathbf{x} \\ x_j \leq \hat{c}_j, j \in \mathcal{I} \setminus i}} \sum_{j \in \mathcal{I} \setminus i} (\alpha_j \hat{f}_j + \hat{u}_j x_j) \right] \\ &= \arg \max_{\theta_i \in \Theta_i} \left[(\hat{\alpha}_i^*(\hat{f}_i - f_i) + (\hat{u}_i - u_i)\hat{x}_i^*) \right. \\ &\quad \left. - \min_{\substack{\mathbf{x} \\ x_j \leq \hat{c}_j, j \in \mathcal{I}}} \sum_{j \in \mathcal{I}} (\alpha_j \hat{f}_j + \hat{u}_j x_j) \right] \end{aligned}$$

The first part of the maximisation is the gain or loss that an agent makes by misreporting its type, whereas the second part is the effect that this misreporting has on the allocation and the global cost. Hence any misreport on its type is cancelled out by the effect on the global cost. The important point to note here is that the minimisation is not carried out by the agent but by a centre that is only aware of $\hat{\theta}_i$. Hence, in order to maximise the term in $[\cdot]$ above, an agent should report $\hat{\theta}_i = (f_i, u_i, \hat{c}_i)$. That is, truthtelling in (f_i, u_i) is a weakly dominant strategy (it is only weakly so because in certain cases an agent on lying would derive a utility which is equivalent to what it derives if it told the truth). Thus, we have proved that the mechanism is strategyproof in (f_i, u_i) . Furthermore, we know that an agent will not report a lower capacity (as per the discussion in section IV-A).

Now, we prove that under the penalty scheme the agent will not report a capacity higher than its actual one. The utility of an agent i , given that it has reported a higher capacity, is

the sum of its cost, transfer and penalty. We now prove that overreporting its capacity is a weakly dominated strategy for an active agent (i.e. overreporting one's capacity is never better than stating one's capacity truthfully). From equations 4 and 7, the utility of an agent would then be:

$$\begin{aligned} U_i(\cdot) &= \max_{\theta_i \in \Theta_i} \left[(\hat{\alpha}_i^*(\hat{f}_i - f_i) + (\hat{u}_i - u_i)\hat{x}_i^*) \right. \\ &\quad \left. - \min_{\substack{\mathbf{x} \\ x_j \leq \bar{x}_j, j \in \mathcal{I}}} \sum_{j \in \mathcal{I}} (\alpha_j \hat{f}_j + \hat{u}_j x_j) \right] - \delta\beta \\ &< \max_{\theta_i \in \Theta_i} \left[(\hat{\alpha}_i^*(\hat{f}_i - f_i) + (\hat{u}_i - u_i)\hat{x}_i^*) \right. \\ &\quad \left. - \min_{\substack{\mathbf{x} \\ x_j \leq \hat{c}_j, j \in \mathcal{I} \\ \hat{c}_i = c_i}} \sum_{j \in \mathcal{I}} (\alpha_j \hat{f}_j + \hat{u}_j x_j) \right] \end{aligned}$$

Thus, together with the fact that an agent would not report a lower capacity (since such a report would mean that its resulting allocation is less or equal to the one when it reports truthfully), the above proves that an agent will always report its truthful capacity c_i . Hence we have that the agent always reports truthfully about its type θ_i . ■

Proposition 2: The mechanism is efficient.

This implies that the centre finds the outcome given by equation 2.

Proof: The above is a result of the strategyproofness of the mechanism. Since the goal of the centre is to achieve efficiency, then given truthful reports, the centre will achieve efficiency. ■

Proposition 3: The mechanism is individually-rational.

A mechanism is *individually rational* if there is an incentive for agents to join it rather than opting out of it. We begin by assuming that the utility an agent derives from not joining the mechanism is 0. Then, we need to prove that the utility an agent derives in the mechanism is always ≥ 0 .

Proof: Given the strategyproofness of the mechanism, the utility of an agent is:

$$\begin{aligned} U_i(u_i, f_i, c_i) &= - \min_{\substack{\mathbf{x} \\ x_j \leq c_j, j \in \mathcal{I}}} \sum_{j \in \mathcal{I}} (\alpha_j f_j + u_j x_j) \\ &\quad + \min_{\substack{\mathbf{x} \\ x_j \leq c_j, j \in \mathcal{I} \setminus i}} \sum_{j \in \mathcal{I} \setminus i} (\alpha_j f_j + u_j x_j) \end{aligned}$$

The first minimisation is over a larger set than the second one. Thus:

$$\min_{\substack{\mathbf{x} \\ x_j \leq c_j, j \in \mathcal{I}}} \sum_{j \in \mathcal{I}} (\alpha_j f_j + u_j x_j) \leq \min_{\substack{\mathbf{x} \\ x_j \leq c_j, j \in \mathcal{I} \setminus i}} \sum_{j \in \mathcal{I} \setminus i} (\alpha_j f_j + u_j x_j)$$

Hence, $U_i(u_i, f_i, c_i) \geq 0$. ■

Proposition 4: The mechanism is robust to uncertainties about the capacity of agents.

In this case, we impose less stringent information requirements on the agents when reporting their capacity. So far, we have considered the case where prior to revealing its type an agent is aware of its capacity. However, we believe that this may not be always practical since the capacity of a supplier may depend on numerous external factors (as discussed in section I). We therefore relax this requirement and consider the case where an agent is aware of only the probability

distribution function (pdf) relating to its capacity. We next prove that the designer can, via the setting of δ , force the agent to either report safe values (i.e. the agent is nearly certain that it will produce at least this capacity) or more risky but potentially more profitable ones.

Proof: We start by looking at the expected utility of an agent given that the probability distribution function of its capacity, $pdf(c_i)$, ranges from a lower bound \underline{c}_i to an upper bound \bar{c}_i :

$$\begin{aligned} E[U_i(c_i, f_i, u_i)] &= E \left[\min_{\mathbf{x}} \sum_{x_j \leq \hat{c}_j} (\alpha_j f_j + u_j x_j) \right. \\ &\quad \left. - \min_{\mathbf{x}} \sum_{x_j \leq \hat{c}_j} (\alpha_j f_j + u_j x_j) - \delta \beta_i \right] \\ &= \min_{\mathbf{x}} \sum_{x_j \leq \hat{c}_j} (\alpha_j f_j + u_j x_j) \\ &\quad - \int_{\hat{c}_i}^{\bar{c}_i} \min_{\mathbf{x}} \sum_{x_j \leq \hat{c}_j} (\alpha_j f_j + u_j x_j) pdf(c_i) dc_i \\ &\quad - \delta P(c_i < \hat{c}_i) \end{aligned}$$

Now, let us analyse how the reports of the agents impact on their utility. The safest report is the minimum report \underline{c}_i . Reporting a higher capacity would then yield a gain of:

$$\begin{aligned} \Delta E[U_i(c_i, f_i, u_i)] &= -\delta P(c_i < \hat{c}_i) + \left[\min_{\mathbf{x}} \sum_{x_j \leq \underline{c}_j} (\alpha_j f_j \right. \\ &\quad \left. + u_j x_j) - \int_{\hat{c}_i}^{\bar{c}_i} \min_{\mathbf{x}} \sum_{x_j \leq \hat{c}_j} (\alpha_j f_j \right. \\ &\quad \left. + u_j x_j) pdf(c_i) dc_i \right] \end{aligned} \quad (8)$$

The agents would then try to maximise the above gain given a certain δ . Thus, the setting of δ would then depend on how certain we want the agents to be about being able to satisfy their capacity. Hence, given $P(c_i \geq \hat{c}_i)$, setting δ as:

$$\begin{aligned} \delta &= \left[\min_{\mathbf{x}} \sum_{x_j \leq \underline{c}_j} (\alpha_j f_j + u_j x_j) - \min_{\mathbf{x}} \sum_{x_j \leq \hat{c}_j} (\alpha_j f_j \right. \\ &\quad \left. + u_j x_j) P(c_i \geq \hat{c}_i) \right] / (1 - P(c_i \geq \hat{c}_i)) \end{aligned}$$

results in no expected gain for the agent. In fact from equation 8, if we consider a fixed δ , then as \hat{c}_i increases, the part in $[\cdot]$ increases while $-\delta P(c_i < \hat{c}_i)$ decreases. Thus there is a \hat{c}_i for a fixed δ that results in a maximum gain. We can therefore conclude that as δ increases, $\hat{c}_i \rightarrow \underline{c}_i$ and as δ decreases, $\hat{c}_i \rightarrow \bar{c}_i$. ■

The second part of the robustness is that even if the agent realises after reporting \hat{c}_i that $c_i < \hat{x}_i^*$ (and it is asked to produce $c_i < x_i^* \leq \hat{c}_i$), it will still produce upto c_i as a result of the payment and penalty scheme.

Proof: This is evident from the way the centre pays the agents. The agents get a higher utility with a higher production since the transfer depends on how much they produce (i.e. \bar{x}) after the allocation. Specifically, consider an agent i that has

overestimated its capacity in such a way that it affects the efficient allocation (i.e. in the efficient allocation calculated by the centre from the reported types, $c_i < x_i^* \leq \hat{c}_i$). Then, that agent derives a utility of:

$$\begin{aligned} U_i(\cdot) &= t_i(x_i \leq \bar{x}_i) - f_i - u_i \bar{x}_i - \delta \\ &= \min_{\mathbf{x}} \sum_{x_j \leq \hat{c}_j} (\alpha_j \hat{f}_j + \hat{u}_j x_j) \\ &\quad - \min_{\substack{\mathbf{x} \\ x_j \leq \hat{c}_j \\ c_i = \bar{x}_i}} \sum_{j \in \mathcal{I}} (\alpha_j \hat{f}_j + \hat{u}_j x_j) - \delta \end{aligned} \quad (9)$$

Only the second term of equation 9 can be affected by agent i varying its production amount \bar{x}_i . Since agent i wants to maximise its utility, it would want this second term to be as small as possible and therefore make \bar{x}_i as large as possible, which is achieved by making $\bar{x}_i = c_i$. ■

So far, we have discussed the use of a uniform δ which is chosen according to how critically demand has to be met. This approach penalises the agent during the current interaction. An alternative approach to dealing with uncertainties in seller capacity would be to instead penalise the agent over-reporting its capacity during future interaction. A possible way of implementing this is via the use of a trust-based mechanism (TBM)[9]. Whilst detailing such a mechanism is outside the scope of this paper, we provide in the appendix an intuitive explanation on how such a mechanism would differ from the one proposed here.

Proposition 5: The optimal task allocation to the agents can be computed exactly by the centre in pseudo-polynomial time.

Proof: The centre can calculate the task allocation to the agents exactly using dynamic programming. Specifically, we wish to calculate $C[n, D]$ — the minimum total cost to satisfy a demand of D with access to n producers. This can be solved using the recursive expressions:

$$\begin{aligned} C[0, d] &= \begin{cases} 0 & \text{if } d = 0 \\ \infty & \text{if } d > 0 \end{cases} \\ C[i, d] &= \min_x \begin{cases} C[i-1, d] \\ C[i-1, d-x] + f_i + x u_i \end{cases} \end{aligned}$$

such that $0 < x \leq c_i$. As the production allocated to each producer is in indivisible units, we can calculate $C[n, D]$ by evaluating all nD possible values. This results in an algorithm which operates in pseudo-polynomial time.

In particular, a simple algorithm for this solution is presented in figure 1. Here we calculate all the values of the array, $C[n, D]$, starting from the known case $C[0, 0] = 0$ and using the recursive expressions above to calculate subsequent values. A more efficient solution may be found using primal-dual algorithms [28]. However for the size of problem tackled here, the above solution is extremely efficient. Moreover, the same approach can then be used to calculate the resulting task allocation to the agents. ■

V. A DECENTRALISED MECHANISM

So far we have considered a centralised mechanism in order to deal with our task allocation problem. However, as discussed


```

Calculate initial row of matrix  $C$ 
 $C[0,0] \leftarrow 0$ 
for  $d = 1$  to  $D$  do  $C[0,d] \leftarrow \infty$ 
Loop through the total number of producers
for  $i = 1$  to  $n$  do
  Loop through the total demand
  for  $d = 0$  to  $D$  do
     $C[i,d] \leftarrow C[i-1,d]$ 
    Loop through the total capacity of producer  $i$ 
    for  $x = 1$  to  $\min\{d, c_i\}$  do
      Compare the previous result to the current
      result and select the minimum of the two
       $C[i,d] \leftarrow \min\{C[i,d], C[i-1,d-x] + f_i + xu_i\}$ 
Return the final result
return  $C[n,D]$ 

```

Fig. 1. Pseudo-code representing the dynamic programming solution to find the optimum centralised solution in pseudo-polynomial time.

in section I, we sometimes require a mechanism for task allocation in which there is no centre that governs the allocations. Therefore, in this section, we consider the CDA as the second class of MBC task allocation mechanism.

Our task allocation problem involves multiple suppliers and multiple buyers, and the matching of the two is determined by the sellers and buyers who successfully transact with one another. As discussed in sections I and II, the most common CDA format assumes buyers and sellers have an increasing marginal cost and no startup cost and the offers in the trade are via price alone. However, in our case, the total production cost depends on both the startup cost and the number of units to be sold (given the marginal cost). In fact, since the startup cost is distributed over the sale quantity, the cost price is not fixed for different numbers of units sold. As a result, the supplier cannot firmly decide on an asking price (based on the production cost per unit or *cost price*) that would allow it to be profitable and to participate in the task allocation (by transacting with potential buyers). This is because the sale quantity cannot be known *a priori*. To overcome this, we assume that it is possible for the supplier to make a prediction about the amount of units it expects to sell (since exact demand can only be estimated)⁹. Now, in traditional cost settings, a supplier can start making bids for a low quantity and slowly ramp up his price so as to ensure he does not make a loss. However, in our scenario, low quantities correspond to higher unit prices. Thus the supplier is faced with the problem that reducing its price may not guarantee that it transacts and in certain cases may lead to a loss (if a buyer specifies a demand such that the ask price becomes lower than the cost price). We therefore allow sellers to communicate the amount they wish to sell to the market via a multidimensional bid consisting of both quantity and price. We also specify in our clearing rules that a transaction only occurs when a buyer makes a bid for this amount.

Given this background, a key objective for the decentralised mechanism is to be *individually rational* (as defined in section

IV-D). In this case, this means ensuring the suppliers can be profitable in the market so that they are incentivised to enter it in the first place. Furthermore, while the mechanism has to be individually rational, our global objective is to achieve the most efficient outcome (task allocation) that we can. Now as we discussed in section III, this is equivalent to finding the allocation that minimises total cost. In a typical CDA mechanism, the optimal task allocation occurs when the total profit of all buyers and all sellers is maximised [43] and this occurs when the combined cost of sellers is minimised on the sell side¹⁰, as the sellers with the lowest cost would be successful.

However, given our additional constraints of limited capacity and a startup cost, the seller's strategic behaviour would be more complex than that of the buyer, since, as we mention before, it additionally has to strategise over the quantity it is expected to sell. In this context, we cannot achieve full efficiency because no agent has complete information about every other agent in the market (unlike in section IV where the centre is aware of everyone's cost functions and capacities) and the sellers do not have increasing marginal costs which would guarantee an equilibrium price for trade [25].

Given this, our aim is to design a protocol that achieves a level of efficiency that is reasonably close to the optimal solution given by our centralised mechanism. To do this, we now outline our protocol, and then go on to compare its performance with its centralised counterpart in terms of task allocation efficiency.

A. The Mechanism

The protocol we propose is a variant of the multi-unit CDA. Buyers and sellers can submit offers to buy and sell multiple units of the resource, respectively, and those orders are queued in an order book which is cleared continuously (with additional constraints as a result of buyers' inelastic demands). The protocol proceeds as follows:

- Buyer i submits an offer, $bid(q, p, i)$, to buy exactly q ($q \geq 1$) units of the good at the unit price p . The utility of buyer i for a quantity other than q is 0.
- Conversely, supplier S_j submits an offer, $ask(q, p, j)$, to sell a maximum of q ($q \geq 1$) units at unit price p .
- These bids and asks are queued in an *orderbook*, which is a publicly observable board listing all the bids and asks submitted to the market (see table II). The bids in the order book are sorted in decreasing order of price and the asks are in increasing order (higher bids and lower asks are more likely to result in transactions).
- The *clearing rule* in the market is as follows. Whenever a new bid or ask is submitted, an attempt is made at clearing the order book. The orderbook is cleared *whenever* a transaction can occur (that is, when the lowest asking price is higher than the highest bidding price *and* any bidding offer can be cleared completely and the bidding quantity for each offer is completely satisfied by the supply to be cleared). The transaction price is set at

⁹In fact, in CDA scenarios demand cannot be known even after the bids have been submitted [6]. This is why sellers try to predict the demand in order to be more profitable [19].

¹⁰Sell side refers to the market from the sellers' perspective.

the bidding price which we experimentally find to result in the total market profits being equally divided between the sell side and the buy side¹¹ [43].

Order Book	
Bids (quantity, price, buyer)	Asks (quantity, price, seller)
(30, 2.95, 2)	(60, 2.20, 3)
(40, 2.75, 5)	(25, 2.60, 1)
(30, 2.70, 1)	(40, 3.22, 2)
(24, 2.16, 3)	(100, 3.50, 5)
	(25, 3.69, 7)
	...

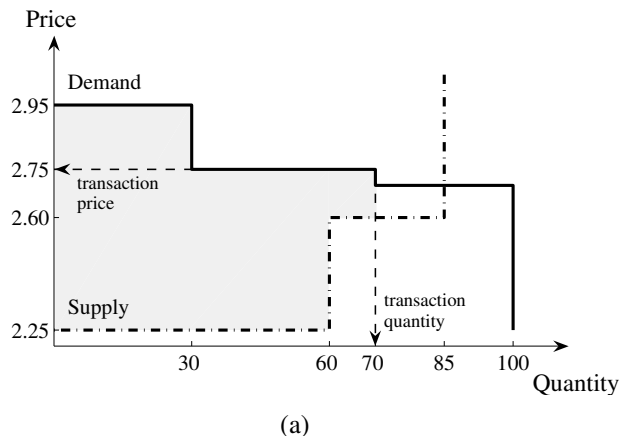
TABLE II
MULTI-UNIT CDA ORDER BOOK - BEFORE CLEARING

To further illustrate this process, we present a graphical representation of the clearing rule in figure 2. As can be seen, the offers queued in the orderbook are used to build demand and supply curves. All bids with a unit price lower than the lowest unit ask price and, similarly, all asks with a unit price higher than the highest unit bid price, cannot result in any transaction and are not represented in the figure. The transaction price and quantity are clearly shown in the figure (2.75 and 70 respectively), as the point where the demand curve crosses the supply curve under the additional constraint that bid offers are not divisible. At this transaction price, the total profit of all buyers and sellers that transact is maximised with all constraints specified by our protocols satisfied. The orderbook in table II can thus be cleared as shown in figure 2 resulting in the new orderbook given in table III. The market clearing is then similar to solving an optimisation problem where the objective is to maximise the total profit of buyers and sellers that will transact given that cleared demand must be equal to cleared supply and no partial clearing of bid is allowed.¹²

Now in order to compare the efficiency of this protocol with that of the centralised mechanism, we assume that the buyers have high limit prices (this represents price inelasticity because buyers are willing to pay any price to acquire the goods and is equal to an arbitrary maximum price that a bid or an ask can be submitted at). Furthermore, we adopt the approach of Gode and Sunder [18] in employing a zero-intelligence strategy in order to find the underlying efficiency of our market. To this end, we next present the ZI2 that is tailored to the bidding structure of our CDA protocol, before we detail the actual

¹¹We chose this option because a mechanism where most of the profits in the market were distributed among sellers would be less appealing to buyers than one where a larger share of profits were distributed among buyers. Thus, with a similar preference among sellers (who will join a market where more profit is distributed among the sell side), a mechanism that equally distributes market profits among the buy and sell side is the rational preference for both buyers and sellers.

¹²We note that other clearing rules are also possible, for example to maximise the number of transactions or to maximise profits of the sellers only. However, the aim of a market mechanism is to maximise social welfare by maximising the total profit extracted in the market, and it is achieved through the simple ordering order books that publicly shows which buyers (with highest valuation of the goods) can transact with which suppliers (with the lowest ask prices).



Order Book	
Bids (quantity, price, buyer)	Asks (quantity, price, seller)
(30, 2.95, 2)	(60, 2.20, 3)
(40, 2.75, 5)	(25, 2.60, 1)
(30, 2.70, 1)	(40, 3.22, 2)
(24, 2.16, 3)	(100, 3.50, 5)
	(25, 3.69, 7)
	...

(b)

Fig. 2. Panel (a) shows the demand and supply (curves) of the order book, with the shaded region representing allocations. Panel (b) points out the *clearable bids and asks* in the order book (shaded area in panel (a)).

Order Book	
Bids (quantity, price, buyer)	Asks (quantity, price, seller)
(30, 2.70, 1)	(15, 2.60, 1)
(24, 2.16, 3)	(40, 3.22, 2)
	(100, 3.50, 5)
	(25, 3.69, 7)

TABLE III
MULTI-UNIT CDA ORDER BOOK - AFTER CLEARING

evaluation.

B. The ZI2 Strategy

One of the principal concerns in developing a market mechanism is to ensure that it is efficient even when the participants adopt a simple strategic behaviour. The underlying intuition here is that by considering such behaviour, we are able to establish a lower bound on the efficiency of the mechanism and we can consider the extent to which the market mechanism itself affects the efficiency of the market. Thus, the ZI strategy [18] is widely used for this purpose since it is not motivated by trading profit and effectively ignores the state of the market and past experience when forming a bid or an ask. It simply draws its offer price from a uniform distribution over a given range.

Since in our mechanism, the asks consist of price and quantity, we extend the ZI strategy to our *ZI2 strategy* that randomises over both price and quantity. As discussed earlier, any sophisticated strategy, on the sell side, would make some form of prediction on the number of units it is likely to sell as part of its price formation process (because information about the actual demand is not available and there is uncertainty as to whether the agent is more competitive than the other participating suppliers). Our ZI2 supplier j , instead, randomises over the expected transaction quantity to form a limit price ℓ_j which is used as in the original ZI strategy. Thus the ZI2 strategy is¹³:

For buyer i ,

$$\begin{aligned} p_i &\sim \mathcal{U}(0, \ell_i) \\ \text{offer} &= \text{bid}(q_i, p_i, i) \end{aligned} \quad (10)$$

For seller j ,

$$\begin{aligned} \hat{q}_j &\sim \mathcal{U}(0, c_j) \\ \ell_j &= (f_j + \hat{q}_j u_j) / \hat{q}_j \\ p_j &\sim \mathcal{U}(\ell_j, \text{max}) \\ \text{offer} &= \text{ask}(c_j, p_j, j) \end{aligned} \quad (11)$$

Buyers are endowed with high limit prices at the beginning of the auction (because they have inelastic demand), while sellers are endowed with their cost functions and capacities (collectively referred to as the production function). Buyer i submits offers to buy the quantity q_i it requires at a unit price drawn from a uniform distribution ranging from 0 to its limit price ℓ_i (see equation 10). Conversely, seller j submits an ask between its limit price and *max* as per equation 11, where c_j is its production capacity, f_j is its startup cost, u_j is its marginal cost.

C. Empirical Evaluation of the Mechanism

In order to perform empirical evaluations, we have developed an implementation of this distributed mechanism based on the protocol and strategies described here¹⁴. As the experimental setup, we ran the simulations over 2000 rounds¹⁵ for two different markets, more specifically a small market with 3 buyers and 3 sellers (*market A*) and a larger market with 15 buyers and 15 sellers (*market B*). We consider both the small and large markets so as to demonstrate the scalability of our mechanism.

In each market, each seller was given a production function (supply for market A is given in table I), while each buyer was required to procure an exact quantity of units with a

relatively *high* limit price. We ran different simulations for each market, with different total demands ranging from 1 to the maximum production quantity. The total demand, D , was distributed among the buyers (see table IV for the demand in market A, where $D = \sum_i q_i, D \in [1, 425]$ given the sellers' production functions in table I). Thus, the total demand in market A was varied from 1 to 425 (the maximum supply quantity of market A), while in market B the total demand ranged from 1 to 2400.

Buyers' Demand			
B_i	B_1	B_2	B_3
allocation 1	100	150	50
...
allocation n	q_1	q_2	q_3

TABLE IV
A SET OF THREE BUYERS WITH DIFFERENT DEMANDS.

In order to empirically evaluate the efficiency of the mechanism, in terms of minimising the total cost of production, we measure this property and compare it to the optimal solution found in the centralised mechanism. Given each total demand, the mean efficiency of the market (averaged over 2000 independent rounds) is shown in figure 4, where the optimal production cost is normalised to 1, while the total production cost of the centralised and the decentralised mechanisms are shown in figure 3. As can be seen, the mechanism is efficient with an average efficiency of 83% (and a minimum efficiency of 53% when demand is relatively low) for the market B and an average efficiency of 86% (and a minimum efficiency of 67%) for market A. In both cases, the minimum efficiency case occurs when the demand is split amongst many more suppliers than are actually needed (with respect to the optimal allocation). This increases the overall cost of supply as a result of the fixed cost of the extraneous suppliers. However, in the typical CDA, the worst-case analysis considers the average efficiency of ZI agents [18]. This is because although it is theoretically possible for an allocation of very low efficiency to occur, in almost every run (higher than 99% of the time), the CDA implemented with agents employing the ZI strategy has a high efficiency. Thus, it is the zero-intelligence nature of the strategy which provides a lower bound on measuring efficiency and, we expect the average efficiency with a more informed strategy to be better [6], [19], [43]. We therefore adopt this approach in discussing the inherent efficiency of our CDA mechanism.

In the experiments with each market, we observe an increasing trend whereby the market efficiency increases as total demand approaches the maximum capacity of the sellers (see figure 4). It can also be seen that there is a high variance when the total demand is relatively low. Considering specifically the set of experiments with market A, the intuitions behind these observations are as follows. The variance of the market efficiency is generally higher when the total demand is low. This is because the optimal allocation for a total demand of 100 is completely covered by seller 1 (with a marginal cost of 1.5 and a startup cost of 100). However, our market mechanism does not ensure that only seller 1 will trade and, thus, sellers 2

¹³ $X \sim \mathcal{U}(A, B)$ describes a discrete uniform distribution between A and B, with steps of 0.01.

¹⁴Available at <http://www.ecs.soton.ac.uk/~pv03r/simulator>

¹⁵The results were validated using a students t-test with two samples of 2000 runs, assuming equal variance with means $\mu_1 = 0.7198$ and $\mu_2 = 0.7218$ and p-value $p = 0.3660$. This means that the difference between the means is not significant.

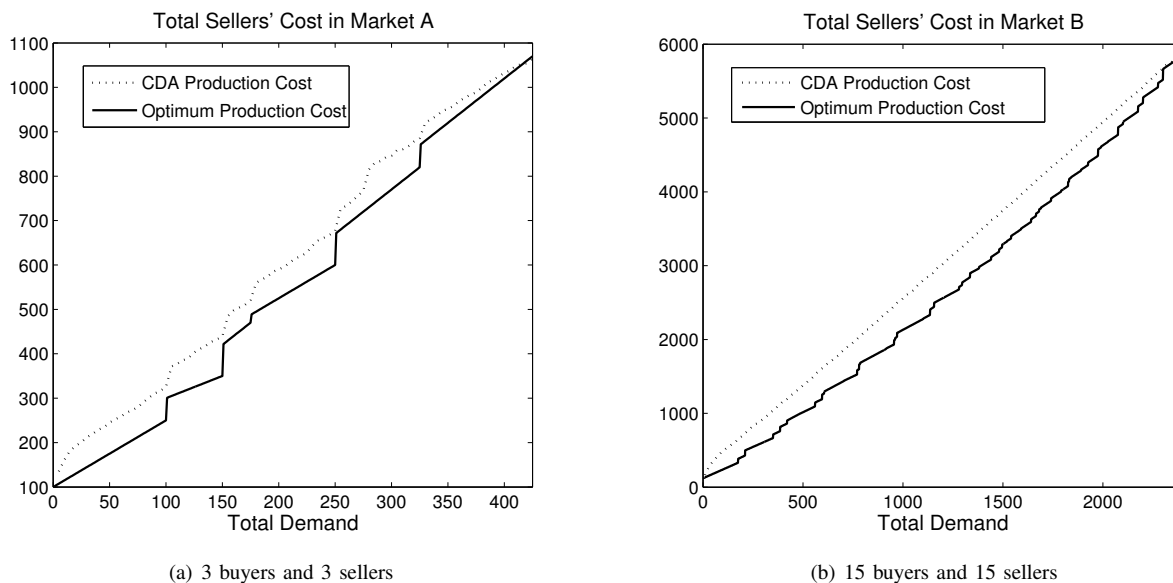


Fig. 3. Optimal and CDA production cost

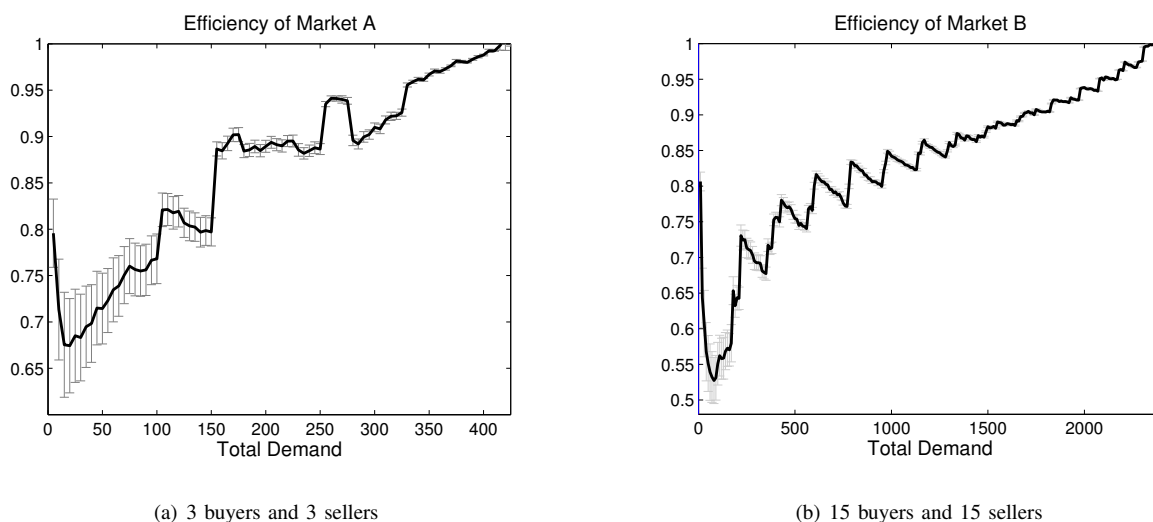


Fig. 4. Average market efficiency

and 3 may also be part of this allocation for the total demand of 100. The high variance is principally an artefact of the additional startup costs if more than one seller were to trade. As the total demand increases past 175, the optimal allocation is covered by at least two sellers. Again, the variance past the demand of 175 is the result of sellers supplying different numbers of units at different marginal costs, with at most one additional startup cost. When the total demand is very high, close to the total capacity, all the sellers participate in the allocation, and the small variance is solely due to the sellers providing different numbers of units (a difference which is relatively low compared to the total startup cost). The observations in the set of experiments with market B can also be explained by the same reasoning, with the higher variance occurring when demand that can be covered by a single seller is distributed among multiple sellers.

Furthermore, we can explain the increasing trend of the market efficiency seen in figure 3. Considering market A, a demand of up to 175 can be provided by only 1 seller. The jumps in figure 3 correspond to the optimal allocation changing between a combination of one to three sellers. For example, jumps at 100 and 150 correspond to the optimal allocation starting with seller 1, changing to seller 2 and finally to seller 3. The increase in efficiency as total demand increases is the result of the number of sellers involved in the optimal allocation, changing from a single seller (up to a total demand of 175) to three sellers (past a total demand of 325 which is the highest demand any two sellers can cover). However, in our market, any number of sellers can trade at any time. Thus, as total demand increases, the loss in efficiency that arises from the extra startup costs (compared to the optimal allocation) decreases which in turn explains the generally increasing trend.

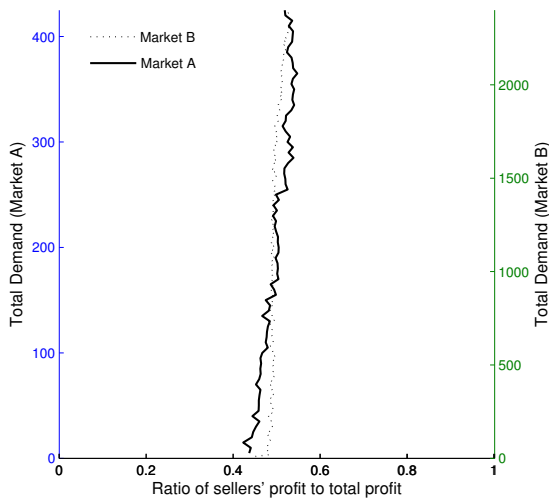


Fig. 5. The sellers' total profit given different demands (for market A with 3 buyers and 3 sellers and market B with 15 buyers and 15 sellers).

In the simulations with market B, a similar trend can be observed, with a lower efficiency when demand is lower than the minimum sellers' capacity (210). As in market A, there are more inefficient allocations that can arise when demand is low (and can be satisfied by a single seller), which would decrease the average efficiency much more than it would given a smaller number of inefficient allocations. Here, we use the same reasoning as in market A to explain the jumps, which are larger in number given the larger number of participants.

As well as being efficient, the simulation results in figure 5 show that, broadly the sellers and buyers do indeed equally share the market profits (the ratio of sellers' profits to total market profit is approximately equal to 0.5 in both cases). This fair division of profits arises from the design of the clearing rule (see section V-A). This is important because this profitability means that the agents are incentivised to enter the market which means our distributed mechanism can be viewed as being individually rational.

Having analysed two different markets (A+B) in detail, we now examine how the efficiency of our mechanism scales up over different markets (see figure 6). In order to do so, we find the average efficiency of markets as the number of buyers and the number of sellers are respectively varied from two¹⁶ to twenty. We run the auctions over 500 iterations with sellers randomly allocated their supply and buyers having a demand ranging from 1 upto the total supply divided by number of buyers. As can be seen, the average efficiency of the mechanism is maintained as the size of the market increases. The average efficiency ranges between 0.64 and 0.89 with no correlation to the market size¹⁷, which implies that it is unaffected by the size of the market i.e. the market scales.

¹⁶A minimum of two sellers and two buyers is required for a double auction.

¹⁷The correlation between average efficiency and number of sellers is 0.1 and the correlation between average efficiency and number of buyers is -0.05.

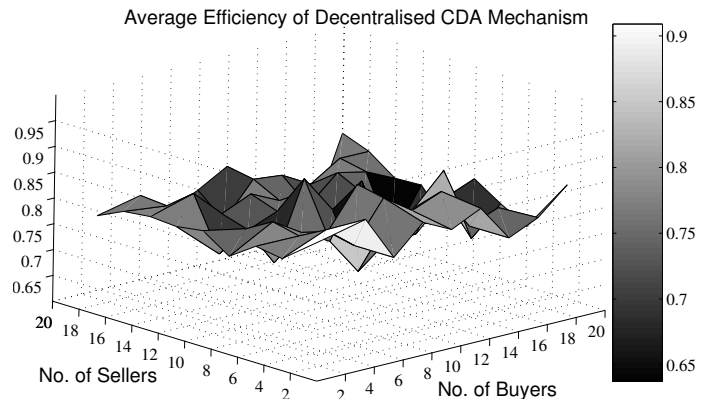


Fig. 6. Performance of Decentralised Mechanism in different markets with different number of buyers and sellers.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we have presented our work on the development of two complementary mechanisms for task allocation. We considered a scenario where production costs are characterised by a cost function composed of a fixed cost, a constant marginal cost and a limited capacity and where we were seeking the minimal total production cost that satisfies demand.

Specifically, in the first mechanism we extend the standard VCG mechanism to our problem domain in order to incentivise selfish agents to report truthfully about their types and thereby enabling the mechanism to find the efficient allocation. This required a novel penalty scheme to ensure that the mechanism is strategyproof for agents misreporting both their cost and their capacities. Individual rationality is conserved under this new mechanism and we show how this mechanism is robust to uncertainties in the capacities of the agents. We then presented a dynamic programming algorithm, that solves the task allocation problem of the centre in pseudo-polynomial time.

In the second mechanism, we extend the standard format of a CDA so as to develop a decentralised mechanism for resource allocation in the same context. We find that this mechanism has a high inherent average efficiency (over 86% in the examples we study) by testing it with a variant of the ZI strategy.

When taken together, we find that these mechanisms represent a trade-off in terms of efficiency and the decentralisation of a mechanism (in the examples we consider, the loss in efficiency can range from 0% to 50% depending on the demand and number of buyers and sellers in the market). However, both mechanisms still ensure that the participants derive a profit by joining the mechanism, thereby justifying their use with selfish agents.

As future work, we first intend to extend these mechanisms to deal with iterated allocations (i.e. ones in which new demand continuously appears) since in several of the cases we consider it is conceivable that the agents can observe and learn about the behaviours of other agents in the system. Our

centralised mechanism would still work in such situations if we consider myopic agents (i.e. agents that cannot strategise over more than one round of allocation [29]) since then these agents will not strategise over rounds. However, this assumption might be too restrictive in some settings. Also we wish to further investigate the link between task allocation protocols which are efficient and those that are robust (i.e. protocols in which it is highly likely that agents will fulfill their assigned task despite being uncertain about their capabilities when revealing their type). The link has been revealed here via the penalty scheme and the connection of the penalty scheme to a trust-based scheme has been discussed. However, a deeper study is required to formally establish the consequence of requiring robust mechanisms on the efficiency of the resultant mechanism. We believe that the hybrid approach combining trust and penalties would be a very interesting field to pursue. Finally, we aim to develop more sophisticated strategies for the decentralised mechanism in order to enhance the efficiency of the system, whilst ensuring that these sophisticated strategies derive higher profit than their simpler counterparts. This has been shown to be achievable in simple CDAs [6], [17], [43] and we believe it is also achievable in our modified CDA protocol. Such developments will enable us to more effectively find the set of agents who can perform the required task at the lowest cost (i.e. the efficiency will be increased).

Acknowledgements: Rajdeep K. Dash is funded by a BAE Systems studentship and an ORS scholarship. Perukrishnen Vytelingum and Esther David are sponsored by the DIF-DTC project (8.6) for Agent-Based Control Systems. Alex Rogers is supported by the ARGUS II DARP (Defence and Aerospace Research Partnership) project. The authors thanks the anonymous reviewers for their comments which improved the paper.

REFERENCES

- [1] Multiagent systems: Challenges and opportunities for decision-theoretic planning. *AI Magazine*, 20(4):35–43, 1999.
- [2] E. Adar and B. Huberman. Free riding on gnutella. *First Monday*, 5(10), October 2000.
- [3] G. R. Bitran and H. H. Yanasse. Computational complexity of the capacitated lot size problem. Technical Report Sloan WP No. 1271-81, Alfred P. Sloan School of Management, Massachusetts Institute of Technology, 1981.
- [4] J. Bredin, D. Kotz, and D. Rus. Market-based resource control for mobile agents. In *Proceedings of the Second International Conference on Autonomous Agents (Agents '98)*, pages 197–204. ACM Press, 1998.
- [5] S. Clearwater, editor. *Market-Based Control- A paradigm for distributed resource allocation*. World Scientific, 1996.
- [6] D. Cliff and J. Bruten. Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical Report HPL-97-91, 1997.
- [7] S. E. Conry, K. Kuwabara, V. R. Lesser, and R. A. Meyer. Multistage negotiation for distributed constraint satisfaction. *IEEE Transactions on Systems, Man and Cybernetics*, 21(6):1462–1477, 1991.
- [8] R. K. Dash, D. C. Parkes, and N. R. Jennings. Computational mechanism design: A call to arms. *IEEE Intelligent Systems*, 18(6):40–47, 2003.
- [9] R. K. Dash, S. Ramchurn, and N. R. Jennings. Trust-based mechanism design. In *Proc. 3rd Int. Conf. on Autonomous Agents and Multi-Agent Systems*, pages 726–753, New York, USA, 2004.
- [10] E.H.Durfee and V.Lesser. Negotiating task decomposition and allocation using partial global planning. In L.Gasser and M.Huhns, editors, *Distributed Artificial Intelligence*, volume II, pages 229–244. Pitman Publishing : London and Morgan Kauffman: San Mateo, CA, 1989.
- [11] M. Eso, S. Ghosh, J. R. Kalagnanam, and L. Ladanyi. Bid evaluation in procurement auctions with piecewise linear supply curves. Technical Report RC 22219, IBM Research, Yorktown Heights, NY, 10598, 2001.
- [12] M. Florian, J. K. Lenstra, and H. G. Rinnooy Kan. Deterministic production planning: Algorithms and complexity. *Management Science*, 26:669–679, 1980.
- [13] D. Friedman and J. Rust, editors. *The Double Auction Market: Institutions, Theories and Evidence*. Addison-Wesley, New York, 1992.
- [14] M. R. Garey and D. S. Johnson. *Computers and Intractability - A Guide to the Theory of NP-Completeness*. Freeman, San Francisco, 1979.
- [15] B. P. Gerkey and M. J. Mataric. Sold!: Auction methods for multi-robot coordination. *IEEE Transactions on Robotics and Automation, Special Issue on Multi-robot Systems*, 18(5):758 – 768, October 2002.
- [16] A. Giovannucci, J. A. Rodriguez-Aguilar, A. Reyes-Moro, F. X. Noria, and J. Cerquides. Towards automated procurement via agent-aware support. In *Proc. 3rd Int. Conf. on Autonomous Agents and Multi-Agent Systems*, pages 244–251, New York, USA, 2004.
- [17] S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22:1–29, 1998.
- [18] D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.
- [19] M. He, H. F. Leung, and N. R. Jennings. A fuzzy logic based bidding strategy for autonomous agents in continuous double auctions. *IEEE Trans on Knowledge and Data Engineering*, 15 (6):1345–1363, 2003.
- [20] N. R. Jennings and S. Bussmann. Agent-based control systems. *IEEE Control Systems Magazine*, 23(3):61–74, 2003.
- [21] A. Kothari, D. C. Parkes, and S. Suri. Approximately-strategyproof and tractable multi-unit auctions. In *4th ACM Conf. on Electronic Commerce (EC'03)*, pages 166–175, 2003.
- [22] V. Krishna. *Auction Theory*. Academic Press, 2002.
- [23] V. R. Lesser and D. D. Corkill. Functionally accurate, cooperative distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 11(1):81–96, 1981.
- [24] S. Martello and P. Toth. *Knapsack Problems, Algorithms and Computer Implementations*. John Wiley and Sons Ltd, England, 1990.
- [25] A. MasColell, M. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [26] K. A. McCabe, S. J. Rassenti, and V. L. Smith. Designing call auction institution: Is double dutch the best? *The Economic Journal*, 102(410):9–23, 1992.
- [27] J. Nicolaisen, V. Petrov, and L. Tesfatsion. Market power and efficiency in a computational electricity market with discriminatory double-auction pricing. *IEEE Transactions on Evolutionary Computation*, 5(5):504–523, October 2001.
- [28] C. H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall, 1982.
- [29] D. Parkes. *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*. PhD thesis, University of Pennsylvania, May 2001.
- [30] R. Porter, A. Ronen, Y. Shoham, and M. Tennenholtz. Mechanism design with execution uncertainty. In *Proceedings of UAI-02*, pages 414–421, 2002.
- [31] K. Ramamritham, J. A. Stankovic, and W. Zhao. Distributed scheduling of tasks with deadlines and resource requirements. *IEEE Trans. Comput.*, 38(8):1110–1123, 1989.
- [32] A. Ronen. *Solving Optimisation Problems among Selfish Players*. PhD thesis, Institute of Computer Science, Hebrew University of Jerusalem, 2000.
- [33] J. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, 1994.
- [34] A. E. R. Roth and I. Erev. Learning in extensive form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behaviour*, 8:164–212, 1995.
- [35] M. H. Rothkopf, A. Pekec, and R. M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44:1131–1147, 1998.
- [36] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 1995.
- [37] T. Sandholm. emediator: A next generation electronic commerce server. *Computational Intelligence*, 18(4):656–676, 2002.
- [38] T. Sandholm. Making markets and democracy work: A story of incentives and computing. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1649–1671, 2003.
- [39] J. Shneidman and D. Parkes. Rationality and self-interest in peer to peer networks. In *2nd International Workshop on Peer to Peer Systems (IPTPS)*, pages 139–148, 2003.

- [40] V. L. Smith. An experimental study of competitive market behaviour. *Journal of Political Economy*, 70:111–137, 1962.
- [41] M. Tennenholtz. Some tractable combinatorial auctions. In *Proc. of the 17th National Conf. on Artificial Intelligence (AAAI/IAAI)*, pages 98–103, 2000.
- [42] G. Tesauro and J. L. Bredin. Strategic sequential bidding in auctions using dynamic programming. In *Proceedings of the First International Joint Conference on Autonomous Agents and MultiAgent Systems (AA-MAS)*, pages 591–598. ACM Press, 2002.
- [43] P. Vytelingum, R. K. Dash, E. David, and N. R. Jennings. A risk-based bidding strategy for continuous double auctions. In *Proc. 16th European Conference on Artificial Intelligence (ECAI)*, pages 79–83, 2004.
- [44] R. Wolski, J. S. Plank, J. Brevik, and T. Bryan. Analyzing market-based resource allocation strategies for the computational Grid. *The International Journal of High Performance Computing Applications*, 15(3):258–281, Fall 2001.
- [45] M. Xia, J. Stallaert, and A. Whinston. Solving the combinatorial double auction problem. *European Journal of Operational Research*, 164(1):239–251, 2004.
- [46] M. Yokoo, E.H.Durfee, T. Ishida, and K. Kuwabara. Distributed constraint satisfaction problems: Formalization and algorithms. *IEEE Trans. on Knowledge and Data Engineering*, 10(5):673–685, 1998.

APPENDIX

A TBM would differ from the current scheme in two main respects¹⁸:

- 1) An agent i would report an extra dimension, rep_i , which is its self-reputation that determines how often it succeeds in providing upto its reported capacity, c_i .
- 2) The centre determines the optimal allocation, K^* , from the reports of the agents (u_i, f_i, c_i, rep_i) , whereby the feasible allocation is determined by a combination of rep_i and c_i .

It is interesting to note that as a result of the way the payments are conditioned in the TBM, a penalty would still be applied if the agent does not produce its capacity but the penalty is not the same. In TBM, there is no incentive for an agent which has been over-optimistic of its capacity to produce the maximum it can. Also, TBM does not consider the case when the pdf from which the (expected) capacity of the agent is determined changes (e.g. seasonal changes affecting the capacity of a solar generator, or failures in components of a machine reducing the capacity of a job-shop machine). However, an interesting hybrid approach would be to have a specific δ_i conditioned on the reputation of each agent i . This would provide us the ability to condition the penalty of an agent dependent upon its past performance and not just its present performance. The investigation of such a mechanism is outside the scope of this paper and is left for future work.

¹⁸We here consider the simplification of TBM in which agents report about their own reputation. This is similar to the mechanism developed in [30].