

Multi-Attribute Auctions with Different Types of Attributes: Enacting Properties in Multi-Attribute Auctions

Albert Pla^a, Beatriz López^a, Javier Murillo^b, Nicolas Maudet^c

^a*Universitat de Girona, Spain*
{albert.pla,beatriz.lopez}@udg.edu

^b*Newronia, Spain*

javier.murillo@newronia.edu

^c*Université Pierre et Marie Curie (Paris-6), France*
nicolas.maudet@lip6.fr

Abstract

Multi-attribute auctions allow agents to sell and purchase goods and services taking into account more attributes besides the price (e.g. service time, tolerances, qualities, etc.). In this paper we analyze attributes involved during the auction process and propose to classify them between verifiable attributes, unverifiable attributes and auctioneer provided attributes. According to this classification we present VMA2, a new Vickrey-based reverse multi-attribute auction mechanism which, taking into account the different types of attributes involved in the auction, allows the auction customization in order to suit the auctioneer needs. On the one hand, the use of auctioneer provided attributes enables the inclusion of different auction concepts such as social welfare, trust or robustness whilst, on the other hand, the use of verifiable attributes guarantee truthful bidding. The paper exemplifies the behaviour of VMA2 describing how an egalitarian allocation can be achieved. The mechanism is then tested in a simulated manufacturing environment and compared with other existing auction allocation methods.

Keywords: Multi-agent systems, Mechanism design, Multiple criteria analysis, Supply chain management, Multi-attribute auctions

1. Introduction

Auctions are gaining relevance in today's economy. They are increasingly being used in fields such as industry procurement (Strecker, 2004), supply chain management (Sadeh and Sun, 2003), electronic advertisement (Fu et al., 2012), service allocation (Pla et al., 2012c; Tavana et Al., 2013) and electronic commerce (Shih et Al., 2007). For instance, using auctions, manufacturers can obtain feedstock from their providers at the best economic conditions, while popular web sites can sell their ad space to the marketing companies which offer the most profitable revenue. In an auction, the auctioneer (e.g. manufacturer) that wishes to buy some good (e.g. feedstock) announces in a call for proposals what he wants to sell. Interested agents, bid for selling the good at the price they consider to obtain some payoff. Then, the auctioneer clear the auction, by determining a winner, and then setting up the price of the good.

Auctioneers wish to optimize their revenue, and for that purposes, auction mechanism designers have to provide mechanisms that ensure incentive compatibility, meaning that bidders are encourage to provide bids with a price equal to their true value of the good (Myerson, 1979). Moreover, companies want to optimize their revenue at the long and mid-term more than in a single transaction. Revenue is directly related to the number of participants, it has been proved that the revenue of an auctioneer when k bidders are participating in an auction is at least as higher than with a subset of $k-1$ bidders (Murillo et al., 2008; Lematre et al., 2003). Thus, keeping the number of participants to an efficient number is something that auction designers should face in order to optimize auctioneer revenues. For doing so, global satisfaction and preferences of the auction participants should be represented, e.g. they can be aggregated as social welfare measures do (Chevaleyre et al., 2006). This aggregation can be conducted in different ways, as maximizing the sum of all of the revenues and payoffs (utilitarian), or following an egalitarian approach, so as pursuing a higher equity between agents revenues (Chevaleyre et al., 2006; Endriss et al., 2003).

Most of the auction research has been concentrated on incentive compatibility and social welfare of mechanism regarding the price of the goods. Other properties such as quality of service and delivering time, have been considered as side constraints on the auction process, but not in the core procedure of clearing the auction (Zhao et al., 2011). For example, in certain auctions, all the bids which do not achieve a minimum quality level are filtered out and discarded before determining the winner. Multi-attribute auctions have been designed to deal with different attributes in the winner determination process. In this case, the auctioneer is posed with the

problem of choosing among a set of Pareto-optimal solutions, a single one. For example, in a multi-attribute auction in which an auctioneer is wishing to externalize a task, the auctioneer can characterize the task according to different attributes, such as price and quality. The auctioneer desires the highest quality at a good price. He will receive different bids of different quality and price; it is reasonable to think that quality and price they are often conflicting criteria, and several combinations are equally optimal. Thus, the auctioneer should make a decision and pick a combination of attributes. Depending on how this selection is done, truthful bidding based exclusively on price is no longer valid. E.g. a bidder which is encouraged to provide its true value for the price, may provide a fake quality, resulting in harmful consequences for the auctioneer.

To preserve truthful bidding in scenarios where money it is not the only key issue (e.g. pay-per-click advertising) new auction mechanisms have been designed. An example of them are position auctions used in advertisement scenarios such as Google Ads (Varian, 2007) where auctions occur in a repetitive way among a set of bidders. Due to repetition, auctioneers are able to qualify bidders according to some attributes. When clearing the auction, the auctioneer unifies the qualification attributes and the received bids in order to determine the winner. By doing so, this kind of auctioneers guarantees incentive compatibility, nonetheless, the attributes this mechanism is dealing with are not manipulable by bidders as they are provided by the auctioneer itself.

In this paper we are going one step forward, by distinguishing different types of attributes regarding the good to be sold, verifiable and unverifiable attributes, in addition to attributes qualifying the buyers. Verifiable attributes are the ones that can be checked upon the reception of the good, as for example, its quality. Conversely, unverifiable attributes, cannot be checked due to its subjectiveness, as the price of the object (a good can be expensive or not depending on the utility provided to the buyer). The third kind of attributes, auctioneer provided attributes, concerns information about good itself but also regarding the bidders (e.g. the reliability of a bidder).

Thus, our first contribution is the classifying of attributes according to the ownership (bidder or auctioneer provided) and their types (verifiable or unverifiable). Second, and using attributes of this categorization we design a new multi-attribute auction mechanism that extends previous multi-attribute auctions taking ideas from position auctions. Our mechanism lets the auctioneer obtain some desirable properties as incentive compatibility, so that the bidders are encouraged to bid with the true properties they can pro-

vide (duration, quality of service) at the best price. Third, the consideration of attributes regarding the bidders' behaviour, enables the configuration of different social welfare approaches (utilitarian, egalitarian) and opens up a range of possibilities for the incorporation of other auction properties such as robustness. To illustrate and test the auction mechanism, we simulate an industrial environment where different agents auction services which must be carried out by external service providers while trying to obtain a service in a reasonable price, time and quality.

This paper is organized as follows: first we introduce some basic concepts regarding auctions. In section 3 we present a brief state of the art of multi-attribute auctions. In section 4 we analyze and classify the different attributes that can take part in a multi-attribute auctions. Based on this classification, in section 5 we present a Multi-attribute auction mechanism to enact different auction properties. Then, we empirically test the mechanism in a simulator to illustrate its working, to compare it with some baseline allocation methods. Finally, in Section 7 we present the conclusions of our work and we point some possible future lines of research.

2. Auctions Background

An auction is a method for buying and selling goods using a bid system in which the winner bids obtain the auctioned goods (Krishna, 2002). In manufacturing domains, as we are used to, auctions often follow a reverse schema: the auctioneer wish to buy a task to be done, and the bidders are selling their working capacity at a given price. For the sake of clarity, we will follow that approach along the paper.

According to that approach (reverse auction), the main steps of an auction are the following:

1. Auction set up: The auctioneer (a_0) sends a call of proposals to all of the participant agents ($a_1, \dots a_n$), in which the item(s) it to be bought (e.g. a task to be done) is described.
2. Bidding: bidders ($a_1, \dots a_n$) return a bid b_i to the auctioneer with the price which they are offering for the good being auctioned, (where i identifies the agent who submitted the bid). The price may be equal to the bidders valuation of the item, but not necessarily. The valuation which a bidder a_i gives to an item can be defined using a function $v_i(it)$. In an incentive compatible mechanism it is expected that bidders bid truthfully, so $v_i(it) = b_i$.




Figure 1: Auction classification according to bidding sides, sold resources, bid composition and the role of the participants.

3. Winner determination problem: the auctioneer decides the winner(s) that maximizes its utility (Lehman et al., 2006). The *utility* is the measurement of the satisfaction received by the participants of an auction, either the bidders or the auctioneers (Neumann and Morgenstern, 1944). Auctioneer’s utility is related to the *revenue*, $u_o(it, b_i) = v_o(it) - p$. He pays an amount of p for an item it when the winner bid is b_i , and the auctioneer has a value for the item, $v_o(it)$, related to its interest on getting it. Accordingly, the WDP can be formulated as:

$$\operatorname{argmax}_{i>0}(u_o(it, b_i)) \quad (1)$$

And, since the payment p is conditioned by b_i and it is constant, it can be simplified to:

$$\operatorname{argmin}_i(b_i) \quad (2)$$

4. Payment: the auctioneer pays the money to the winner(s). The utility or *profit* of the winner bidder is $u_i(it, b_i) = p - v_i(it)$, where b_i is the bid he submits, and p the payment received. Observe that p is not necessarily equal to b_i since it will depend on the payment mechanism established in the auction, as explained in the auction taxonomy below.

As Figure 1 shows, there are many auction types depending on several criteria (Parsons et al., 2011). First, when a participant can play the role of auctioneer and bidder at the same time, we are talking about double side auctions; conversely, in a one-side auction, a participant can act just as auctioneer or bidder. Second, the number of goods or services that are being sold in an auction determines if an auction is a single-item or a multi-item auction. Third, the number of attributes of the bid defines if an auction is a single-dimension (also known as uni-attribute) or a multi-attribute auction. On the one hand, in single-dimension auctions the bid is composed just by the price while, on the other hand, in multi-attribute auctions bidders provide other attributes beside the bid amount (e.g. time or quality). Fourth, if the

winner of an auction pays the exact amount it has provided in its bid, the auction is a first price auction; otherwise, if the price is conditioned by the non-winning bids, is a second price auction. Finally, and according to the role of the participants, an auction can be classified as forward or reverse. As stated above, whilst in forward auctions the auctioneer is selling goods or services and the bidders compete for them, in reverse auctions the auctioneer is the buyer and the bidders compete to sell their services.

In addition to that classification, auctions can be open-cry or sealed-bid (Bichler et al., 1999). In open-cry auctions, participants bid openly against each other, making their bids public and revealing their preferences to the rest of participants. In this mechanisms, agents privacy is sacrificed for the sake of transparency. A typical example of these kind of auctions are English and Dutch auctions. Conversely, in sealed-bid auctions, privacy of the agents is preserved since all bidders simultaneously submit sealed bids to the auctioneer. In this way, no bidder knows how much the other participants have bid.

A desirable property that an auction mechanism should provide is to ensure that bidders provide truthful bids (incentive compatible mechanism). This means that bidders obtain a better profit by revealing their real attributes than by providing other values or trying to cheat. Vickrey auctions are sealed bid, single-item, uni-dimensional, second-price auctions, which guarantees incentive compatibility (Vickrey, 1961). In this kind of auctions, the winner bid pays the second best price. Similarly, Vickrey-Clarke-Groves guarantees this property for multi-item auctions (Bartal et al., 2002).

An auction mechanism can be measured in different terms such as the auctioneers' revenue or social welfare. On the one hand, the highest the revenue, the better the auction. On the other hand, regarding social welfare, there is no unique measure, but any aggregation of individual agents preferences can be used for that purpose. Thus the utilitarian approach tries to maximize the sum of all of the individual utilities, that is:

$$maximize \left(\sum_{i=0}^n u_i(it) \right) \quad (3)$$

An egalitarian approach, could consists on maximizing the minimum utility.

$$maximize (min_i(u_i(it))) \quad (4)$$

Other egalitarian approaches could be defined upon the minimization of the difference between the maximum and minimum utilities (Equation 5) or the minimization of the sum of difference between agents utilities and the mean

utility (Equation 6).

$$\text{minimize } (\max_i(u_i(it)) - \min_i(u_i(it))) \quad (5)$$

$$\text{minimize } \left(\sum_{i=0}^n \left| \frac{\sum_{j=0}^n u_j(it)}{n} - u_i(it) \right| \right) \quad (6)$$

In general, fair mechanisms maintain agents interested in the auction after running several auctions, and maintaining the number of participants increases the revenue of the auctioneer at the mid and long term, (see, for example, Murillo et al. (2012)’s work for an experimental approach).

3. Related Work

Multi-attribute auctions have been used in the electronic advertisement markets (Athey and Ellison, 2011), e.g. Varian (2007) proposes to include an extra attribute provided by the auctioneer itself, not by the bidder. This approach is similar to trust-based approaches as Ramchurn et al. (2009) where, to enhance robust allocations, agents describe other agents reliability using a trust parameter and including this parameter in the winner determination phase. In the same line, in priority auctions (PA) (Murillo et al., 2012) the auctioneer uses a priority attribute in order to obtain egalitarian allocations. In our approach, attributes may come from the auctioneer but bidders can also express quality values in multiple attributes, so, in some way, our proposal could be considered as an extension of these kinds of auctions (we propose to use one or more attributes incorporated by the auctioneer but also multiple attributes by the bidders side).

In multi-attribute auctions, a key work is Che (1993). In it, the author describes different scenarios regarding the payment rule and demonstrate that, in order to be incentive compatible, the evaluation of the payment attributes should match the evaluation of the second best bid, but not necessarily with the same combination of attributes. David et al. (2002) also follows aggregations and evaluations of Che (1993) but using a first-price sealed bid auction. Conversely to ours, we follow a second-price auction in order to obtain an incentive compatible mechanism.

In a posterior work, Parkes and Kalagnanam (2005) propose an adaptation of the VCG algorithm (MacKie-Mason and Varian, 1994) for multi-attribute auctions under an iterative schema (similar to English auctions). Meaning that bidders are allowed to modify their bids in response to the bids from other agents. Similar to this work, a framework to develop different kinds of English auctions under a non-linear multi-criteria preference

schema is presented in Bellosta et al. (2008, 2011). In our approach, we do not allow iteration as we want to achieve a fast allocation with a controlled communication load. These kind of approaches are more suitable for long-term contracts with companies, but not at the operational level as we do.

Mahr and de Weerd (2006) present a multi-attribute auctioning mechanism where agents use preference orders to express their bids. In a later work, Harrenstein et al. (2009) generalize the approach on what is known Qualitative Vickrey auctions. However, the authors focus on defining a mechanism so that bidders can express qualitative preferences. We are interested in having a set of different attributes per bid. Thus, handling different attributes increases the number of combinations to manage in a preference order solution. Finally, Suyama and Yokoo (2004) propose a multi-attribute mechanism where the attributes vary depending on the resource bundle assigned to tasks instead of the resource by itself. This is a complementary approach towards we need to extend our work in a future research as currently we are only considering a constant set of attributes.

Other approaches consider attributes different than price as constraints. E.g., Zhao et al. (2011) present a mechanism for auctions with temporal constraints based on VCG with a new payment method. Time constraints are used to filter the participating bids, but time is not considered when evaluating the bids, leaving aside whether time improves a bid or not. With our method, we take into account all of the attributes (thus also time) to optimize the allocation.

Auctions and multi-attribute auctions have been widely used to allocate resources in workflow environments (Prodan et al., 2011; Puustjärvi, 2005; Zhang et al., 2010). E.g, Pla et al. (2011) compare the efficiency of resource allocation in workflows when auctioning different types of attributes, however, the different attributes are not merged or mixed in any stage of the process. In this line, a framework for the use of reverse multi-attribute auctions for resource allocation in business processes is proposed (Talluri and Ragatz, 2004). Following the work of these researches, we illustrate and test our mechanism in the workflow management domain (Pla et al., 2012a).

4. Attribute Typologies in Multi-attribute Auctions

In multi-attribute auctions, the item which is sold is defined due to a set of attributes, in addition to the price, that it is determined in the auction process. Multi-attribute auctions may involve many attributes regarding the item, as for example, the size and quality requested by the buyers, and the price offered by the sellers. Moreover, in some uni-attribute auctions

the auction is cleared by taking into account other attributes besides the ones offered by the bidder. For instance, information about bidders' past performance Ramchurn et al. (2009) (reliable or not), previous auction results Murillo et al. (2012) (the number of auctions which an agent has won) or opinions which the auctioneer has in regard to bidders Varian (2007) can be taken into account when deciding the forthcoming winners.

In this section we consider all of the above attributes, as they all influence the winner determination problem and the payment mechanism of multi-attribute auctions. In doing so we distinguish two main criteria: attribute *ownership*, and *verifiability*. Ownership means that attributes can be characterized according to the information source (the auctioneer itself or the participating bidders). On the other hand verifiability concerns the capability of an attribute to be verified by an agent.

Regarding ownership, we can classify attributes as auctioneer-provided or bidder-provided:

- **Auctioneer-provided attributes:** This is the set of attributes which the auctioneer uses to extend the bidder's bid. They can be used to qualify bidders or their bids on behalf of a given objective criteria. These attributes can involve: number of withdraws on behalf of a given winner, number of times the bidder has won an auction, number of times the bidder reported a lower quality than the one it bid for, etc. Moreover, these attributes can be used to express the subjective auctioneer beliefs regarding bidders or their bids. For instance, they can be used to define reputations or how popular a bidder is according to the auctioneer's perspective. Since it can be assumed that an auctioneer will not try to deceive itself, we can say that auctioneer-provided attributes are trustable and the auction mechanism does not need to concern about their reliability. Those attributes can then be used to include fairness or other properties in the auction mechanism.
- **Bidder-provided attributes:** they are the set of attributes that bidders express regarding the qualification of the good to be served.

Bidder-provided attributes can be, in turn, classified according to the capability of the auctioneer to verify them. We distinguish two kinds of attributes :

- **Unverifiable bidder attributes:** These are the set of attributes defined by an agent whose true values are only known by the agent itself. These attributes are also the ones which define the auction currency. A typical example of this kind of attribute is the economic value which a

bidder offers or asks to obtain an item or for providing a service. The bidder knows its true value, however the auctioneer has no way to know the true value of the attribute neither before nor after delivering the auctioned item. Similarly, bidders could trade permits for generating certain amounts of CO_2 emissions (e.g. as happens with emissions trading between countries) Hoesch-Klohe et al. (2010).

Despite the fact that there can be more than one unverifiable attribute, auctions are typically designed using only one unverifiable attribute due to the complexity of introducing more than one. In cases where there is more than one currency attribute Rackson et al. (2002) (e.g. international auctions), all the attributes are translated into a unique attribute or currency so the auctioneer only deals with one type of unverifiable attribute (for example a currency which acts as a standard monetary unit or a virtual currency). Therefore, despite the existence of more than one unverifiable attribute during the bidding process, the auction is performed considering just one single unverifiable attribute. From this point of view, in this paper we assume that there can only be one currency in an auction and when we refer to an unverifiable attribute we mean the unique attribute which acts as currency for the auction (namely the attribute itself or the virtual currency which aggregates more than one unverifiable attributes).

- **Verifiable bidder attributes:** these are the set of attributes which are defined by agents whose true value can be known and checked by another agent. It does not matter if they can be verified before or after delivering the auctioned item, however, this verification must always be completed before the payment is performed. Examples of this type of attribute are delivery times, electricity consumption or other physical specifications. For example when auctioning a task which will be carried out by a bidder, the bidder can specify a certain delivery time t . Once the task has been completed the auctioneer can then check if the final delivery time t' was as specified during the bidding process. As these attributes can be checked, they can be used to adjust payment to bidders, to evaluate if the item delivery has been successful or to establish parameters to describe bidder qualities (auctioneer-provided attributes).

In Figure 2 we illustrate a simple multi-attribute auction in which the three different types of attributes are used. Auctioneer A calls an auction in order to find an agent to perform a certain task. Bidders send bids containing the economic cost they expect to charge for the task (attribute b) and the delivery time they propose (dt). Once the bidders have sent their proposals, A will include a reliability attribute r which rates his satisfaction regarding

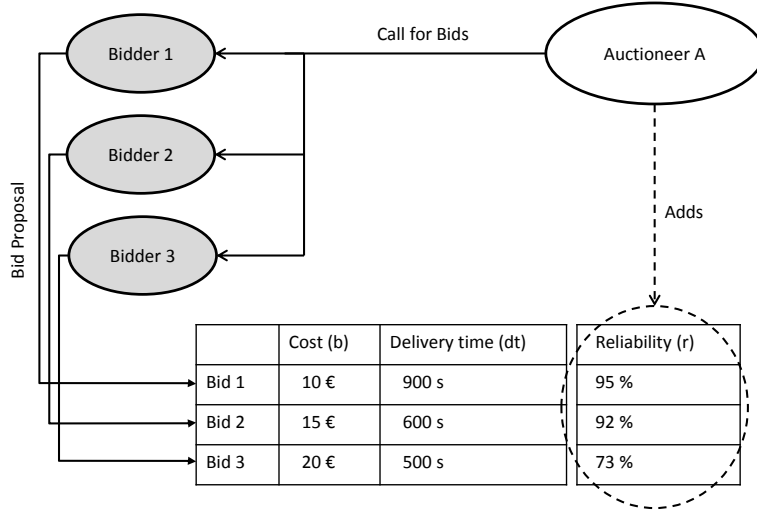


Figure 2: Example of the attribute types used in the bidding process: Cost is a unverifiable bidder provider attribute, delivery time a verifiable bidder provided attribute and reliability is an auctioneer provided attribute.

previous deals with the different bidders. Finally, the winner of the auction is computed using every attribute involved in the auction process (b, dt, r) . In this example b is a unverifiable bidder attribute as the auctioneer cannot know the true value of the cost attribute for each bid, at most, it can estimate it; dt is a verifiable bidder attribute as, once the task is finished, the auctioneer can compare the real delivery time dt' with the one which was provided in the bid; finally, r is an auctioneer-provided attribute as it is added into the bid by the auctioneer itself and it is also used to determine the winner of the auction.

5. VMA2: Multi-attribute auction based on verifiable, unverifiable, and auctioneer-provided attributes

Using the classification of attributes provided in the previous section we present the VMA2 multi-attribute auction mechanism. VMA2 follows a Vickrey-based approach to guarantee the incentive compatibility of the mechanism (Pla et al., 2013).

Let be (a_1, \dots, a_n) the bidders, and a_0 the auctioneer, who wants to buy an item it . This time, however, the auctioneer in addition to the price, requires additional attributes from the item, as for example, price and quality. Each agent provides one bid according to their expected utility, b_i . The bid is multi-dimensional, including unverifiable attributes A_i^u , as well as veri-

fiable ones, A_i^v . The auctioneer, solves the winner determination problem with modified bids, b'_i , so that each original bid b_i has been extended with attributes the auctioneer has about the bidder A_i^p , $b'_i = b_i \oplus A_i^p$ (auctioneer-provided attributes). Then, the problem the auctioneer is faced is to find the value that maximize its utility. The questions to be answered include: how bidders deal with multidimensional item descriptions; is the revenue meaning only the cheapest price? Moreover, and regarding the payment mechanism, what happens if the agents cheat on the verifiable attributes (e.g. the delivery time)? To answer to all of these questions, the different steps of the auction process should be revised, taking into account the nature of the attributes involved.

We assume that auctions are repeated along time there are no externalities in the process, as well as no budget constraints by any agent and that bids are presented under sealed bid (bids presented by the agents are unknown for the rest of the auction participants and are not made public at any time). Regarding attributes, we adopt the following notation:

- Auctioneer provided attributes A^p : $at_1^p \in D_1^p, \dots, at_{m_p}^p \in D_{m_p}^p$
- Verifiable (bidder provided) attributes A^v : $at_1^v \in D_1^v, \dots, at_{m_v}^v \in D_{m_v}^v$
- Unverifiable (bidder provided) attributes A^u : $at^u \in D^u$

5.1. Request for proposal - The role of verifiable attributes

Once an agent wants to buy an item, it summons an auction. An item is not necessarily a physical object. For example, it can be a service externalization, or the need of a resource to deploy a task. The auctioneer tries to obtain the best option at the best possible price. For that purpose, the agent builds a call for proposal *cfp* defining the item to be purchased to the bidders.

$$cfp = \langle taskId, (at_1^v, R_1^v), \dots, (at_{m_v}^v, R_{m_v}^v) \rangle \quad (7)$$

where at_i^v is the id of a verifiable attribute, and $R_i^v \subseteq D_i^v$ the projection of the attribute original domain and m_v the number of verifiable attributes involved in the auction. For example, if at_1^v is the delivery time, defined in \mathbb{N} , with 0 the lowest value (earlier time), a call for proposal could contain (*deliveryTime*, [2..5]) meaning that the auctioneer is requesting a task to be done between 2 and 5 hours. Of course, attribute interpretation should be supported by an ontology, that is out of the scope of this paper.

5.2. Bidding - The role of unverifiable attributes

When an agent a_i receives a call for proposals cfp it analyzes if he is interested in participating in the auction. For that purpose he determines its capacity of attending a request $A_i^v \subseteq R_1^v \times R_{m_v}^v$, with $A_i^v = (at_{1i}^v, \dots, at_{m_v i}^v)$. Since the auction follows a sealed bid schema and the bidder does not know which bids his contenders offered in past auction, the bidder is only concerned about his offer. In the case the bidder is interested in participating, he needs to return a bid he considers it will maximize its expected utility and which includes unverifiable attributes (usually the price he requires).

The true value of the unverifiable attribute is conditioned by the rest of the attributes which the bidders offer, corresponding to the valuation which the bidder makes of the verifiable attributes it bids: $v_i(A_i^v) = A_i^u$. Thus, the resulting bid structure is as follows:

$$b_i = A_i^u \oplus A_i^v = (at_i^u, at_{1i}^v, \dots, at_{m_v i}^v) \quad (8)$$

Cheating agents provide bids with different values than their true valuations. For example, with a bid $b_{i'} = (at_i^{u'}, at_{1i}^{v'}, \dots, at_{m_v i}^{v'})$ where $A_i^v \neq A_i^{v'}$ or $A_i^{u'} \neq v_i(A_i^v)$, the agent could be offering better attributes than its skills or a lower price¹, with the aim of deceiving the mechanism and winning the auction with better economical conditions.

5.3. Winner determination - The role of auctioneer provided attributes

Upon the reception of the bids, the auctioneer extends them with the information recorded about the bidders, $b'_i = b_i \oplus A_i^p$, where A_i^p is the auctioneer provided attributes, $at_{1i}^p, \dots, at_{m_p i}^p$. A_i^p may be used to represent auctioneer opinions (such as trust in the bidder) or to add certain properties to the auction (e.g., fairness can be obtained by the use of a priority attribute based on previous auctions outcomes). As a result, a modified bid which includes the auctioneer provided attributes is obtained:

$$b'_i = (at_i^u, at_{1i}^v, \dots, at_{m_v i}^v, at_{1i}^p, \dots, at_{m_p i}^p) \quad (9)$$

The attributes used to extend the bid will affect the characteristics of the auction and the resulting allocations:

- **Egalitarian social welfare:** A priority attribute which defines bidder's history (number of auctions won, lost and participated) can be used in order to enhance the equity and fairness of the allocations. This type of

¹We assume that bidders are able to estimate their skills accurately

allocation can be useful to keep bidders interested in future auctions and to avoid recurrent auction problems such as the bidders drop problem Murillo et al. (2012); Lematre et al. (2003).

For instance, the auctioneer can assign a priority attribute $w_i \in [0, 1]$ to each bidder, according to the ratio between the number of auctions in which they have participated, and the auctions they have lost within a specific time window (Equation 10). The higher the ratio of lost auctions, the higher the priority of the bidder, meaning that a bidder with high priority should be awarded a soon to prevent the agent from leaving the market due to too low incomes.

$$w_i = 1 - \frac{1 + \text{auctions_won}_i}{1 + \text{auctions_participated}_i} \quad (10)$$

The method used to compute bidders' priorities can affect the outcomes of the auction, as it will condition who the winner of the auction is. Moreover, the use of a priority attribute can condition the arrival of new bidders as newcomers may have less options of winning an auction than bidders that are already inside the market. For instance, the approach commented above would allow the arrival of new bidders to the market, however, they would have a lower priority than bidders with a bad ratio of won/participated auctions; the auctioneer would consider that newcomers are not as susceptible to leave the auction as a bidder with a bad auction record. In this situation a newcomer would have the same priority than an auction which had won all the auctions in which it had participated or than an auction who decided to stay off the last auctions.

- **Robustness:** The inclusion of a trust attribute $\tau_i \in (0, 1]$ defining the reliability of the bidder will increase the robustness of the allocations. τ may express the relation between the number of auctions won and the success of delivering the won items, in this way, not reliable bidders would decrease their chances to win future auctions. The success of the delivery could be measured upon the fulfilment of the verifiable attributes provided in the winner bid (e.g. delivering an item later than agreed would be considered a failure whilst providing it on time would be considered as a success).

It is important to take into account that the use of auctioneer provided attributes may also affect the bidders' behaviour. For instance, the use of a trust attribute can incentive bidders to improve their accuracy when estimating their delivery times but an abuse of this attribute may lead certain bidders to abandon the auction due an unexpected estimation error at the

early stage of the allocation process. Thus, we recommend mechanism designers to do a deep study of what an attribute can imply before adding it to the auction mechanism.

Once the different bids have been extended, the auctioneer needs to determine the winner bid that maximize its expected utility. In single dimensional auctions, the utility of the auctioneer u_0 depends on the item valuation v_0 and the payment performed p (see section2). However, in the multi-dimensional case, there are several possible combinations $at_{1i}^v, \dots, at_{m_v i}^v \in R_1^v \times R_{m_v}^v$ which are feasible. Particularly, the utility can be defined on the basis of the Nadir point (the worst case scenario), $A_w^v = at_{1w}^v, \dots, at_{m_v w}^v$, and any bid b'_i should improve this situation. For example, the worst situation could be defined as the latest time of a task which an agent can accept. For that purpose, we need to define an evaluation function on bids, $f_0(b'_i)$, that provides a real value, such that $u_0(cfp, b'_i) = v_0(A_w^v) - f_0(b'_i)$. The highest the utility, the lowest f_0 . Thus, the evaluation function value must decrease when the bid solution is better.

The evaluation function f_0 should represent the different auctioneers preferences (different auctioneers may have different evaluation functions). For example, low price (unverifiable attribute), high quality (verifiable attribute) and high priority (auctioneer provided attribute) should be combined. For that purpose, a multicriteria function that fulfill a set of conditions within the range of the attributes could be used (Pla et al., 2012b). Such a function, should fulfill certain conditions. First, the functions used for the evaluation must be a **real-valued function**. Given a bid, the evaluation function must return a real number in order to compute the payment as stated above. Second, the evaluation function must be **monotonic**. If one of the attributes of a bid is improved, the result of the evaluation function will change consequently, guarantying that a better bid will not obtain a worse evaluation. Finally, in order to allow the mechanism to calculate the payment, the evaluation function must have a **bijective** behavior regarding the auction currency. Taking this into account, we require f_0 to have a partial anti function f_0^{-1} which, given the result of the evaluation function, the auctioneer provided attributes and the verifiable attributes, returns the value of the unverifiable bidder attribute. Thus, if $f_0(b'_i) = x$, the partial anti-function f_0^{-1} will return the unverifiable attribute which acts as currency, $f_0^{-1}(x, at_{1i}^v, \dots, at_{m_v i}^v, at_{1i}^p, \dots, at_{m_p i}^p) = at_i^u$. Some of the evaluation functions which meet these requirements are the product, some mathematical norms such as the Euclidean norm, the weighted sum or the weighted sum of functions (Pla et al., 2012b).

Given the previously defined utility function of the auctioneer is $u_0(cfp, b'_i) =$

$v_0(A_w^v) - f_0(b'_i)$, the goal of the auctioneer is to minimize f_0 :

$$\operatorname{argmin}_i(f_0(b'_i)) \quad (11)$$

The auctioneer clears the market by sorting the bids ($b'_{\sigma(i)}$ where $\sigma(i)$ indicates the rank of the bid) from the best evaluation (the one with the lowest value) to the worst. The winner of the auction is the first in the ranking: b'_1 (and b'_2 the second best bid).

5.4. Payment - Playing all the attributes together

The payment mechanism is inspired in Vickrey and position auctions but taking into account the peculiarities of reverse multiattribute auctions. In Vickrey auctions, the winner bid receives just the necessary amount to beat the second highest bid (Equation 13). In other words, the payment the winner receives is the price it should have bid to obtain the same evaluation as the second highest bid. The winner bid is $(at_1^u, at_1^v, \dots, at_{m_v}^v, at_1^p, \dots, at_{m_p}^p)$ with a valuation of $f_0(at_1^u, at_1^v, \dots, at_{m_v}^v, at_1^p, \dots, at_{m_p}^p)$. In order to simplify the comprehension of the payment mechanism and to simplify notation, we will use A_i^v and A_i^p to refer to the the bid components regarding bidder verifiable attributes and auctioneer provided attributes, being:

$$f_0(b'_i) = f_0(at_i^u, at_{1i}^v, \dots, at_{m_v i}^v, at_{1i}^p, \dots, at_{m_p i}^p) = f_0(\langle at_i^u, A_i^v, A_i^p \rangle) \quad (12)$$

Thus, the payment (p) to the winner should equal the offer provided by the second best bid according to the auctioneer valuation, $f_0(b'_2)$. That poses the problem of finding the value of p which, combined with the attributes of the winning bid, equals the valuation of the second best bid; that is

$$f_0(\langle p, A_1^v, A_1^p \rangle) = f_0(b'_2) \quad (13)$$

Therefore, the payment is obtained by means of the antifunction f_0^{-1} :

$$p = f_0^{-1}(f_0(b'_2), A_1^v, A_1^p) \quad (14)$$

Where p is the payment which will be received by the single winner of our mechanism.

However, this payment mechanism does not prevent the bidders to lie regarding their verifiable attributes since including a false attribute could increase the chances to win the auction while not being penalized in the payment. For example, a bidder could submit a bid saying that it will finish its task in 10 minutes when actually it knows it will finish the task in 15 minutes. This lie would have increased the chances of the bidder to win the auction.

Thus we adapt the payment mechanism in order to penalize dishonest bidders: when bidders intentionally lie to win the auction, payment is obtained by equating the initial bid evaluation (with the false value of the attribute) to the assessment of real goods (with the true value of the attribute). That is, since the payment is performed after a task or service has been completed or after an item has been delivered, the auctioneer can measure the quality of the received item and realize that instead of receiving the agreed product it got one with characteristics $A_i^{v'} = at_{1i}^{v'}, \dots, at_{m_i}^{v'}$. When this situation occurs, the payment mechanism is based on making the valuation of the obtained qualities same as the bid ones; in other words, we maintain the utility of the auctioneer.

$$f_0(\langle p, A_1^{v'}, A_1^p \rangle) = f_0(b'_1) \quad (15)$$

$$p = f_0^{-1}(f_0(b'_1), A_1^{v'}, A_1^p) \quad (16)$$

Where $A^{v'}$ is the vector of verifiable attributes obtained by the auctioneer after the delivery process (and that we compare with the true attributes of the bidder).

In regard of the auctioneer provided attributes, the mechanism does not have to be concerned about them inasmuch they express a set of beliefs regarding the bidder and the auctioneer itself during the auction process, not the delivered item. Thus, they cannot be changed during the delivery as they concern to a past time instant.

Summing up, our payment proposal is a two case method:

$$p = \begin{cases} f_0^{-1}(f_0(b'_2), A_1^v, A_1^p) & \text{if } A_1^{v'} \succeq A_1^v \\ f_0^{-1}(f_0(b'_1), A_1^{v'}, A_1^p) & \text{if } A_1^{v'} \prec A_1^v \end{cases} \quad (17)$$

where operator \succeq means *the same or better than* and \prec *worse than*.

In this way, the bidder is encouraged to bid truthfully: regarding unverifiable attributes, overbidding attributes will reduce his chances of winning the auction without increasing its payment whilst underbidding will increase its risk of working below its real price (resulting in a negative utility). Respecting verifiable attributes, the bidder is also encouraged to bid truthfully: if the bidder provides its service or good as it was indicated in the bid, he will receive the economical amount that he was expecting or even more (the exact value to beat the following highest bid), increasing its utility. However, if the bidder lies and delivers its product in worst conditions than the ones agreed it will not receive the payment he expected, reducing its utility. Extended studies regarding the strategy proofness can be found in (Pla et al., 2013).

5.5. A simple example

To illustrate the behaviour of the VMA2 we present a simple example in which an auctioneer needs to externalize a service S_0 which must be finished before a certain time t_0 and must fulfill a minimum quality q_0 . As the auctioneer must accomplish other services which depended on S_0 , the earlier S_0 is finished, the higher will be its satisfaction or utility as it will have more time to dedicate to the rest of the tasks (e.g. hiring slower but cheaper providers). In the same way, a higher quality and a lower price will increase the auctioneer's utility.

In the first step the auctioneer would provide a call for proposals defining the minimum required quality $q_0 = 1$ and the maximum delivery time $t_0 = 100$, where q_0 and t_0 are verifiable attributes (A_0^v). Delivering the service too early (before 60 minutes) may cause inventory problems, whilst delivering the task too late may difficult the development of depending tasks, therefore. Therefore, the auctioneer defines the domain of the verifiable attributes as $q \in [1, 2]$ and $t \in (60, 100]$. The auctioneer values the worst case allocation q_w, t_w with a score of 8000 ($v_0(q_w, t_w) = 8000$). Thus, the resulting call for proposals is the following:

$$cfp = \langle S_0, (q, [1, 2]), (t, [60, 100]) \rangle \quad (18)$$

When a bidder receives the cfp , it checks if he is able to perform the desired task with the asked or better conditions. If so, he calculates the attributes he can provide and the cost it will result from them and offers his bid b_i . E.g, if a bidder a_a is able to perform the desired task in 80 minutes ($t_a = 80$) with a quality of 1.2 ($q_a = 1.2$) with an economic cost (ec) of 100 he would provide a bid $b_a = (100, 1.2, 80)$ where 100 is the economic cost of the bid, an unverifiable attribute. In the case that a_a tried to win the auction by cheating, he could have submitted a bid $(90, 1.2, 75)$, $(100, 1.5, 25)$ or any other bid with false attributes.

The auctioneer could receive other bids from other agents (a_b, a_c , see Table 1). Next, it extends the bids with the information concerning bidders (auctioneer provided attributes A^p) resulting in a new extended bid b'_i . In this case, the auctioneer just adds a priority w based on past auctions results in order to achieve an egalitarian social welfare. Then, it decides the winner of the auction using the an evaluation function which derives from his utility function, in this case a product of all the attributes (Equation 19)

$$f_o(b'_i) = f_0(\langle ec_i, q_i, t_i, w_i \rangle) = ec_i * \frac{1}{q_i} * t_i * w_i \quad (19)$$

The auctioneer clears the auction by selecting the bid which obtains the best evaluation, in this case, bidder a_a .

		A^u		A^v		A^p		
B_i	b_i	ec_i	q_i	t_i	w_i	b'_i	$f(b'_i)$	rank
a_a	$((100),(1.2,80))$	100	1.2	80	0.8	$(100,1.2,80,0.8)$	5333,33	1
a_b	$((95),(1.2,90))$	95	1.2	90	0.8	$(95,1.2,90,0.8)$	5700,00	3
a_c	$((100),(1.3,85))$	100	1.3	85	0.85	$(100,1.3,85,0.85)$	5557,69	2

Table 1: List of bids and their corresponding ranks and evaluations.

Finally, when the bidder finishes S_0 , the auctioneer evaluates the obtained result and proceeds to pay the bidder. As the payment mechanism considers two situations, we illustrate two different payment scenarios: when the delivered items meet the auction agreement and when the bidder breaks it.

- In the case bidder a_a finishes S_0 with t_a and q_a , it will receive the necessary amount it should have bid to beat the second best bid (bidder a_c). So the payment would be:

$$f_0(\langle p, q_a, t_a, w_a \rangle) = f_0(\langle ec_c, q_c, t_c, w_c \rangle) \quad (20)$$

$$p = \frac{100 * \frac{1}{1.3} * 85 * 0.85}{\frac{1}{1.2} * 80 * 0.8} \quad (21)$$

$$p = 104.21 \quad (22)$$

In consequence, the utility of the auctioneer is:

$$v_0(q_w, t_w) - f_0(104.21, 1.2, 80, 0.8) \quad (23)$$

$$u_0 = 8000 - 5557,87 = 2442,13 \quad (24)$$

- In the other hand, if A finishes S_0 with $t'_a \prec t_a$ (delayed²) or $q'_a \prec q_a$, it will receive the necessary amount it should have bid to achieve the same evaluation he obtained with the original bid. To illustrate that, we will consider $t'_a = 90$ and $q'_a = 1.15$.

$$f_0(\langle p, q'_a, t'_a, w_a \rangle) = f_0(\langle ec_a, q_a, t_a, w_a \rangle) \quad (25)$$

$$p = \frac{100 * \frac{1}{1.2} * 80 * 0.8}{\frac{1}{1.15} * 90 * 0.8} \quad (26)$$

²notice that \prec means *worse than* (see Equation 17)

$$p = 85.19 \tag{27}$$

In consequence, the utility of the auctioneer is:

$$v_0(q_w, t_w) = f_0(85.19, 1.15, 90, 0.8) \tag{28}$$

$$u_0 = 8000 - 5333, 33 = 2666, 66 \tag{29}$$

In this way, it can be seen how the utility of the auctioneer does not decrease when the bidder fails in delivering the agreed service.

6. Experimentation

To test and illustrate the behaviour of the presented mechanism, we simulated the processes and the services required by a company during a certain time period.

6.1. Simulation Environment

The simulation tool used to test the mechanisms described reproduces the operation of an industrial environment where companies need to solve different issues ? (e.g. unforeseen faults, local system crashes, customer problems, etc.) by outsourcing certain tasks. When an incident is detected, a preliminary diagnosis process determines the incident typology, the task (or set of tasks) which must be carried in order to solve the incidence and its maximum deadline. It, then, determines the order in which the tasks must be performed and the tasks which need to be assigned to an external provider.

In the simulation, each entity (which can correspond to a company, a team, a department, etc.) is represented by an agent which acts selfishly, working solely for its own benefit. The participants in the simulation can play two different roles: they can act as service agent *SA* (an agent which controls a workflow and needs to externalize a task) or as a resource provider agent *RP* which wants to perform a task for a payment. Despite the fact an entity can perform both roles (but not within the same auction), for the sake of simplicity and to make the experiments more comprehensible, agents can take only one role during the simulation. Thus, service agents will always act as the auctioneers whilst resource provider agents will always be bidders (Figure 3).

The data used to model the simulation has been extracted from the records of an industrial organization ³. These records contain information

³Due to confidentiality agreements we cannot reveal the name of the company

regarding unexpected issues (unexpected faults, customer complaints, revisions, etc.) the company has had to deal with. The data was collected over two years and is composed of information concerning the tasks which were performed in order to solve these issues (the type of task performed, the emergency level of each task, tasks deadlines, the partner or team which developed each task, task durations, etc.). Using this records we have estimated the probability distribution function that represents the task occurrence during an average day. The same has been done to model the urgency parameter distribution (which defines the deadline of each task) and the task service durations of each resource provider (each resource provider has different time distributions for each kind of task). These distributions are then used to execute the simulations⁴.

In the simulation there are 7 service agents are in charge of managing different types of tasks (one type per agent) and 8 resource providers which compete to perform those tasks (each resource provider can perform different tasks according to Table 2). The agents which take part in the simulation are defined as follows:

- **Service Agents (SA):** When a service agent detects an incidence, a preliminary diagnosis process determines its typology, its maximum delivery time and other possible attributes required. Then, the agent calls for an an auction in order to outsource the task which will resolve the incidence. Each task is defined by its typology (which determines the type of resources required to perform the task) and its delivery time.

Service agents are characterized by:

- The type of the service task of which it is in charge (1 to 7). Each type of task can only be performed by certain types of resource provider (see Table 2).
- The probability distribution function that defines the task occurrence. In the simulation, tasks appear following this probability function.
- The probability distribution function that defines the emergency level of the tasks. Every time the service agent wishes to outsource a task it assigns an emergency level to it according to its emergency parameter distribution function. The emergency level

⁴The specific probability distributions for each type of resource and task can be found in Torrent and Pla (2013).

	S1	S2	S3	S4	S5	S6	S7		S1	S2	S3	S4	S5	S6	S7
RP1	✓	✓	✓	✓	✓	✓	✓	RP5			✓		✓		✓
RP2	✓	✓	✓	✓	✓	✓	✓	RP6					✓		
RP3	✓	✓	✓	✓	✓	✓	✓	RP7		✓	✓		✓		✓
RP4	✓		✓	✓	✓	✓	✓	RP8		✓	✓		✓		

Table 2: Capacity of a resource provider (RP) to develop a service (S).

Emergency Level	1	2	3	4
Maximum duration (minutes)	30	120	300	600

Table 3: Relationship between the urgency of a service and its maximum execution time.

is used to establish the execution deadline for each task (see Table 3).

- **Resource Provider Agents (RP):** When a resource provider receives a call for proposals it evaluates whether it can perform the task required (given its capacity and type) and its true values (its cost and expected delivery time). With this information the agent can decide if it wants to participate into the auction and defines the bid to be submitted.

Resource providers are defined by:

- The type of the resource provider (1 to 8). It determines the tasks which the agent can perform (see Table 2).
- Its capacity: each resource provider can perform only one task at a time.
- A probability distribution function specifying the time it takes to perform each task. If the resource provider can perform n tasks, the resource provider will have n different time distributions.
- The cost the agent incurs for performing each task
- Its bidding strategy (which can be adaptative bidding, honest bidding or cheating).

6.2. Scenarios

We test our approach in three different scenarios:




Figure 3: Multi-Agent System architecture for the simulation

6.2.1. Scenario 1: Efficiency

The goal of the first scenario is to analyze the behaviour of the auction mechanism regarding the attributes in an independent way (economic cost for unverifiable attributes and delays for verifiable ones). For that purpose, efficiency achieved with VMA2 is compared with previous approaches:

- **First price uni-attribute auction (bid amount):** The resource provider with the cheapest price is the bidder which will realize the service. Its payment will be the economic amount indicated in the bid.
- **First price uni-attribute auction (delivery time):** The resource provider with the fastest delivery time is the bidder which will carry out the service. Its payment will be the economic amount indicated in the bid.
- **Vickrey multi-attribute auction:** The winner determination problem is computed as in VMA2 but payment is calculated as in Equation 14. Thus, payment is performed independently of the outcomes of the task. Product is used as evaluation function ($f_0(b, t) = b * t$)⁵.
- **Verifiable and Unverifiable Multi-attribute Auction VMA2:** The winner and his payment is determined using the mechanism presented in Section 5. The main difference with the Vickrey multi-attribute auction is that payment is now computed with the received verifiable attributes as explained in Equation 17. Product is used as evaluation function ($f_0(b, t) = b * t$).

⁵Previous studies have shown that product offers more equilibrated allocations than other evaluation functions (Pla et al., 2012b)

In this scenario all the bidders follow an adaptative bidding strategy, adapting the bids they offer to increase the chances of winning the auction and maximizing their benefits (Lee and Szymanski, 2005).

6.2.2. Scenario2: Strategy Proofness

The main objective of this scenario is to analyze the incentive compatibility of the mechanism. Moreover, we point how VMA2 preserves auctioneers utility when bidders try to obtain higher benefits by manipulating the attributes of their bids. For that purpose we compare a simulation where all the agents bid following an honest strategy (they bid the attributes they think they will be able to deliver) with a simulation where four resource providers (RP5 to RP8) try to manipulate the bid attributes and the rest bid honestly.

6.2.3. Scenario 3: Handling auctioneer provided attributes

The goal of the third scenario is to show how social welfare goals can be implemented by means of auctioneer provided attributes. For that purpose, the experiment realized in Scenario 1 is repeated using VMA2 and an auctioneer provided priority attribute w . Moreover, the mechanism is compared with a multi-attribute adaptation of the PA mechanism (Murillo et al., 2008) where w is only used to determine the winner but not to compute the payment.

In both cases (VMA2 and PA), the priority is calculated using the number of auctions bidders have participated since the last time they won (Equation 30) where a lower w means that they have more chances to win the next auction, as follows:

$$w_i = \frac{1}{1 + lw_i} \quad (30)$$

where lw_i corresponds to the number of auctions bidder a_i has participated since the last time it won and w_i its priority.

6.3. Results and Discussion

In this subsection the results of the proposed scenarios are shown and discussed. In the first scenario, as the objective of the experiment is to evaluate the attributes independently, the results are evaluated in terms of economic cost and delays (when a service is delivered after its maximum execution time, see Table 3). In the second scenario, in addition to economic cost and delays, utility is used to evaluate the resulting allocations when considering all the attributes. Finally, in the last scenario, the results are also in terms of social welfare using the Gini's index (Gastwirth, 1972).




Figure 4: Total task allocation cost for each Service Agent using different allocation mechanisms: Vickrey multi-attribute auction (blue), VMA2 (red), choosing the cheapest resources provider (green) and choosing the fastest resource provider (purple).

6.3.1. Scenario 1

This scenario analyzes the performance of VMA2 against other three auction mechanisms in terms of economic cost and delays. Figure 4 shows the economic amount the 7 different service agents had to spent in order to allocate their tasks to a resource provider. Yellow and red bars show the economic cost when assigning the tasks to resource using uni-attribute auctions, following a policy of assigning the cheapest or the fastest resource whilst blue bars present the economic expense when using an auction mechanism: clear blue for the Vickrey multi-attribute mechanism and dark blue for the mechanism presented in this paper, VMA2. The chart clearly shows that following the strategy of picking the fastest resource provider leads to a higher economic cost for the service execution. This is a logical result as this strategy ignores the economic cost of the resources. Regarding the other three strategies it is not clear which one is the most efficient in economic terms as the most inexpensive strategy varies according to the kind of service that has been allocated. E.g. we can see how for SA7 the Vickrey auction results in a lower economic investment, however, for SA2 the Vickrey auction results in the worst strategy in economic terms (excluding the fastest resource strategy). If we focus on VMA2 we can see that this strategy is the one which gives best economic results to four service agents (SA1, SA2, SA5 and SA6), the second best choice for SA4 and SA7 and the worst for SA3.

To complete the information given in the previous paragraph, Figure 5 shows a box plot of the mean cost of a service in the different simulations, offering a global view of a service cost. The chart shows that VMA2 reduces




Figure 5: Mean cost for a service allocation when using different allocation methods: Vickrey multi-attribute auction (1), VMA2 (2), choosing the cheapest resource provider (3) and choosing the fastest provider (4).




Figure 6: Delays produced in the simulation when using different allocation methods: Vickrey multi-attribute auction (1), VMA2 (2), choosing the cheapest resource provider (3) and choosing the fastest provider (4).

the cost for the service allocation, this fact is probably caused by the payment reduction the service agents apply when the checkable attributes are not respected (when the item is not delivered as it was specified).

Regarding the delays produced in the system (Figure 6), we can see how the allocation strategy which has obtained a lower number of delays is the fastest resource choice while the worst results have been achieved by the cheapest resource strategy. This two facts are not surprising as the first policy is focused exclusively in time efficiency while the second one ignores this parameter. Regarding VMA2 and the multi-attribute Vickrey auction, it can be seen that they obtain a very similar number of delays. This is a logical fact as both use the same evaluation function to determine the winner of the auction, favoring the same kind of resource provider. They are clearly less sensitive to delays than the best price choice and the number of delays produced is not much higher than the delays produced using the fastest resource choice.

This experiment shows how VMA2 can provide a good compromise among the different attributes which take part in the resource allocation process. It can achieve as competitive prices as electing the cheaper resource provider without jeopardizing the quality of the rest of the attributes. Comparing our approach with a classical Vickrey multi-attribute auction we can see how, if using the same evaluation functions, the results regarding the quality of the auction are analogue, however, when comparing the economic efficiency against a classic Vickrey multi-attribute auction, VMA2 reduces costs for the auctioneer as it pays less to bidders which breaches the submitted bid. In other words, VMA2 preserves the auctioneer utility when bidders break the auctioneer agreement. It is important to remark that in an ideal world




Figure 7: Economic amount spent by agents to allocate tasks to resources using different task allocation methods: Vickrey multi-attribute auction (blue), VMA2 (red), choosing the cheapest resource provider (green) and choosing the fastest provider (purple).




Figure 8: Auctioneers utilities obtained by the resulting task allocation when using different task allocation methods: Vickrey multi-attribute auction (blue), VMA2 (red), choosing the cheapest resource provider (green) and choosing the fastest provider (purple).

where all the bidders respect the submitted bids, VMA2 and Vickrey Multi-attribute auctions would act exactly in the same way and would obtain the same results (both in economic and in attribute efficiency).

6.3.2. Scenario2

Now, we analyze the behaviour of VMA2 in front of cheating agents. Figures 7 8 and 9 show the results of the second experiment. The first one shows the economic cost of each service using the Vickrey multi-attribute auction and the VMA2 when all the bidders follow an honest strategy (blue and red lines) and when four of them try to manipulate the attributes (green and purple). Figure 8 shows the auctioneers utilities in the same contexts (following the same colour pattern). The first remarkable thing of these charts is that when all the bidders act following an honest strategy both the Vickrey auction and VMA2 obtain similar results, the service cost is quite similar and the same for the auctioneers utilities. As mentioned above, this reflects the fact that when all the bidders follow an honest strategy the Vickrey auction and VMA2 are analogue and we can blame the small difference between the obtained results to the randomness of the simulation. If focusing on the simulations where bidders 5 to 8 have tried to manipulate the bid, we can see how when using the Vickrey auction the auctioneers have paid a small overprice as they evaluated the offered bids as very good items. However, when the received attributes where not the ones in the




Figure 9: Resource providers (bidders) revenue when using different task allocation methods: Vickrey multi-attribute auction (blue), VMA2 (red), choosing the cheapest resource provider (green) and choosing the fastest provider (purple).

auction agreement, lowering their utility. The opposite happens when using VMA2, Figure 7 shows that auctioneers have paid a lower price as when they detected that the bidden attributes where not true. This has made its utility to be maintained despite having received worse attributes than expected (Figure 8).

Regarding the bidders point of view, Figure 9 clearly points that, when using VMA2, bidders 5 to 8 would have obtained higher incomes by revealing their real attributes (red bars) than trying to manipulate the auction (purple bars). This does not happens when using Vickrey auction as allocation method since bidders obtained higher benefits when they bidden false attributes (green bars) than when revealing the real ones (blue bars).

6.3.3. Scenario3

Finally, this subsection analyzes the performance of VMA2 when including an auctioneer provided priority attribute in the auction mechanism in order to improve fairness in the auction process. Moreover, VMA2 is compared with PA. The results are shown in terms of agents revenue (Figures 10 and 11) while fairness is evaluated using Gini’s index.

The bar chart in Figure 10 shows that when using a priority attribute, almost all the service agents have not needed to spend as much money to allocate their services than in the previous scenarios (without priority attributes). This is clearly seen in the expenditure done by SA2. In cases where the costs have not been reduced (SA 1 and 4) the cost for the service allocations has been very similar to the one resulting from not using a priority attribute. Thus we can say that the inclusion of fairness into the mechanism can reduce the costs for the auctioneers. When comparing PA and VMA2




Figure 10: Service Agents (auctioneers) and Resource Providers (bidders) revenue when using VMA2 without a priority attribute (blue), when using a multi-attribute adaptation of the PA algorithm (red) and when using VMA2 with a priority attribute (green).

from the auctioneers point of view, we can see how the service costs have been similar. Regarding the benefits of the service providers, it can be seen that the use of a priority attribute considerably changes the distribution of the revenues. The differences between the richest and the poorest service providers are significantly reduced when including fairness into the mechanism (e.g. RP2 and RP3 against RP6). Another remarkable fact is that bidders which obtained poor revenues (e.g. RP6 and RP7) increase their benefits. However, it is also important to notice that the inclusion of fairness into the mechanism has modified the ranking of the wealthiest service providers (e.g. RP3 goes from being the richest agent to the third richest one).

To evaluate the fairness of the allocations revenue, Figure 11 presents a box plot of the gini's index (Gastwirth, 1972) corresponding to the 200 executions of the scenario. It clearly shows that VMA2 provides the lower inequality coefficient (0.102), this fact can also be observed on the bar chart where the green bars corresponding to the bidders present a lower gap between the richest and the poorest agents. The inequality coefficient obtained by PA is also lower than the coefficient obtained when no priority is used (0.224 against 0.469). Thus, by observing both plots we can conclude that the use of priorities as an auctioneer provided attribute can improve the fairness of the task allocation by assigning some services to the weaker agents without increasing the global cost of the allocation, however by using this methodology, some of the strongest bidders may experience envy as their incomes could be reduced and their position in the wealth rank altered for the




Figure 11: Gini’s index for the resulting allocations when using VMA2 with different priority attributes:using VMA2 without a priority attribute (1), when using a multi-attribute adaptation of the PA algorithm (2) and when using VMA2 with a priority attribute (3).

sake of equity. Further research should involve the use of different priority calculation methods in order to overcome this problem (e.g, by using less aggressive functions or stochastic priorities).

7. Conclusions

In this paper the need of reviewing multi-attribute auction mechanisms named VMA2 based on the different kinds of attributes is posed. Particularly, the attributes are divided between verifiable bidder provided attributes, which define the physical characteristics of the auctioned item, unverifiable bidder attributes, which are related with the bidders internal valuation of the auctioned item, and auctioneer provided attributes, which are used to express evaluations and opinions of the auctioneers regarding the bidders.

From the attribute classification, a new reverse multi-attribute mechanism based on the role of the attributes is presented. On the one hand, the distinction between different attribute roles ensures incentive compatibility during the bidding process as offering untruthful unverifiable attributes may reduce the chances of winning the auction (underbidding) or in a revenue-lost (overbidding). On the other hand, verifiable attributes can be used to compare the delivered item with the one which was offered during the auction process, allowing auctioneers to readjust their payments when bidders do not respect the auction agreement. Furthermore, the use of auctioneer provided attributes (e.g. priorities) allows to include different social welfare configurations as well as to add other concepts such as trust or robustness to the mechanism. Particularly, in this paper we have illustrated how an egalitarian

allocation can be achieved.

The presented mechanism has been tested in a simulator which reproduces the operation of an industry organization in which workflows are trying to externalize tasks using auctions (price and time are the attributes involved during the process). In the simulations, VMA2, has been compared with other auction mechanisms such as uni-attribute auctions and a Vickrey adaptation for multi-attribute auctions. The experiments have shown that VMA2 can achieve an optimal revenue for the auctioneers without jeopardizing the quality of the attributes (represented as delays during the simulation), and that the presented mechanism encourages truthful bidding as bidders obtain higher utilities when they choose to reveal their real preferences. The tests also show how, with VMA2, the utility of the auctioneers is preserved when bidders do not deliver their items in the agreed conditions thanks to the use of verifiable attributes. These results are compared to a Vickrey multi-attribute auction where under these circumstances, the auctioneers utility dramatically decreases. Regarding the analysis of social welfare through the use of auctioneer-provided attributes (priorities), the experiments show the difference with previous mechanism that follow ad hoc approaches which include priorities on the winner determination problem but not in the payment method. Thus, handling this kind of information (priority) as auctioneer provided attributes, as we are proposing, offers the possibility of dealing with the attributes in a holistic approach along all of the mechanism steps.

The work presented on this paper has been developed but it is not limited to a supply chain management domain. An interesting work would be to export and to test VMA2 to new and sizzling domains as the smart electric grid where we are conducting our future research.

References

- Athey, S., Ellison, G., 2011. Position auctions with consumer search. forthcoming Quarterly Journal of Economics 126 (3), 1213–1270.
- Bartal, Y., Gonen, R., Nisan, N., 2002. Incentive compatible multi unit combinatorial auctions. In: In TARK 03. pp. 72–87.
- Bellosta, M.-J., Kornman, S., Vanderpooten, D., 2008. A unified framework for multiple criteria auction mechanisms. Web Intelligence and Agent Systems 6 (4), 401–419.
- Bellosta, M.-J., Kornman, S., Vanderpooten, D., 2011. Preference-based english reverse auctions. Artificial Intelligence 175 (78), 1449 – 1467, <http://www.sciencedirect.com/science/article/pii/S0004370210002043>.
- Bichler, M., Kaukal, M., Segev, A., March 1999. Multi-attribute auctions for electronic procurement. In: Proceedings of the First IBM IAC Workshop on Internet Based Negotiation Technologies. Yorktown Heights, NY.
- Che, Y.-K., 1993. Design competition through multidimensional auctions. The RAND Journal of Economics 24 (4), 668–680.
URL <http://doi.wiley.com/10.2307/2555752>
- Chevaleyre, Y., Dunne, P., Endriss, U., Lang, J., Lematre, M., Maudet, N., Padget, J., Phelps, S., Rodriguez-aguilar, J., Sousa, P., 2006. Issues in multiagent resource allocation. Informatica 30, 2006.
- David, E., Azoulay-Schwartz, R., Kraus, S., 2002. Protocols and strategies for automated multi-attribute auctions. In: Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 1. AAMAS '02. ACM, New York, NY, USA, pp. 77–85.
URL <http://doi.acm.org/10.1145/544741.544761>
- Endriss, U., Maudet, N., Sadri, F., Toni, F., May 2003. Resource allocation in egalitarian agent societies. In: Herzig, A., Chaib-draa, B., Mathieu, P. (Eds.), Secondes Journées Francophones sur les Modèles Formels d’Interaction (MFI-2003). Cépaduès-Éditions, pp. 101–110.
- Fu, H., Jordan, P. R., Mahdian, M., Nadav, U., Talgam-Cohen, I., Vassilvitskii, S., 2012. Ad auctions with data. In: INFOCOM Workshops. pp. 184–189.
- Gastwirth, J. L., August 1972. The estimation of the lorenz curve and gini index. The Review of Economics and Statistics 54 (3), 306–16.
URL <http://ideas.repec.org/a/tpr/restat/v54y1972i3p306-16.html>

- Harrenstein, B. P., de Weerd, M. M., Conitzer, V., 2009. A qualitative vickrey auction. In: Proceedings of the 10th ACM conference on Electronic commerce. EC '09. ACM, New York, NY, USA, pp. 197–206.
URL <http://doi.acm.org/10.1145/1566374.1566403>
- Hoesch-Klohe, K., Ghose, A., L, L.-S., 2010. Towards green business process management. In: IEEE SCC. IEEE Computer Society, pp. 386–393.
URL <http://dblp.uni-trier.de/db/conf/IEEEsc/scc2010.html>
- Krishna, V., Mar. 2002. Auction Theory. Academic Press.
- Lee, J., Szymanski, B., 2005. A novel auction mechanism for selling time-sensitive e-services. In: IEEE Conference on ECommerce Technology (CEC'05). Press, pp. 75–82.
- Lehman, D., Muller, R., Sandholm, T., 2006. The Winner Determination Problem (Chap. 12). MIT Press.
- Lematre, M., Verfaillie, G., Fargier, H., Lang, J., Bataille, N., m. Lachiver, J., 2003. Equitable allocation of earth observing satellites resources. In: In Proc of 5th ONERA-DLR Aerospace Symposium (ODAS03).
- MacKie-Mason, J. K., Varian, H. R., 1994. Generalized vickrey auctions.
- Mahr, T., de Weerd, M. M., 2006. Multi-attribute vickrey auctions when utility functions are unknown. In: Schobbens, P.-Y., Vanhoof, W., Schwanen, G. (Eds.), Proceedings of the Belgium-Dutch Conference on Artificial Intelligence (BNAIC). BNVKI, pp. 221–227.
URL <http://www.st.ewi.tudelft.nl/mathijs/publications/mahr06.pdf>
- Murillo, J., Lpez, B., Muoz, V., Busquets, D., 2012. Fairness in recurrent auctions with competing markets and supply fluctuations. Computational Intelligence 28 (1), 24–50.
URL <http://dblp.uni-trier.de/db/journals/ci/ci28.html>
- Murillo, J., Munoz, V., Lopez, B., Busquets, D., 2008. A fair mechanism for recurrent multi-unit auctions. In: Bergmann, R., Lindemann, G., Kirn, S., Pchouek, M. (Eds.), Multiagent System Technologies. Vol. 5244 of Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 147–158.
- Myerson, R. B., 1979. Incentive Compatibility and the Bargaining Problem. Econometrica 47 (1), 61–74.
URL <http://dx.doi.org/10.2307/1912346>
- Neumann, J. V., Morgenstern, O., 1944. Theory of Games and Economic Behavior. Princeton University Press.
URL <http://jmvidal.cse.sc.edu/library/neumann44a.pdf>

- Parkes, D. C., Kalagnanam, J., 2005. Iterative multiattribute vickrey auctions. *Management Science* 51, 435451.
- Parsons, S., Rodriguez-Aguilar, J. A., Klein, M., Feb. 2011. Auctions and bidding: A guide for computer scientists. *ACM Comput. Surv.* 43 (2), 10:1–10:59.
URL <http://doi.acm.org/10.1145/1883612.1883617>
- Pla, A., Gay, P., Melndez, J., Lpez, B., 2012a. Petri net-based process monitoring: a workflow management system for process modelling and monitoring. *Journal of Intelligent Manufacturing*, 1–16.
URL <http://dx.doi.org/10.1007/s10845-012-0704-z>
- Pla, A., Lopez, B., Murillo, J., 2012b. Multi criteria operators for multi-attribute auctions. In: Torra, V., Narukawa, Y., Lpez, B., Villaret, M. (Eds.), *Modeling Decisions for Artificial Intelligence. Lecture Notes in Computer Science*. Springer Berlin Heidelberg, pp. 318–328.
- Pla, A., López, B., Murillo, j., 2013. Proving incentive compatibility in multi-attribute auctions. In: *13th International Conference of the Catalan Association for Artificial Intelligence*. p. submitted.
- Pla, A., Lpez, B., Murillo, J., 2011. Workflow resource allocation through auctions. In: *Proceedings of the 2nd Workshop on Artificial Intelligence and Logistics (AILog in IJCAI)*. pp. 55–60.
- Pla, A., Lpez, B., Murillo, J., 2012c. Multi-attribute auction mechanism for supporting resource allocation in business process enactment. In: Kersting, K., Toussaint, M. (Eds.), *STAIRS*. Vol. 241 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, pp. 228–239.
URL <http://dblp.uni-trier.de/db/conf/stairs/stairs2012.html>
- Prodan, R., Wiecek, M., Fard, H. M., 2011. Double auction-based scheduling of scientific applications in distributed grid and cloud environments. *J. Grid Comput.* 9 (4), 531–548.
- Puustjärvi, J., 2005. Coordinating electronic auctions by workflow systems. In: *Proceedings of the IADIS International Conference WWW/Internet*. pp. 139–142.
- Rackson, R. I., Krane, J. A., Trevisani, P. J., 2002. U.S. Patent No. 6,415,270. Washington, DC: U.S. Patent and Trademark Office.
- Ramchurn, S. D., Mezzetti, C., Giovannucci, A., Rodriguez-Aguilar, J. A., Dash, R. K., Jennings, N. R., June 2009. Trust-based mechanisms for robust and efficient task allocation in the presence of execution uncertainty. *J. Artif. Int. Res.* 35, 119–159.
URL <http://dl.acm.org/citation.cfm?id=1641503.1641506>

- Sadeh, N., Sun, J., 2003. Multi-attribute supply chain negotiation: coordinating reverse auctions subject to finite capacity considerations. In: Proceedings of the 5th international conference on Electronic commerce. ICEC '03. ACM, New York, NY, USA, pp. 53–60.
URL <http://doi.acm.org/10.1145/948005.948013>
- Shih, D.H., Lin, B., Huang, S.Y., 2007. MoRVAM: A reverse Vickrey auction system for mobile commerce. *Expert Systems with Applications*. V32-4, 1113-1123.
URL <http://dx.doi.org/10.1016/j.eswa.2006.02.012>.
- Strecker, S., 2004. Multiattribute auctions in electronic procurement: Theory and experiment. Dissertation, Universitt Karlsruhe (TH).
- Suyama, T., Yokoo, M., 2004. Strategy/false-name proof protocols for combinatorial multi-attribute procurement auction. In: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 1. AAMAS '04. IEEE Computer Society, Washington, DC, USA, pp. 160–167.
URL <http://dx.doi.org/10.1109/AAMAS.2004.255>
- Talluri, S., Ragatz, G. L., 2004. Multi-attribute reverse auctions in b2b exchanges: A framework for design and implementation. *Journal of Supply Chain Management* 40 (1), 52–60.
- Tavana, M., Abtahi, A.M., Khalili-Damghani, K., 2013. A new multi-objective multi-mode model for solving preemptive timecostquality trade-off project scheduling problems, *Expert Systems with Applications*, in press.
URL <http://dx.doi.org/10.1016/j.eswa.2013.08.081>
- Torrent, F., Pla, A., 2013. Production factory data set. University of Girona.
URL <http://eia.udg.es/~apla/fac.data/>
- Varian, H. R., 2007. Position auctions. *International Journal of Industrial Organization* 25 (6), 1163 – 1178.
URL <http://www.sciencedirect.com/science/article/pii/S0167718706001317>
- Vickrey, W., 1961. Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance*, 8–37.
- Zhang, Y., Huang, G. Q., Qu, T., Ho, O., 2010. Agent-based workflow management for rfid-enabled real-time reconfigurable manufacturing. *International Journal of Computer Integrated Manufacturing* 23 (2), 101–112.
- Zhao, D., Zhang, D., Perrussel, L., 2011. Mechanism design for double auctions with temporal constraints. In: *IJCAI*. pp. 472–477.