Sequential Mixed Auctions

Boris Mikhaylov
Artificial Intelligence Research Institute, IIIA
Spanish National Research Council, CSIC
08193 Bellaterra, Spain
boris@@iiia.csic.es

Jesus Cerquides
Artificial Intelligence Research Institute, IIIA
Spanish National Research Council, CSIC
08193 Bellaterra, Spain
cerquide@@iiia.csic.es

Juan A. Rodriguez-Aguilar
Artificial Intelligence Research Institute, IIIA
Spanish National Research Council, CSIC
08193 Bellaterra, Spain
jar@@iiia.csic.es

ABSTRACT
Mixed multi-unit combinatorial auctions (MMUCAs) offer a high potential to be employed for the automated assembly of supply chains of agents. However, in order for mixed auctions to be effectively applied to supply chain formation, we must ensure computational tractability and reduce bidders' uncertainty. With this aim, we introduce Sequential Mixed Auctions (SMAs), a novel auction model conceived to help bidders collaboratively discover supply chain structures. Thus, an SMA allows bidders progressively build a supply chain structure through successive auction rounds. Moreover, the incremental nature of an SMA provides its participants with valuable information at the end of each auction round to guide their bidding. Finally, we empirically show that SMAs significantly reduce the computational effort required by MMUCA at the expense of a slight decrease in the auctioneer's revenue.

Categories and Subject Descriptors
J.2 [**Computer Applications**]: Social and Behavioral Sciences; I.2.11 [**Computing Methodologies**]: Artificial IntelligenceDistributed Artificial Intelligence

General Terms
Algorithms, Experimentation, Performance

Keywords
Supply chain formation, collaboration, combinatorial auctions, mixed auctions, sequential auctions

1. INTRODUCTION
According to [7], "Supply Chain Formation (SCF) is the process of determining the participants in a supply chain, who will exchange what with whom, and the terms of the exchanges". Combinatorial Auctions (CAs) [2] are a negotiation mechanism well suited to deal with complementarities among the goods at trade. Since production technologies often have to deal with strong complementarities, SCF automation appears as a very promising application area for CAs. However, whilst in CAs the complementarities can be simply represented as relationships among goods, in SCF the complementarities involve not only goods, but also *transformations* (production relationships) along several levels of the supply chain.

The first attempt to deal with the SCF problem by means of Combinatorial Auctions (CA) was done by Walsh et al. in [7]. Later on, mixed multi-unit combinatorial auctions (MMUCAs), a generalization of the standard model of CAs, are introduced in [1]. Rather than negotiating over goods, in MMUCAs the auctioneer and the bidders can negotiate over *transformations*, each one characterized by a set of input goods and a set of output goods. A bidder offering a transformation is willing to produce its output goods after having received its input goods along with the payment specified in the bid. While in standard combinatorial auctions, a solution to the winner determination problem (WDP) is a set of atomic bids to accept, in MMUCAs, the *order* in which the auctioneer "uses" the accepted transformations matters. Thus, a *solution* to the WDP is a *sequence of transformations*. For instance, if bidder *Joe* offers to make dough if provided with butter and eggs, and bidder *Lou* offers to bake a cake if provided with enough dough, the auctioneer can accept both bids whenever he uses Joe's transformation before Lou's to obtain cakes.

Unfortunately, the MMUCA WDP has been proved to be NP-complete [1]. Although reasonably fast solvers have been introduced [3], MMUCA still turns out to be impractical in real-world procurement scenarios. Furthermore, a bidder in MMUCA only knows the desired outcome of the supply chain and the currently stocked goods. Hence, it is difficult, specially for providers placed in the intermediate levels of the supply chain, to decide what to bid for. Therefore, in order for mixed auctions to be effectively applied to supply chain formation, we must ensure computational tractability and reduce bidders' uncertainty. With this aim, we introduce Sequential Mixed Auctions (SMAs), a novel auction model conceived to help bidders collaboratively discover supply chain structures.

SMAs propose to solve a supply chain formation problem by means of a sequence of auctions. The first auctioning round starts with the desired outcome of the supply chain as requested goods and the stocked goods as available goods. During the first auction bidders are only allowed to bid for transformations that either (i) produce goods in the set of requested goods or (ii) consume goods from the available goods. After selecting the best set of transformations, the auctioneer updates the set of requested and available goods after the execution of these transformations and then it will start a new auction. The process will continue until no bids can be found that improve the supply chain.

Each auction involves only a small part of the whole supply chain. Which means that auction's WDP is over a small subsets of bidders, goods and transformations of former MMUCA. These auctions can be solved faster, while we keep the ability to scale our model. The resulting solution is constructed from auctions' WDP solutions. This method copes with complexity of bidding protocol of former MMUCA as well. With our approach manufacturers know which goods are required at any given auction. Consequently, they bid only on those transformations, which provide these goods. That makes all former complexity and uncertainty go away. A more detailed elaboration of sequential MMUCA and its WDP is presented in following sections. Section 2 provides some background of mixed auctions, whereas section 3 formally states the WDP and discusses the means to solve it. Section 4 analyses some initial empirical results comparing SMAs with MMUCAs in terms of solution quality and solving time. Finally, section 5 draws conclusions and discusses future research.

2. BACKGROUND: MIXED AUCTIONS

Next we summarize the work in [1], which introduces mixed multi-unit combinatorial auctions as a generalization of the standard model of combinatorial auctions (CA) and discusses the issues of bidding and winner determination.

Let G be the finite set of all the types of goods. A *transformation* is a pair of multi-sets 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 *multi-sets 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 can deliver a (no input required) and that it can deliver c if provided with b. Note that this is not the same as $\{(\{b\}, \{a, c\})\}$. In the former case, if another agent can produce b if provided with a, we can use the second transformation and get c; in the latter case this would not work.

In a MMUCA, agents negotiate over bundles of transformations. Hence, a *valuation* $v : \mathbb{N}^{(\mathbb{N}^G \times \mathbb{N}^G)} \to \mathbb{R}$ is a (typically partial) mapping from multi-sets of transformations to 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 $v(\{(\{\mathit{line}, \mathit{ring}, \mathit{head}, 6 \cdot \mathit{screws}, \mathit{screwdriver}\}, \{\mathit{cylinder}, \mathit{screwdriver}\})\}) = -10$ 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 done.[1]

An *atomic bid* $b = (\{(\mathcal{I}^1, \mathcal{O}^1), \ldots, (\mathcal{I}^n, \mathcal{O}^n)\}, p)$ specifies a finite multi-set of finite transformations and a price p. A *bidding language* allows a bidder to encode choices between alternative bids and the like [4]. Informally, an OR-combination of several bids means 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 XOR-language is known to be fully expressive for MMUCAs [1].

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). Thus, transformations are the building blocks composing bids. We can classify the types of transformations over which agents bid as follows:

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 $\mathcal{O}$ after having received $\mathcal{I}$.* They can model a wide range of different processes in real-world situations (e.g. assembly, transformation, or exchange).

The *input* to the WDP consists of a complex bid expression for each bidder, a multi-set $\mathcal{U}_{in}$ of (stock) goods the auctioneer holds to begin with, and a multi-set $\mathcal{U}_{out}$ of (required) goods the auctioneer expects to end up with.

In standard CAs, a solution to the WDP is a set of atomic bids to accept. As to MMUCAs, however, the *order* in which the auctioneer "uses" the accepted transformations matters. For instance, if the auctioneer holds a to begin with, then checking whether accepting the two bids $Bid_1 = (\{(\{a\}, \{b\})\}, 10, id_1)$ and $Bid_2 = (\{(\{b\}, \{c\})\}, 20, id_2)$ is feasible involves realizing that we have to use Bid_1 before Bid_2. Thus, a *solution* to the WDP will be a *sequence of transformations*. A *valid* solution has to meet two conditions:

(1) *Bidder constraints:* The multi-set of transformations in the sequence has to *respect the bids* submitted by the bidders. This is a standard requirement. For instance, if a bidder submits an XOR-combination of transformations, at

[1]We use $6 \cdot \mathit{screws}$ as a shorthand to represent six identical elements in the multi-set.

most one of them may be accepted. With no transformation free disposal, if a bidder submits an offer over a bundle of *transformations*, all of them must be employed in the transformation sequence, whereas in the case of transformation free disposal any number of the transformations in the bundle can be included into the solution sequence, but the price to be paid is the total price of the bid.
(2) *Auctioneer constraints:* The sequence of transformations has to be *implementable:* (a) check that $\mathcal{U}_{in}$ is a superset of the input set of the first transformation; (b) then update the set of goods held by the auctioneer after each transformation and check that it is a superset of the input set of the next transformation; (c) finally check that the set of items held by the auctioneer in the end is a superset (the same set in the case of no good free disposal) of $\mathcal{U}_{out}$.

An *optimal* solution is a valid solution that maximizes the sum of prices associated with the atomic bids selected.

The WDP for MMUCAs is a complex computational problem. In fact, one of the fundamental issues limiting the applicability of MMUCAs to real-world scenarios is the computational complexity of the WDP, which is proved in [1] to be $\mathcal{NP}$-complete. Although [3] introduces an integer program to efficiently solve the WDP that drastically outperforms the original IP described in [1], the computational complexity impedes scalability. The next section introduces a new mixed auction model that allows to tame complexity while reducing bidders' uncertainty.

3. SEQUENTIAL MIXED AUCTIONS

3.1 An informal introduction

An SMA proposes to solve a supply chain formation problem by means of a sequence of auctions. The first auction in the sequence starts with the desired outcome of the supply chain as requested goods and the stocked goods as available goods. During the first auction, bidders are only allowed to bid for transformations that either: (i) produce goods in the set of requested goods; or (ii) consume goods from the available goods. After selecting the winning bids (the best set of transformations), the auctioneer updates the set of requested and available goods after the execution of these transformations. Moreover, the winning bids are included as part of the supply chain. Thereafter, the auctioneer starts a new auction in the sequence. The process continues until no bids can improve the supply chain. Hence, the purpose of the auctioneer is to use a sequence of auctions to progressively build the structure of the supply chain.

Figure 1 illustrates the operation of an SMA. Say that a cocktail bar intends to form a supply chain using an SMA to produce a gin & lemon cocktail. Assume that the bar knows approximate market prices for a gin & lemon cocktail as well as for its ingredients. The auctioneer starts the first auction in the SMA issuing a request for quotation (RFQ) for a gin & lemon cocktail. Figure 1a depicts the RFQ along with each good's market price in brackets (e.g. the expected market price of 1 liter of gin is 4€). During the first auction, the auctioneer received two bids: one offering to deliver a cocktail for 9€ (figure 1b); and another one to make a cocktail for 1€ when provided with lemon and gin (figure 1c). The auctioneer must now choose the winning bid out of the bids in figure 1d. However, notice that the bid in figure 1c

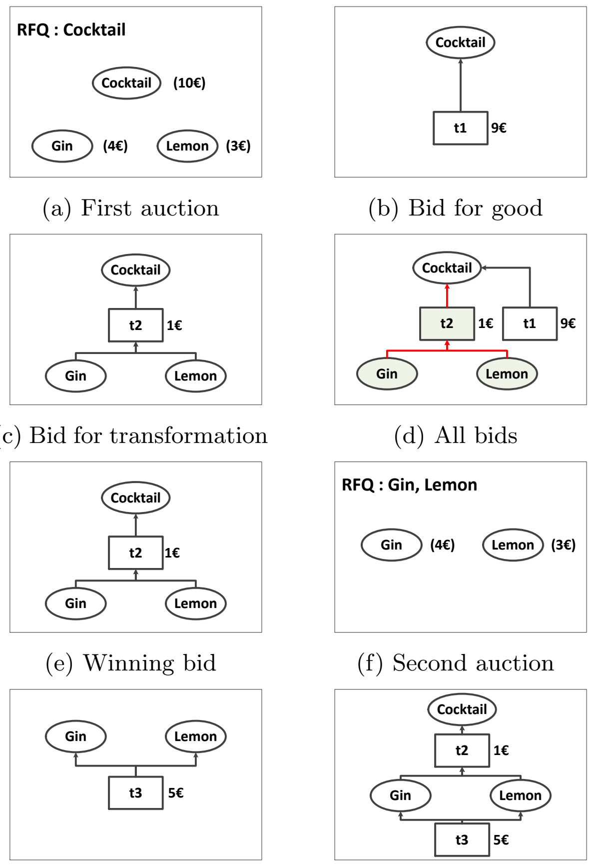

(a) First auction (b) Bid for good (c) Bid for transformation (d) All bids (e) Winning bid (f) Second auction (g) Combinatorial bid for goods (h) Resulting supply chain

Figure 1: Example of sequential mixed auction.

can only be used whenever some provider(s) offer gin and lemon. Thus, the auctioneer assesses the *expected price* of the bid using the market prices of gin (4€) and lemon (3€). Since the expected price is $8(= 1 + 4 + 3)$€, the auctioneer chooses this bid as the winning bid and discards the bid in figure 1b, namely buying the cocktail for 9€.

At this point, the structure of the supply chain is the one depicted in figure 1e. Nonetheless, in order to run the supply chain, the auctioneer must still find providers of gin and lemon. With this aim, the auctioneer starts a new auction of the SMA by issuing an RFQ for gin and lemon (figure 1f). According to our example, this time the auctioneer only receives the combinatorial bid in figure 1g, which offers both lemon and gin for 5€. Since the bid is cheaper than the overall market price of both gin and lemon (4€+3€), this bid is selected as the winning bid of the second auction. Figure 1h shows the resulting structure of the supply chain after the second auction. Since there are no further goods to allocate, the auctioneer closes the SMA. The resulting supply chain produces a cocktail at the cost of 6€.

Although the SMA in this example obtains the optimal solution, this is not always the case. In general, at the end of each auction the auctioneer discards some bids because other bids are *expected* to lead to cheaper solutions. For instance, the bid in figure 1b is discarded to favour the bid in figure 1c. Therefore, since discarded bids might eventually lead to better solutions during subsequent auctions, unlike an MMUCA, an SMA is not guaranteed to obtain an optimal solution (sequence of transformations). Although SMAs may lose optimality, the example anticipates how an SMA help cope with computational complexity and bidders' uncertainty. Firstly, an SMA breaks the formation of a supply chain into several auctions, instead of running a single auction with all bids as MMUCA does. Secondly, after each auction in an SMA, bidders are informed about the needs of the supply chain. Therefore, the auctioneer guides bidders after each tier of the supply chain is formed, hence reducing their uncertainty with respect to participating in MMUCAs (MMUCA bidders only know the expected final outcome of the supply chain!).

3.2 Defining and solving the winner determination problem

As mentioned above, an SMA is essentially a sequence of auctions. For instance, the SMA in figure 1 is composed of two consecutive auctions. Each auction in the sequence receives a set of stock goods and final goods along with the bids submitted by bidders. Then the auctioneer solves the WDP to assess the winning bids as well as the remaining stock goods and required final goods, which are passed on to the next auction in the sequence. When solving the WDP, we assume that the auctioneer is aware of the market prices of goods (recall the example in figure 1) so that it can compute the expected price of bids when necessary.

The sequence of auctions continues till an auction either: (i) obtains a set of winning bids that produce the required goods while consuming all the stock goods; or (ii) does not receive any bids that can improve the supply chain (there are no winning bids). At this point, each auction in the sequence has assessed its winning bids. Thereafter, the transformations in the bids must be merged to compose a solution (a sequence of transformations) for the SMA. The merging must ensure that each transformation in the solution sequence has enough input goods available.

Next, we focus on: (i) formally defining the WDP faced in each auction; (ii) solving the WDP by casting it as an MMUCA WDP. Moreover, we prove that it is possible to guarantee that an SMA solution can be composed from the solutions of the auctions in the sequence. To provide such guarantee, we observe that the winnig bids in each auction define an order on goods that subsequent auctions must fulfill. For instance, consider the example in figure 1. After the clearing of the first auction, a *cocktail* is obtained after *gin*. Thus, during the second auction, we disallow that bidders offer to produce *gin* out of a *cocktail*. Enforcing that the bids in each auction abide by the goods ordering established in previous auctions will allow us to merge auction solutions into an SMA solution.

For the formal definition of the WDP, we restrict ourselves to bids in the XOR-language, which is known to be fully expressive. For each bidder i, let Bid_{ij} be the jth atomic bid occurring within the XOR-bid submitted by i. Recall that each atomic bid consists of a multi-set of transformations and a price: $Bid_{ij} = (\mathcal{D}_{ij}, p_{ij})$, where $\mathcal{D}_{ij} \in \mathbb{N}^{(\mathbb{N}^G \times \mathbb{N}^G)}$ and $p_{ij} \in \mathbb{R}$. For each Bid_{ij}, let t_{ijk} be a unique label for the kth transformation in $\mathcal{D}_{ij}$ (for some arbitrary but fixed ordering of $\mathcal{D}_{ij}$). Let $(\mathcal{I}_{ijk}, \mathcal{O}_{ijk})$ be the actual transformation labeled by t_{ijk}. Finally, we assume that each auction knows the expected market prices at which goods can be bought ($P^- : G \to \mathbb{R}$) and sold ($P^+ : G \to \mathbb{R}$).

We are now ready to define under what circumstances a sequence of transformations constitutes a valid solution for an auction in a sequence:

Definition 1. Valid solution (l-th auction). Given a multi-set $\mathcal{U}^l_{in}$ of available goods and a multi-set $\mathcal{U}^l_{out}$ of required goods, and a partial order on the goods $\prec$, an allocation sequence Σ^l for a given set of XOR-bids over transformations t_{ijk} is said to be a valid solution iff:

1. Σ^l either contains all or none of the transformations belonging to the same atomic bid. That is, the semantics of the BID-operator is being respected:
$$t_{ijk} \in \Sigma^l \;\Rightarrow\; t_{ijk'} \in \Sigma^l$$
2. Σ^l does not contain two transformations belonging to different atomic bids by the same bidder. That is, the semantics of the XOR-operator is being respected:
$$t_{ijk}, t_{ij'k'} \in \Sigma^l \;\Rightarrow\; j = j'$$
3. Each transformation in Σ^l fulfills the partial order on goods $\prec$.
4. For each transformation in Σ^l, either its input goods are in $\mathcal{U}^l_{in}$, its output goods are in $\mathcal{U}^l_{out}$, or both.

In order to assess the expected revenue of a valid solution Σ^l, we must first compute the goods that the auctioneer should buy and sell in the market to implement the solution, namely to *use* all the transformations in the sequence. First, we calculate the units of each good produced by a sequence Σ^l as:

$$\mathcal{P}(g) = \mathcal{U}^l_{in}(g) + \sum_{t_{ijk} \in \Sigma^l} [\mathcal{O}_{ijk}(g) - \mathcal{I}_{ijk}(g)] \tag{1}$$

Hence, we assess the units of each good to buy or sell in the market as:

$$\mathcal{U}^{l+1}_{out}(g) = \mathcal{U}^l_{out}(g) - \mathcal{P}(g)$$

$$\mathcal{U}^{l+1}_{in}(g) = \mathcal{P}(g) - \mathcal{U}^l_{out}(g)$$

Notice that in fact $\mathcal{U}^{l+1}_{out}$ and $\mathcal{U}^{l+1}_{in}$ stand for the remaining required goods and available stock goods that auction l passes on to the next auction.

Definition 2. Expected revenue (l-th auction). The expected revenue for the auctioneer associated with a valid solution Σ^l is the sum of the prices associated with the selected atomic bids plus the expected prices of the goods that are required to be bought ($\mathcal{U}_{out}^{l+1}$) and sold ($\mathcal{U}_{in}^{l+1}$) in the market, namely:

$$\sum\{p_{ij} \mid \exists k : t_{ijk} \in \Sigma^l\} + \sum_{g \in G}[\mathcal{U}_{out}^{l+1}(g)P^-(g) + \mathcal{U}_{in}^{l+1}(g)P^+(g)]$$

Definition 3. WDP (l-th auction). Given a set of XOR-bids, multi-sets $\mathcal{U}_{in}^l$ and $\mathcal{U}_{out}^l$ of initial and final goods, respectively, and a partial order on the goods $\prec$, the winner determination problem is the problem of finding a valid solution Σ^l that maximizes expected revenue for the auctioneer.

Once an optimal valid solution is obtained for the l-th auction, the partial order on goods must be updated to be fed into the next auction. Thus, for each transformation $(\mathcal{I}_{ijk}, \mathcal{O}_{ijk}) \in \Sigma^l$, and for every pair of goods (g, g') such that $g \in \mathcal{I}_{ijk}$ and $g' \in \mathcal{O}_{ijk}$, the relationship $g \prec g'$ is added to the partial order.

We say that a sequence of auctions is *complete* when the last auction in the sequence either: (i) has a valid solution that produces all the required goods and consumes all the stock goods ($\mathcal{U}_{out}^{l+1} = \{\}$ and $\mathcal{U}_{in}^{l+1} = \{\}$); or (ii) does not obtain any valid solution ($\Sigma^l = \{\}$). At this point, the solution for the SMA must be computed from the valid solutions of the auctions in the sequence. First, we characterize a valid solution for an SMA.

Definition 4. Valid solution (SMA) Given a multi-set $\mathcal{U}_{in}$ of available goods and a multi-set $\mathcal{U}_{out}$ of required goods, and a complete sequence of m auctions, an allocation sequence Σ over the transformations in the solutions $\Sigma^1, \ldots, \Sigma^m$ is said to be a valid solution iff:

1. Every transformation in Σ has enough input goods available, and hence can be used, when the sequence starts with the multi-set $\mathcal{U}_{in} + \mathcal{U}_{out}^{m+1}$ of available goods, where $\mathcal{U}_{out}^{m+1}$ stands for remaining required goods after the last auction.
2. The sequence produces the multi-set $\mathcal{U}_{out} + \mathcal{U}_{in}^{m+1}$ of goods, where $\mathcal{U}_{in}^{m+1}$ stands for the remaining available goods after the last auction.

The following lemma ensures that we can compose a solution for an SMA from the solutions of its auctions. Because of lack of space, the proof of the lemma is detailed in [5].

LEMMA THEOREM 1. *The transformations in the valid solutions of a complete sequence of auctions of an SMA can be ordered into a solution for the SMA.*

Finally, we focus on solving the WDP for an auction in an SMA. Firstly, in practice, to ensure that the bids submitted by bidders can only produce valid solutions, we discard bids that do not comply with the partial order on goods. Secondly, to allow a WDP solver to assess the expected price of bids we generate a collection of *phantom* bids. The purpose of phantom bids is to guarantee that there are enough input goods and enough output goods for all the bids received by the auctioneer. Thus, they are meant to stand for potential market bids. In general, phantom bids will be of the type $BID((\{\},\{g\}), P^-(g))$ and $BID((\{g'\},\{\}), P^+(g))$, representing offers to sell some required good and buy some stock good respectively. At this point, the set of bids (received bids as well as phantom bids), the required goods and the input goods can be fed into a WDP solver for MMUCA, such as the computationally-efficient solver in [3].

4. EMPIRICAL ANALYSIS

One of the reasons to introduce SMAs is to reduce the computational complexity of solving MMUCAs. Opposite to MMUCAs, SMAs provide a non-optimal solution to the MMUCA problem. In this section we evaluate the trade-off between solution quality and solving time provided by SMAs with respect to a state-of-the-art optimal MMUCA solver as the size of the problems increases.

4.1 Experimental settings

Since no real world examples of MMUCA problems are available we have artificially generated MMUCA problems following the recommendations in [6]. For an initial evaluation of SMAs, we generated auctions in simple scenarios. Particularly, we generated problems with: (i) only one atomic bid in every XOR-bid, (ii) only one transformation in each atomic bid, (iii) six goods, and (iv) only one unit of each good.

We compare CCIP, an optimal state-of-the-art MMUCA solver introduced in [3], with our recently introduced sequential solver SMA. Note that SMA takes advantage of the existence of a market for goods. The offers coming from this market are incorporated into the auction via phantom bids. For some problems, the existence of phantom bids leads to better solutions. Hence, to make a fair comparison, phantom bids are also included into the CCIP MMUCA solver.

To analyze the evolution of performance as the problem size increases, we increase the number of transformations from 60 to 200 in steps of 10. We generate 100 different problem instances for each different problem size.

4.2 Results

In order to compare solving times, in Figure 2a we plot the median, 75% percentile and 25% percentile of the solving times across the 100 instances generated for each problem size. The speed-up ranges between 4 and 6 times, growing larger as the number of supply chain operations (and hence the problem size) increases.

To benchmark how close the quality of the approximate solution provided by SMA (x_{SMA}) is to that of the optimal one (x^*), we need to identify a baseline solution. In our case, this baseline (x_B) is the solution where every requested good is bought at market price and every stocked and unrequested good is sold at market price. The quality of the solution

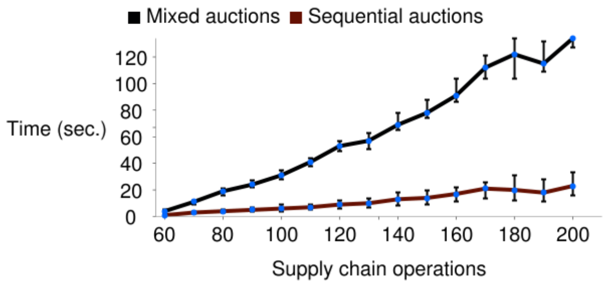

(a) CCIP and SMA times

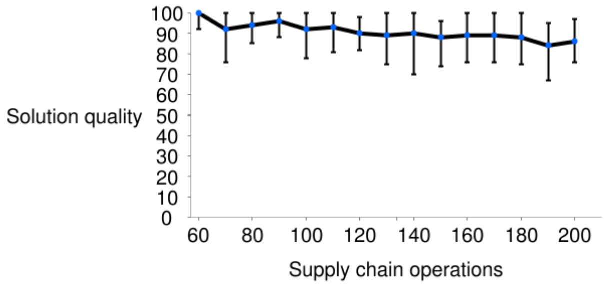

(b) Quality of SMA solution

Figure 2: Experimental results

found by SMA is assessed as

$$Quality(x_{SMA}) = 100\% \cdot \frac{R(x_{SMA}) - R(x_B)}{R(x^*) - R(x_B)}. \quad (2)$$

Thus, the quality of the optimal solution is 100% and the quality of our baseline solution is 0%.

Figure 2b shows the median, 75% percentile and 25% percentile of the solution quality for SMA. We observe that the median quality never falls below 85% in our experiments. These results, although preliminary, show that SMAs decrease computational complexity while keeping a reasonably high accuracy. Further empirical analysis needs to be done to ascertain the impact of the spread of market prices and of the number of goods in the reported trade-off.

5. CONCLUSIONS AND FUTURE WORK

In this work, we have attempted to improve supply chain formation mechanism of Mixed Multi-Unit Combinatorial Auctions (MMUCAs) to make it applicable to real-world procurement scenarios.

To cope with extensive computing times of MMUCA and bidder's uncertainties we moved from single auction to a sequence of auctions that form a supply chain. At each auction bidders are only allowed to bid on transformations that consume available goods or produce requested goods. After selecting the best set of transformations, the auctioneer updates the set of requested and available goods. The sequence ends when supply chain cannot be further improved.

As at each auction bidders are aware of which goods are required and available at current supply chain stage, they only get involved when their services are in need and form more efficient bids. There is no more uncertainty for providers, especially at intermediate levels.

Each auction deals with just a small part of supply chain. Consequently, while solving WDP for individual auction we deal with small subsets of bidders, goods and transformations of former MMUCA. Preliminary tests have shown decrease of solution times up to 6 times.

One of our concerns about SMA mechanism was the solution quality. However, our experimental work showed that median quality does not fall below 85% .

As part of our future work we will explore the possibility to keep n solutions from auction i to auction $i+1$. This should increase solution accuracy significantly.

Our next task in future is to extend empirical analysis. It is important to test how different supply chain structures affect SMA solutions, which should help identify in what scenarios auctioneer benefits the most.

Also, as currently both MMUCA and SMA solvers are based on integer programming, it might be very interesting to analyze how the picture changes when we implement heuristic algorithms to improve computing times.

Acknowledgements

This work has been funded by projects AT (CSD2007-0022) and EVE (TIN2009-14702-C02-01).

6. REFERENCES

[1] Jesus Cerquides, Ulle Endriss, Andrea Giovannucci, and Juan Antonio Rodriguez-Aguilar. Bidding languages and winner determination for mixed multi-unit combinatorial auctions. In *IJCAI*, pages 1221–1226, Hyderabad, India, 2007.

[2] Peter Cramton, Yoav Shoham, and Richard Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2006.

[3] Andrea Giovannucci, M. Vinyals, J. Cerquides, and J. A. Rodriguez-Aguilar. Computationally-efficient winner determination for mixed multi-unit combinatorial auctions. In *AAMAS*, pages 1071–1078, Estoril, Portugal, May 12-16 2008.

[4] Noam Nisan. *Bidding Languages for Combinatorial Auctions*, chapter 9. Combinatorial Auctions. MIT Press, 2006.

[5] Juan Antonio Rodriguez-Aguilar, Boris Mikhaylov, and Jesus Cerquides. Sequential mixed auctions for supply chain formation. Technical report, IIIA, 2011. Available at *http://mmuca.com/report.pdf*.

[6] Meritxell Vinyals, Andrea Giovannucci, Jesús Cerquides, Pedro Meseguer, and Juan A. Rodriguez-Aguilar. A test suite for the evaluation of mixed multi-unit combinatorial auctions. *Journal of Algorithms*, 63(1-3):130 – 150, 2008.

[7] W.E. Walsh and M.P. Wellman. Decentralized supply chain formation: A market protocol and competitive equilibrium analysis. *Journal of Artificial Intelligence Research*, 19:513–567, 2003.