

On the Empirical Evaluation of Mixed Multi-Unit Combinatorial Auctions

Meritxell Vinyals and Jesús Cerquides

IIIA, Artificial Intelligence Research Institute
CSIC, Spanish National Research Council
{meritxell,cerquide}@iiia.csic.es

Abstract. Mixed Multi-Unit Combinatorial Auctions extend and generalise all the preceding types of combinatorial auctions. In this paper, we try to make headway on the practical application of MMUCAs by: (1) providing an algorithm to generate artificial data that is representative of the sort of scenarios a winner determination algorithm is likely to encounter; and (2) subsequently assessing the performance of an Integer Programming implementation of MMUCA on CPLEX.

1 Introduction

A combinatorial auction (CA) is an auction where bidders can buy (or sell) entire bundles of goods in a single transaction ([1]). Although computationally very complex, selling items in bundles has the great advantage of eliminating the risk for a bidder of not being able to obtain complementary items at a reasonable price in a follow-up auction (think of a combinatorial auction for a pair of shoes, as opposed to two consecutive single-item auctions for each of the individual shoes). The study of the mathematical, game-theoretical and algorithmic properties of combinatorial auctions has recently become a popular research topic in AI. This is due not only to their relevance to important application areas such as electronic commerce or supply chain management, but also to the range of deep research questions raised by this auction model. In particular, supply chain formation (SCF) appears as a very promising application area where strong complementarities arise as discussed in [6].

The work in [3] introduces a generalisation of the standard model of CA. This new auction extends and generalises all the preceding types of combinatorial auctions: single-unit CAs, multi-unit CA, double CAs, and supply chain formation CAs. It provides a bidding language that can express several types of complex bids, and allows for bids on combinations of production processes, as well as a general winner determination problem (WDP) solver working on any network topology. This auction model is called mixed multi-unit combinatorial auction (MMUCA). Consider as an example the assembly of a car's engine, whose structure is depicted in Fig.

(a) in table 1. Notice that each part in the diagram, in turn, is produced from further components or raw materials. For instance, a cylinder ring (part 8) is produced by transforming some amount of stainless steel with the aid of an appropriate machine. Therefore, there are several production levels involved in the making of a car’s engine. A MMUCA allows to run a supply chain formation auction where bidders can bid over bundles of parts, bundles of transformations, or any combination of parts and transformations. Despite its potential for application, and unlike CAs, little is

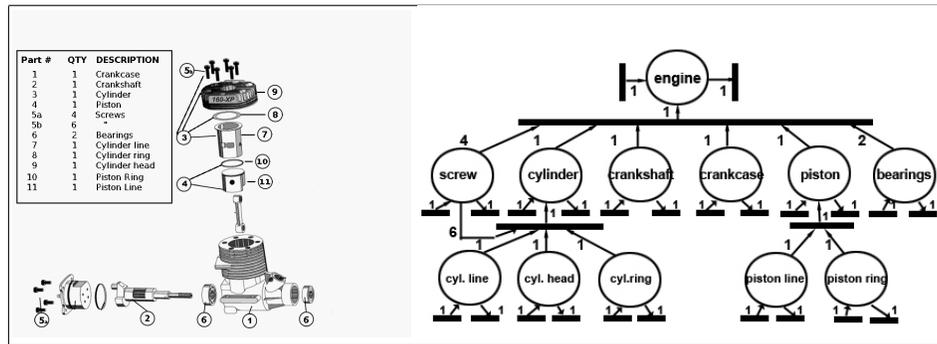


Table 1. (a) Components of a car engine. (b) Market transformations for a car’s engine.

known about the practical application of MMUCAs since no empirical results have been reported on any winner determination algorithms. These results are unlikely to come up unless, and along the lines of the research effort carried out in CAs [4], researchers are provided with algorithms or test suites to generate artificial data that is representative of the auction scenarios a winner determination algorithm is likely to encounter. Hence, winner determination algorithms could be accurately tested, compared, and improved. In this paper, we try to contribute to the practical application of MMUCAs along two directions. Firstly, we provide an algorithm to generate artificial data sets that are representative of the sort of scenarios a winner determination (WD) algorithm is likely to encounter. Secondly, we employ such algorithm to generate artificial data and subsequently assess the performance of an Integer Programming (IP) implementation of MMUCA on CPLEX. Based on our empirical results, we argue that there is a need for special-purpose WD algorithms for MMUCAs if these are intended to be employed in large scenarios.

The paper is structured as follows. In section 2 we provide some background on MMUCAs. Next, in section 3 we analyse the required features of an artificial data set generator for MMUCAs whose algorithm is detailed in section 4. In sections 5 and 6, we analyse some early, empirical

results of an IP formulation of the WDP, draw some conclusions and outline paths to future research.

2 Background

Next, we introduce MMUCA by summarising the work in [3, 2]. Let G be the finite set of all the types of goods. A *transformation* is a pair of multisets over G : $(\mathcal{I}, \mathcal{O}) \in \mathbb{N}^G \times \mathbb{N}^G$. An agent offering the transformation $(\mathcal{I}, \mathcal{O})$ declares that it can deliver \mathcal{O} *after* having received \mathcal{I} . In our setting, bidders can offer any number of such transformations, including several copies of the same transformation. That is, agents will be negotiating over *multisets of transformations* $\mathcal{D} \in \mathbb{N}^{(\mathbb{N}^G \times \mathbb{N}^G)}$. For example, $\{(\{\}, \{a\}), (\{b\}, \{c\})\}$ means that the agent in question is able to deliver a (no input required) and that it is able to deliver c if provided with b .

In an MMUCA, agents negotiate over bundles of transformations. Hence, a *valuation* $v : \mathbb{N}^{(\mathbb{N}^G \times \mathbb{N}^G)} \rightarrow \mathbb{R}$ is a (typically partial) mapping from multisets of transformations to the real numbers. Intuitively, $v(\mathcal{D}) = p$ means that the agent equipped with valuation v is willing to make a payment of p in return for being allocated all the transformations in \mathcal{D} (in case p is a negative number, this means that the agent will accept the deal if it *receives* an amount of $|p|$). For instance, valuation $v(\{(\{line, ring, head, 6 \cdot screws, screwdriver\}, \{cylinder, screwdriver\})\}) = -10$ means that some agent can assemble a cylinder for \$10 when provided with a cylinder line, a cylinder ring, a cylinder head, six screws, and a screwdriver, and returns the screwdriver once done¹.

An *atomic bid* $BID(\{(\mathcal{I}^1, \mathcal{O}^1), \dots, (\mathcal{I}^n, \mathcal{O}^n)\}, p)$ specifies a finite multiset of finite transformations and a price. Informally, an OR-combination of several bids signifies that the bidder would be happy to accept any number of the sub-bids specified, if paid the sum of the associated prices. An XOR-combination of bids expresses that the bidder is prepared to accept at most one of them.

The *input* to the WDP consists of a complex bid expression for each bidder, a multiset \mathcal{U}_{in} of goods the auctioneer holds to begin with, and a multiset \mathcal{U}_{out} of goods the auctioneer expects to end up with. A *valid solution* to the WDP will be a *sequence of transformations* satisfying that: (1) the multiset of transformations in the sequence has to *respect the bids* submitted by the bidders; and (2) the set of goods held by the auctioneer in the end is a superset of \mathcal{U}_{out} .

For the formal definition of the WDP, we restrict ourselves to bids in the XOR-language, which is known to be fully expressive (as proved

¹ We use $6 \cdot screws$ as a shorthand to represent six identical elements in the multiset.

by Cerquides et al. [3]). Therefore, solving the WDP for MMUCAs with XOR-bids amounts to maximise $\sum_{b \in B} x_b \cdot p_b$, while fulfilling the constraints informally defined above, where B stands for the set of all atomic bids, x_b is binary decision variable indicating whether bid b is selected or not, and p_b stands for the price of bid b . As noticed in [3], the number of decision variables of an IP to solve a MMUCA WDP is of the order of $|T|^2$, where $|T|$ is the overall number of transformations mentioned anywhere in the bids. As the reader may notice, this represents a serious computational cost as the number of transformations grow.

3 Bid Generator Requirements

In order to test and compare MMUCA WD algorithms, researchers must be provided with algorithms or test suites to generate artificial data that is representative of the auction scenarios a WD algorithm is likely to encounter. Hence, WD algorithms can be accurately tested, compared, and improved. Unfortunately, we cannot benefit from any previous results in the literature since they do not take into account the notion of transformation introduced in [3, 2]. In this section we make explicit the requirements for a bid generation technique considering that in MMUCA agents trade transformations instead of goods.

3.1 A Taxonomy of Transformations

Bids in MMUCAs are composed of transformations. Each transformation expresses either an offer to buy, to sell, or to transform some good(s) into (an)other good(s). Since transformations are the building blocks composing bids, we must firstly characterise the types of transformations a bid generator may need to construct in order to produce bids. Our analysis of transformations has led to a classification into three types, namely:

1. **Output transformations** are those with no input good(s). Thus, an O-transformation represents a bidder's offer to sell some good(s). Besides, an O-transformation is equivalent to a bid in a reverse CA.
2. **Input transformations** are those with no output good(s). Thus, an I-transformation represents a bidder's offer to buy some good(s). Notice that an I-transformation is equivalent to a bid in a direct CA.
3. **Input-Output transformations** are those whose input and output good(s) are not empty. An IO-transformation stands for a bidder's offer to deliver some good(s) after receiving some other good(s): *I can deliver \mathcal{I} after having received \mathcal{O}* . They can model a wide range of different processes in real-world situations (e.g. assembly, transformation, or exchange).

Figure (b) in Table 1 presents samples of each transformation type. In the figure, vertical, black bars stand for transformations, circles stand for goods, and directed arrows from goods into or from transformations represent the goods input into or produced out of a transformation. Thus, for instance, we differentiate an I-transformation to consume a piston, an O-transformation to give away a piston, and an IO-transformation giving away a piston after receiving a piston ring and a piston line.

Notice that any bid in a MMUCA results as a combination of transformations of the above-listed types. Therefore, a bid generator for MMUCA must support the generation of transformations of all these types.

3.2 Requirements

It is time to consider how to combine transformations of the above-described types in order to construct bids. Since MMUCAs generalise CAs, as discussed in [3], our approach is to depart from artificial data sets generators for CAs, keeping the requirements summarised in [4], namely: (1) certain goods are more likely to appear together than others; (2) the number of goods in a bundle is often related to which goods compose the bundle; (3) valuations are related to which goods appear in the bundle, and, where appropriate, valuations can be configured to be subadditive, additive or superadditive in the number of goods requested; and (4) sets of XOR'ed bids are constructed in a meaningful way, on a per-bidder basis.

Notice though that the requirements above must be reformulated, and eventually extended, in terms of transformations since a bidder in a MMUCA bids over a bundle of transformations, whereas a bidder in a CA bids over a bundle of goods. This difference leads to pose a fundamental issue: how should an artificial data set generator for MMUCA compose bids? Indeed, notice that a CA generator bundles goods from a given set of goods to construct bids. And hence, analogously, what is the set of transformations from which a MMUCA generator constructs bids? In order to provide a proper answer we must take inspiration on realistic scenarios faced by buyers and providers. If so, within a given market we expect several producers to offer the very same or similar services (transformations) at different prices, as well as several consumers to require the very same or similar services (transformations) valued at different prices. In other words, within a given market we can identify a collection of common services that companies request and offer. For instance, in the example in Fig. (a) in table 1, several providers may offer to assemble a cylinder through the very same transformation: $t = (\{6 \cdot screws, 1 \cdot cylinder_line, 1 \cdot cylinder_rig, 1 \cdot cylinder_head\}, \{cylinder\})$.

Eventually, a provider may either offer to perform such transformation several times (e.g. as many times as cylinders are required), or to bundle it with other transformations, or the two. Whatever the case, we can regard this sample transformation as an *atomic* transformation because it represents the minimum transformation required to deliver a service. Hereafter, we shall consider the common goods and services in a given market to be represented as a collection of atomic transformations that we shall refer to as *market transformations*. Therefore, market transformations represent the "goods" providers and buyers can request and bid for. Hence, bids for MMUCAs shall be composed as combinations of market transformations. More formally, we define atomic transformations and market transformations as follows²:

Definition 1 (Atomic transformation). *Given a set of transformations $T = \{t_1, \dots, t_n\}$, we say that transformation $t_i = (\mathcal{I}_i, \mathcal{O}_i)$ is minimum in T iff $\forall t_j \in T$ satisfying that $\forall g_i \in \mathcal{I}_i, g_j \in \mathcal{I}_j$ $g_i, g_j \in \mathcal{I} \cap \mathcal{I}'$ and $\forall g_i \in \mathcal{O}_i, g_j \in \mathcal{O}_j$ $g_i, g_j \in \mathcal{O} \cap \mathcal{O}'$, the following inequalities hold: (i) $m_{\mathcal{I}_i}(g) \leq m_{\mathcal{I}_j}(g) \forall g \in \mathcal{I}_j$ and; (ii) $m_{\mathcal{O}_i}(g) \leq m_{\mathcal{O}_j}(g) \forall g \in \mathcal{O}_j$.*

Definition 2 (Market transformations). *We say that $T \subseteq \mathbb{N}^G \times \mathbb{N}^G$ is a set of market transformations iff: (i) it is finite; (ii) every transformation $t \in T$ is minimum; and (iii) $(\{g\}, \{\}), (\{\}, \{g\}) \in T \forall g \in G$.*

Notice that the third condition ensures that there are at least two market transformations for every good in G , thus ensuring that every good is individually available to buy and/or sell. Fig. (b) in table 1 depicts a sample of market transformations if intending to build the car engine in Fig. (a) in table 1. From the discussion so far, we shall consider a new requirement: "there is a finite set of market transformations to bid for".

If bids are composed as combinations of market transformations, we must introduce the notion of *transformation multiplicity* as the counterpart of good multiplicity (the number of units of a given good within an offer or a request). Say that in a CA a bidder submits a bid for the goods in multi-set $\{\text{engine}, \text{engine}, \text{piston}\}$. It is clear that the multiplicity in this bundle of good *engine* is two, whereas the multiplicity of good *piston* is one. But things become more complicated when we consider transformations because the multiplicity of a given transformation must be defined in terms of another transformation, which in turn is composed of multiple input and output goods. Intuitively, we say that a

² Hereafter, we consider that the multiplicity (that is, the number of occurrences) of the elements in a multiset \mathcal{S} is provided by a function $m_{\mathcal{S}}$ from \mathcal{S} to \mathbb{N} .

transformation is a multiple of another one if both share the same input and output goods and the former has more input and output goods than the latter but keeping the same ratio between input and output goods. For instance, given transformations $t = (\{6 \cdot \text{screws}, 1 \cdot \text{cylinder_line}, 1 \cdot \text{cylinder_rig}, 1 \cdot \text{cylinder_head}\}, \{\text{cylinder}\})$ and $t' = (\{12 \cdot \text{screws}, 2 \cdot \text{cylinder_line}, 2 \cdot \text{cylinder_rig}, 2 \cdot \text{cylinder_head}\}, \{2 \cdot \text{cylinder}\})$ we say that t' has multiplicity two with respect to t . Put formally:

Definition 3 (Transformation multiplicity). *Let $t = (\mathcal{I}, \mathcal{O})$ and $t' = (\mathcal{I}', \mathcal{O}')$ be transformations such that $\forall g \in \mathcal{I}, g' \in \mathcal{I}' \ g, g' \in \mathcal{I} \cap \mathcal{I}'$ and $\forall g \in \mathcal{O}, g' \in \mathcal{O}' \ g, g' \in \mathcal{O} \cap \mathcal{O}'$. We say that t has multiplicity κ with respect to t' iff $m_{\mathcal{I}}(g) = \kappa \cdot m_{\mathcal{I}'}(g) \ \forall g \in \mathcal{I}$ and $m_{\mathcal{O}}(g) = \kappa \cdot m_{\mathcal{O}'}(g) \ \forall g \in \mathcal{O}$.*

A further issue has to do with the way bidders value transformations. Notice that performing a transformation to assemble the engine in Fig. (a) in table 1 results in a new product that has more market value than its parts. Therefore, a car maker values the transformation according to his expected benefits, namely the difference between the expected market value of the engine and the cost of its parts. Therefore, if the parts cost \$850 and the expected market value of the engine is \$1000, the car maker should be willing to offer to pay less than \$150 for the transformation. On the other hand, any provider is expected to request less than \$150 in order to perform the transformation. In general, buyers and providers in a MMUCA should value a transformation on the basis of the difference between the expected market value of its output goods and the cost of its input goods. Notice though that we are not assuming here that such difference must be always positive. This discussion leads us to a further requirement: *"every transformation valuation is assessed in terms of the surplus resulting from the valuation of its output goods with respect to the valuation of its input goods"*.

Finally, a last requirement stems from the fact that, unlike auctioneers in CAs, not all goods involved in a MMUCA must be requested by the auctioneer. Back to our example of a car maker in need of engines as depicted in Fig. (a) in table 1, it can run a MMUCA only requesting engines. Thereafter, bidders may offer already-assembled engines, or other goods (e.g. parts like crankcases, crankshafts, or screws) that jointly with transformations over such goods help produce the requested goods. Hence, the new requirement goes as follows: *"unrequested goods by the auctioneer may become involved in the auction"*.

Following the analysis above, we can reformulate the requirements for an artificial data set generator for CAs and add the new requirements

derived so far to finally obtain the requirements for an artificial data set generator for MMUCAs: (1) there is a finite set of market transformations to bid for; (2) certain transformations are more likely to appear together than others; (3) the number of transformations in a bundle is often related to which transformations compose the bundle; (4) valuations are related to which transformations appear in the bundle, and, where appropriate, valuations can be configured to be subadditive, additive or superadditive in the number of transformations requested; (5) every transformation valuation is assessed in terms of the difference between the valuation of its output goods with respect to the valuation of its input goods; (6) sets of XOR'ed bids are constructed in a meaningful way, on a per-bidder basis; and (7) unrequested goods by the auctioneer may be involved in the auction.

4 An Algorithm for Artificial Data Set Generation

In what follows we describe a bid generation algorithm that automates the generation of artificial data sets for MMUCA that while capturing the requirements above. The algorithm's purpose is to generate MMUCA WDP (each one composed of a collection of XOR bids and the set of goods available to and requested by the auctioneer) that can be subsequently fed into an MMUCA WD algorithm. The bid generation technique makes explicit which transformations and how many of them to offer/request in a bundle, how to price the bundle, and which bids to combine in an XOR bid. Hence, the algorithm starts by generating the set of goods involved in MMUCA. Next, it generates the goods the auctioneer requests. After that, it creates a subset of atomic transformations, which are the market transformations to employ for bid generation. Thereafter, it generates bids as linear combinations of market transformations, which are subsequently priced according to a pricing policy. The resulting bids are further composed into XOR (mutually exclusive) bids because the XOR language is a fully expressive language allowing bidders to express all their preferences in a single XOR bid [3]³. Hence, our algorithm assumes that each bidder formulates a single XOR bid, being the number of bidders equal to the number of XOR bids.

Good Generation This process requires the number of different goods (n_{goods}) involved in an auction along with the maximum price any good can take on ($maxPrice$). Based on these values, it assesses for each good g : (1) its average market price (μ_g) drawn from a uniform distribution

³ Here we only provide the bid generation algorithm. The interested reader must refer to [5] for a description of all algorithms required by the artificial data set generator.

$U[1, maxPrice]$ where $maxPrice$ stands for the maximum market price any good can take on; and (2) the distribution to assess its multiplicity, or more precisely, the success probability ($g_{geometric}$) of a geometric probability distribution from which the good multiplicity can be drawn.

Requested Goods Generation This process assesses the number of units of each good the auctioneer requests, namely the multiset \mathcal{U}_{out} . Since the auctioneer must not request all goods, this process selects a subset of the goods in G to be part of \mathcal{U}_{out} . Firstly, it determines whether a good g is requested by the auctioneer by comparing the value drawn from a uniform distribution $U[0, 1]$ with $p_{good_requested}$, the probability of adding a new good to \mathcal{U}_{out} . Once included a given good g , the number of units requested for g is drawn from a geometric distribution with the success probability $g_{geometric}$ obtained by the good generation process.

Market Transformations Generation This process generates a finite set of transformations to be employed as the building blocks to subsequently compose bids. For each good, this procedure constructs two market transformations (one I-transformation and one O-transformation). Each transformation, which according to definition 2 is atomic, involves a single good with multiplicity one. For instance, $(\{engine\}, \{\})$ and $(\{\}, \{engine\})$ stand respectively for the I-transformation and O-transformation for good *engine*. After that, the algorithm generates a limited number of market IO-transformations ($n_{IO_market_transformations}$). In order to generate each market IO-transformation, this procedure chooses the goods to include in its input and output set employing the probabilities of adding some good to the input and output set respectively ($p_{good_in_input}$ and $p_{good_in_output}$). Once included a good to either the input or output set, its multiplicity is calculated from a geometric distribution parametrised by $g_{geometric}$.

Finally, there is the matter of attaching to each market transformation a probability distribution to draw its multiplicity. We assume that the bid generation process, detailed by algorithm 1, is to employ a geometric distribution for each market transformation to calculate its multiplicity. Hence, the generation of market transformations assesses the success probability to be employed by such geometric distributions, namely the probability of adding an extra unit of a transformation already included in a bundle bid. Thus, each transformation t is assigned a success probability $t_{geometric}$. However, we cannot randomly generate success probabilities because transformations are defined over multisets of goods, and therefore we must keep consistency with respect to the success probabilities assigned to each good by the good generation process. Therefore, we

propose to set the success probability for each transformation as follows. Given a transformation $t = (\mathcal{I}, \mathcal{O})$, for each good involved in the transformation, g , we assess its probability of having multiplicity $|m_{\mathcal{I}}(g) - m_{\mathcal{O}}(g)|$. We set the success probability of t to the minimum of these probabilities as follows: $t_{geometric} = \min_{g \in \mathcal{G}} g_{geometric}^{|m_{\mathcal{I}}(g) - m_{\mathcal{O}}(g)|}$.

Algorithm 1 Bid Generation($MTS, n_{XOR.bids}, \mu, \sigma_{prices}, \mu_{add_new_XOR_clause}, \sigma_{add_new_transformation}, \mu_{add_new_transformation}, \sigma_{add_new_transformation}, \alpha$)

```

1: for  $g = 1$  to  $n_{goods}$  do
2:   for  $b = 1$  to  $n_{XOR.bids}$  do
3:      $p_{prices.bid}[b, g] \leftarrow \mu[g] \cdot N(1, \sigma_{prices})$ 
4:   end for
5: end for
6:  $Bids \leftarrow \{\}$ 
7: for  $b = 1$  to  $n_{XOR.bids}$  do
8:    $XORBid \leftarrow EmptyXORBid()$ 
9:    $nXORClauses \leftarrow \mathcal{N}(\mu_{add\_new\_XOR\_clause}, \sigma_{add\_new\_XOR\_clause})$ 
10:  for  $x = 1$  to  $nXORClauses$  do
11:     $Bid \leftarrow EmptyCombinatorialBid()$ 
12:     $nTransfBid \leftarrow \mathcal{N}(\mu_{add\_new\_transformation}, \sigma_{add\_new\_transformation})$ 
13:    for  $t = 1$  to  $nTransfBid$  do
14:       $MT \leftarrow$  Randomly select transformation from  $MTS$ 
15:       $multiplicity \leftarrow Geometric(MT.t_{geometric})$ 
16:       $T.inputs \leftarrow MT.inputs \cdot multiplicity$ 
17:       $T.outputs \leftarrow MT.outputs \cdot multiplicity$ 
18:       $T.price \leftarrow \sum_{g \in T.outputs} p_{prices.bid}[b, g] - \sum_{g \in T.inputs} p_{prices.bid}[b, g]$ 
19:       $p_{offer} \leftarrow (T.t_{geometric})^{multiplicity}$ 
20:       $discount \leftarrow \alpha \frac{1 - e^{-p_{offer}}}{1 - e}$ 
21:       $Bid.t \leftarrow Bid.t \cup T$ 
22:       $Bid.price \leftarrow Bid.price + T.price \cdot (1 - discount)$ 
23:    end for
24:     $XORBid \leftarrow XORBid \cup \{Bid\}$ 
25:  end for
26:   $Bids \leftarrow Bids \cup \{XORBid\}$ 
27: end for
28: return Bids

```

Bid Generation The bid generation algorithm (algorithm 1) generates bids that are subsequently combined into XOR bids, each one encoding the offer or request of a bidder. Firstly, for each XOR bid ($XORBid$) the algorithm composes each bid (Bid) by combining the market transformations (MTS) returned by the market transformation generation process. The number of market transformations ($nTransfBid$) to compose each bid is obtained from a normal distribution $\mathcal{N}(\mu_{add_new_transformation}, \sigma_{add_new_transformation})$ (line 12). Market transformations are randomly chosen from the set of market transformations (MTS) (line 14) and their multiplicity in the bundle bid is obtained from a geometric distribution

with success probability $t_{geometric}$ (line 15). Next, the algorithm prices the transformation according to its multiplicity (lines 16-20), and bundles it into the bid under construction (line 21). Finally, after creating the bid, the algorithm adds it to the XOR bid (line 24). At this point, notice that the number of bids that compose an XOR bid is obtained from a normal distribution $\mathcal{N}(\mu_{add_new_XOR_clause}, \sigma_{add_new_XOR_clause})$ (line 9).

And yet there remains the matter of setting a valuation for each bid within an XOR bid via some pricing policy while fulfilling the requirements in section 3. At this aim, a pricing policy must provide the means to price a good, a transformation, multiple units of the very same transformation, and a bundle of transformations in a realistic manner. As to pricing goods, in order to vary prices among bidders, our algorithm generates a price for bidder b for good g , represented as $p_{prices_bid}[b, g]$, from a normal distribution $\mathcal{N}(\mu[g], \sigma_{prices})$, where $\mu[g]$ stands for good g 's average price in the market and σ_{prices} for the variance among bidders' prices (lines 2-4). Thereafter, a transformation's price for bidder b is assessed in terms of the difference from his valuation of its output goods with respect to his valuation of its input goods (line 19) as stated by requirement 5 in section 3. It is time to address bid valuations while keeping in mind requirement 4 in section 3. At this aim, we propose to introduce superadditivity by applying multiplicity-based discounts to transformations. Going back to the example in Fig. (a) in table 1, we observe that screws are usually traded in higher quantities than full engines. Thus, not surprisingly the same (percentage) discount may apply to an offer for 100 screws than to an offer for 5 engines. Hence, an offer to produce more than 5 engines, though more unlikely, should reflect higher discounts. In other words, as a general rule the more unlikely for a transformation to be traded at certain units (multiplicity), the higher the discount to apply to its overall price. In this way we try to capture in a realistic manner the way multiplicity-based (volume-based) discounts are applied in the real world. Therefore, given transformation t , we firstly assess the probability p_{offer} of the transformation to be traded with multiplicity m from a geometric distribution with success probability $t_{geometric}$ as follows: $p_{offer} = t_{geometric}^{multiplicity}$ (line 20). Secondly, we compute the discount to apply (*discount*) as follows: $discount = \alpha \frac{1 - e^{1 - p_{offer}}}{1 - e}$. Indeed, in this way we manage to apply higher discounts to more unlikely offers within the range $[0, \alpha]$. Notice too that setting α to zero leads to no discounts, and thus to no superadditivity. Finally, a bid valuation is obtained by adding the prices of its transformations (line 23).

5 Experimental results

In this section we shall illustrate the computational cost of solving the WDP for MMUCA. At this aim, we present our first experimental results for MMUCA by assessing the performance of an IP implementation on CPLEX.

As explained at the end of section 2, the number of decision variables of an IP to solve a MMUCA WDP depends on the overall number of transformations. Thus, the number of transformations must be considered as one dimension when measuring the time complexity of a WD algorithm for MMUCA. However, transformations are subsequently combined in several ways in order to finally compose bids in the XOR bidding language. Thus, transformations can be bundled into bids (adding AND-constraints), which in turn may be put together into XOR bids (adding XOR-constraints) giving as a result problems with different levels of restriction. Here we list the significant changes that involve the introduction of different constraints in a MMUCA problem : (1) the introduction of AND-relationships reduces the resulting number of bids ; (2) the introduction of XOR-relationships reduces the resulting number of XOR-bids ; (3) the introduction of XOR-relationship reduces the length of the solution sequence and, consequently, the number of decision variables of the IP since the maximum length of a solution sequence is the sum of the number of transformations contained in the largest bid of each XOR bid.

Hence, we believe that the size of bids (transformation bundles) as well as the size of XOR bids must be regarded as further dimensions when measuring the time complexity of a WD algorithm. Considering the dimensions mentioned so far, we propose a first experiment to evaluate an implementation of the IP formulated in [3] when solving artificial data sets in scenarios with: (1) XOR-bids composed of a single bid with a single transformation (neither AND nor XOR constraints); (2) XOR-bids composed of a single bid with two transformations (AND constraints) (3) XOR-bids composed of two bids with one single transformation (XOR constraints); and (4) XOR-bids composed of two bids with two transformations (both AND and XOR constraints).

Considering the above-described experimental scenarios, we have run our experiments as follows. Firstly, we have generated 50 WDP instances for each configuration using a MATLAB implementation of the artificial data set generator detailed in section 4 whose source code is publicly available at http://www.iiia.csic.es/~meritxell/material/MMUCA_problem_generator.zip. We have solved each WDP with an IP implementation of MMUCA on CPLEX 10.1 and recorded the resulting solving times. Notice though that we have

set to 3600 seconds the time deadline to solve each WDP. Furthermore, we have only considered feasible WDP instances to calculate solving times since the time required by CPLEX to prove infeasibility is (usually) significantly lower than time required to find an optimal solution. All our tests have been run on a Dell Precision 490 with double processor Dual-Core Xeon 5060 running at 3.2 GHz with 2Gb RAM on Linux 2.6.

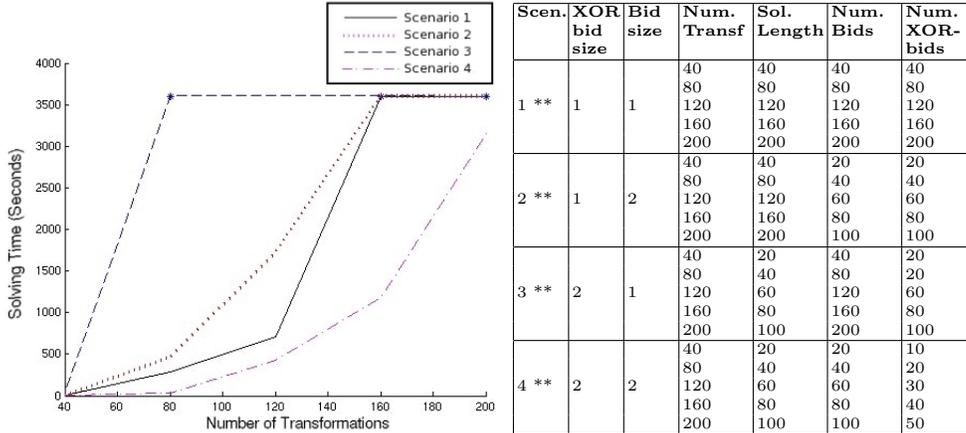


Fig.5 Solving time with respect to number of transformations.

Parameters characterising our experimental scenarios.

** Rest of parameters fixed as: $n_{goods} = 4, n_{IO_market_transformations} = 10, maxPrice = 100, \sigma_{prices} = 0.05, p_{good_requested} = 0.3, p_{good_in_input} = 0.3, p_{good_in_output} = 0.1$

Fig. 5 depicts the median of the solving times resulting when varying the number of transformations for the above-described scenarios. The star(*) symbol stands for the median value exceeding the time limit (3600s). If that is the case, we cannot know the exact median value, but only that it exceeded the time limit. Observe that indeed the MMUCA computational cost increases as the number of transformations grows. Solving times exponentially increase for all scenarios. However, significant differences appear among different scenarios. Observing results of scenario 2 and 3 it stands out that CPLEX solving times increase with the introduction of restrictions with respect to a unconstrained problem (scenario 1). In these scenarios, the space search reduction provided by the addition of the constraints does not compensate CPLEX for considering these restrictions. We also observe that considering XOR-relationships significantly increases the complexity of the problem with respect to considering AND-relationships. However, as we observe in scenario 4, the introduction of constraints can involve important reductions in solving times. We hypothesize that the introduction of a larger number of constraints (in this scenario we introduce AND-constraints and XOR-constraints) leads to a larger reduction of the search space with respect to scenario 2 and 3. In

this case, although CPLEX has to operate with two types of constraints, the reduction of the search space compensates for the introduction of more complexity and produces better results. Nonetheless, it stands out the need for more research in order to accurately understand and explain the significant variations obtained in these experiments through the different scenarios above.

6 Conclusions and future work

In this work, we have attempted at making headway in the practical application of MMUCAs along two directions. Firstly, we have provided an algorithm to generate artificial data sets for MMUCA that are representative of the sort of scenarios a WD algorithm is likely to encounter. A distinguishing feature of the algorithm is that it pursues to capture in a realistic manner how bidders trade transformations. Our algorithm reformulates and extends in terms of transformations the requirements for an artificial data set generator for CAs. Secondly, we provide the first empirical tests for MMUCAs by assessing the performance of a CPLEX IP implementation. These tests assess the computational cost of solving the WDP as transformations grow for different bid expressions in the XOR bidding language. Our results indicate that the scalability of an IP implementation of MMUCA is seriously compromised by the exponential growth of computational cost as the number of transformations increases. Hence, we argue in favour of special-purpose optimal and local algorithms that improve the current performance of an IP implementation.

References

1. P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2006.
2. A. Giovannucci, J. A. Rodríguez-Aguilar, J. Cerquides, and U. Endriss. On the winner determination problem in mixed multi-unit combinatorial auctions. In *Proceedings of the Sixth International Conference on Autonomous Agents and Multiagent Systems*, Honolulu, Hawaii, USA, May 14-18 2007. In press.
3. J. Cerquides, U. Endriss, A. Giovannucci, and J. A. Rodríguez-Aguilar. Bidding languages and winner determination for mixed multi-unit combinatorial auctions. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI 2007)*, pages 1221–1226, India, January 2007.
4. K. Leyton-Brown and Y. Shoham. *A Test Suite for Combinatorial Auctions*, chapter 18. MIT Press, 2006.
5. M. Vinyals, J. Cerquides, and J. A. Rodríguez-Aguilar. On the empirical evaluation of mixed multi-unit combinatorial auctions. Technical Report RR-III A-2007-01, III A-CSIC, February 2007.
6. W. E. Walsh, M. P. Wellman, and F. Ygge. Combinatorial auctions for supply chain formation. In *Proc. of the 2nd ACM Conference on Electronic Commerce*, pages 260–269, Minneapolis, Minnesota, 2000.