Using multi-attribute combinatorial auctions for resource allocation

Ferran Torrent-Fontbona, Albert Pla, and Beatriz Lopez

University of Girona, Girona 17071, Spain,
{ferran.torrent,albert.pla,beatriz.lopez}@udg.edu,
WWW home page: eia.udg.es/~ftorrent

Abstract. Social concerns about the environment and global warming suggest that industries must focus on reducing energy consumption, due to its social impact and changing laws. Furthermore, the smart grid will bring time-dependent tariffs that pose new challenges to the optimisation of resource allocation. In this paper we address the problem of optimising energy consumption in manufacturing processes by means of multi-attribute combinatorial auctions, so that resource price, delivery time, and energy consumed (and therefore environmental impact) are minimised. The proposed mechanism is tested with simulated data based on real examples, showing the impact of incorporating energy into task allocation problems. It is then compared with a sequential auction method.

Keywords: auctions, multi-attribute, smart grid, energy, resource allocation

1 Introduction

In the coming years it will become crucial to incorporate energy into manufacturing process management due to environmental concerns, time-dependent electricity prices (see Figure 1) and new legislation (i.e. legislation based on energy related standards such as ISO:50001). Smart grids will use these variable prices to reduce overall energy requirements by, for example, filling valleys or cutting peaks in the energy load (see Figure 2), contributing to more sustainable use of energy. As a consequence, the problem of allocating resources to tasks needs to be revised from the energy point of view. In this regard, some previous works have claimed that the resource allocation problem is apt to be redesigned to take account of energy use [11], and some researchers have started to look for solutions in market-based frameworks, such as auctions [23]. However, in [23] the authors follow a single criteria optimisation formulation, considering energy consumption but not price. Including energy (not only energy costs but energy consumption and/or environmental footprint) in resource allocation alters the problem from a single criterion to a multi-criteria one, so all involved objectives should be handled at once. Moreover, when there are several tasks involved in the allocation problem, the problem itself becomes a combinatorial problem, not only due to the capacity limitations of the resource agents, but also due to variable energy prices, and modifications in time and energy consumption when an agent is responsible for more than one task.

Our research concerns production scheduling where the arrival of tasks is unknown in advance. Thus, in this paper we propose to allocate tasks under demand in such a way

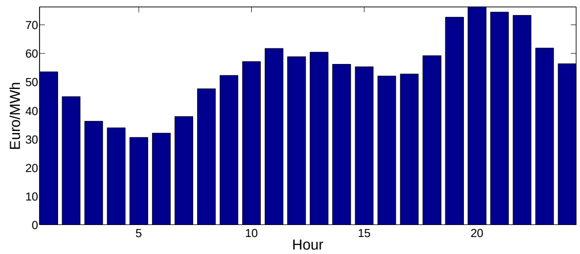

Fig. 1. Average day hourly energy price of the Spanish production market in December 2012 according to [20]. Electricity companies are expected to transfer these variable prices to customers to flatten their energy consumption curves.

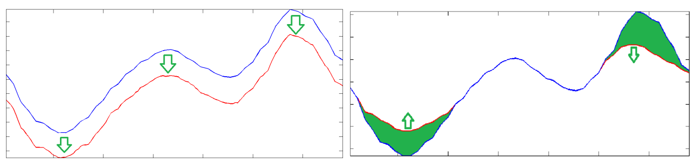

Fig. 2. On the left: illustration of energy efficiency where each curve correspond to a load shape. On the right: illustration of peak reduction and valley filling of an energy consumption curve.

that energy consumption, resource prices and delivery times are taken into account. To this end, our contribution is to solve this allocation problem using an auction mechanism with the following characteristics:

1. Multi-attribute, enabling the auction clearing by handling all the objectives involved: energy, price and delivery time.
2. Combinatorial, as several tasks can be allocated at once whilst bidders can send combinations of bids with different costs. Regarding energy, for example, a bidder could offer two OR bids, one with an energy consumption equals to 5 kWh for deploying task $T1$, and another one with a consumption of 7 kWh for deploying tasks $T1, T2$; however, it is not interested in performing $T2$ alone. This could be because task $T2$ consumes a lot of energy alone (i.e. it requires warming an engine), but its cost diminishes when performed after task $T1$.

Our starting point is the VMA2 auction framework [18], which enables us to deal with multi-attributes auctions, and the combinatorial auction approach with energy issues described in [23]. From these two previous work we define a new auction mechanism that we call a Multi-Attribute Combinatorial Auction (MACA). The main contribution of the paper is to put together these previous works, so as we extend the multi-attribute approach of VMA2 to handle bundles of tasks and introduce different attributes to the combinatorial approach of [23]. As a result, with MACA, we are able to allocate resources to tasks handling variable energy costs and other attributes. We also analyse the tool's performance in a real-world scenario.

This paper is organised as follows. In the following section we review some related work. In Section 3, we explain the auction approach. In section 4 we present the results obtained in experimental testing, and prove a discussion of them. In Section 5 we present our conclusions and propose some future work.

2 Related Work

Resource allocation and job scheduling problems are well-known problems which have been the focus of much research. For example, [1] formalizes the problem, considering multi-skill resources and proposing a Branch and Bound (B&B) algorithm to solve it, minimising the makespan; in [26] the problem is also formalized and the authors propose a heuristic method to solve the problem, optimising the task's execution cost and considering stochastic durations of the tasks. Furthermore, [3,4,18,21] propose solutions to the workflow scheduling problem using auctions due to the distributed nature of the context they consider. This is also the case in the present study; however these earlier studies do not consider combinatorial or multi-attribute auctions as we do.

Concerning multi-attribute auctions, a key work is [2], where the author describes different scenarios regarding the payment rule and demonstrates that to achieve incentive compatibility the payment should be derived by matching the evaluation of the payment and the provided attributes with the evaluation obtained by the second best bid. In a later work, [15] proposes an adaptation of the Vickrey-Clarke-Groves method [13] (VCG) for multi-attribute auctions under an iterative schema (bidders are allowed to modify their bids in response to the bids from other agents). In our work, we use a similar approach to determine the auction winner and its payment, however we do not allow iteration. In practice, bid iteration leads to a slower procedure due to the increase of communications and a possible loss of privacy for bidders, who may not want to reveal their offers to competitors. These drawbacks may be acceptable in cases where auctions appear only occasionally and where losing an auction might lead to a long period without workload for bidders. However, in our problem the allocation of resources to tasks is performed on a continuous basis on the arrival of new tasks; therefore we prefer to use Vickrey auctions, which provide equivalent results in a more straight forward mechanism [22]. Another interesting approach for multi-attribute auctions is VMA2 [18], which allows auctioning tasks and resources based on different kinds of attributes which can be defined by bidders and by attributes. However, VMA2 is intended for auctioning single and isolated tasks whilst we aim to auction bundles of tasks.

Public institutions are making great efforts to design and develop a future smart grid [5,6]. Moreover, many researchers are focusing on developing new household management systems that deal with time-dependent rates, [14], studying consumers' behaviour when faced with variable prices [7,10,14] and studying and designing a new negotiation system between electricity companies (distributors), producers and consumers [24]. Despite this great research, little work has been done relating to workflow management considering time dependent energy rates. In [9] Simonis and Hadzic developed some lower bounds based on cumulative constraints to use with the problem solving algorithm. In [11] the authors consider time-dependent energy rates in a workflow context and propose a solution to the scheduling problem by using reverse auctions, presenting a new formalization of the problem. In any of these previous approaches the consumption agreement is combined with the time-dependent rates, as we are doing here.

3 Methodology

In this paper we deal with job scheduling and resource allocation having in mind the new challenges posed by the smart grid and the environmental impact of the performance of tasks.

In particular, we are dealing with the problem of allocation of resources to tasks assuming a dynamic environment, such that tasks are unknown in advance of their arrival. In this scenario, an agent is in charge of handling task arrivals and assigning appropriate resources to carry out tasks. At a given moment of time, there are multiple tasks to be performed, each with different requirements. Resources that can deploy the tasks are handled by other agents. Resources are allocated to tasks following an auction protocol.

An auction is a method for buying and selling goods or items using a bid system in which the best bids obtain the sold items. In domains where the aim is to allocate or outsource tasks to third party companies it is common to follow a reverse auction schema: an auctioneer needs a task to be done and offers to pay an external provider for carrying it (becoming the buyer who aims to buy a service at the cheapest price) whilst bidders offer their working capacity at a given price (becoming the sellers who compete to offer the best working conditions at the cheapest price). This reverse auction schema is the one followed in this research.

The auction approach is of particular interest when we tackle allocation of energy consuming tasks under variable energy costs. In this case, auctions offer bidders the chance to handle energy costs for tasks, leaving the assignment process to the auctioneer: bidders provide offers to deploy tasks at a given time, at a given price and with the energy costs they would incur; thus, no alternatives other than those provided in the bid would be considered by the auctioneer.

However, the management of multiple attributes other than price (e.g. energy consumption and delivery times) requires a multi-criteria decision. Moreover, the dependencies between attributes and bidder's schedules (e.g. the time when a task is being performed conditions its costs due to variable energy prices) will push bidders to submit multiple bids with different attribute configurations. In consequence, we need to use a combinatorial multi-attribute auction mechanism. The mechanism is described below, according to the 4 main steps of the protocol: call for proposals, bidding, determining the winner, and payment. We consider companies as agents that act from self-interest in order to increase their own utility. They will aim to outsource tasks on the best possible terms (auctioneer agents) or they will aim to sell their resources in order to perform tasks at the highest prices for the lowest effort (bidder agents).

3.1 Call for proposals

When an auctioneer needs to outsource a task it sends a call for proposals indicating the different tasks constraints and the required skills $\mathbf{RQ}_i$ to all the bidders $(a_1 ... a_n)$ inside the market. Each set of tasks is defined as a set of independent tasks $\mathbf{T} = \left\{T_1 \ldots T_{|\mathbf{T}|}\right\}$. Each task is defined as follows:

$$T_i = \langle \left[\underline{s_i}, \overline{s_i}\right], \left[\underline{et_i}, \overline{et_i}\right], \mathbf{RQ}_i \rangle \tag{1}$$

where $\underline{s_i}$ is the task earliest start time , and $\overline{s_i}$ the latest start time; $\underline{et_i}$ the earliest end time and $\overline{et_i}$ the latest end time; and $\mathbf{RQ}_i$ is a list with the resource skills required by the task. All of this parameters ($\left[\underline{s_i}, \overline{s_i}\right], \left[\underline{et_i}, \overline{et_i}\right], \mathbf{RQ}_i$) constitute the task constraints of our problem. Bidders aiming to perform a certain task need to have available resources with the required skills, otherwise they will be unable to perform the task. On the other hand, they should provide actual starting times for tasks and duration that agree with the task time windows $\left[\underline{s_i}, \overline{s_i}\right]$ and $\left[\underline{et_i}, \overline{et_i}\right]$.

3.2 Bidding

Once a bidder receives the auctioneer's proposal, if the bidder is interested in any of the auctioned tasks and is able to provide an offer according to the task's constraints, it offers a bundle of bids where each bid describes possible conditions (price, energy consumption and delivery time) under which the bidder can perform the task. It is worth noting that in doing so, each resource agent has its own energy constraints and resource capacity constraints, which are opaque to other agents, and which are summarised in the bids.

Every bidder can send several bids with different configurations for the same task, because (due to variable energy prices) the cost of performing a task may change depending on the time it is scheduled and on other tasks the bidder could be assigned to perform. This leads to combinatorial auctions, meaning that agents bid bundles of tasks at different prices and conditions. We followed the notation presented in [23] to express combinatorial bids where the kth bid proposed by the jth bidder to perform the ith task is defined as

$$B_{i,j,k} = \langle T_i @ s_{i,j,k} : \left(\mu_{i,j,k}, \epsilon_{i,j,k}, \delta_{i,j,k}\right), M_{i,j,k}, E_{i,j,k}, \varDelta_{i,j,k} \rangle \tag{2}$$

where T_i is the ith task to which the bid is submitted, $s_{i,j,k}$ is the start time proposed by the bidder, $\mu_{i,j,k}$ is the price of the bid, $\epsilon_{i,j,k}$ is the energy consumption and $\delta_{i,j,k}$ is the duration; $M_{i,j,k}$, $E_{i,j,k}$ and $\varDelta_{i,j,k}$ are $N \times 1$ vectors that indicate modifications on the price, energy consumption and duration (respectively) if the bid is accepted together with another bid of the same bidder. In this way $E_{i,j,k}\left(l\right)$ indicates a modification on the energy consumption of $B_{i,j,k}$ if the lth bid of bidder i is also accepted to perform its corresponding task.

In our work we consider three attributes: price, energy and duration. However, this can be generalised to apply more attributes according to the ethos suggested by [18].

3.3 Winner determination problem

Once bidding period is over, the auctioneer must decide which bids maximise its expected utility [19]. For that purpose it calculates the utility of each bid and seeks the optimal combination of bids with the highest utility. The utility of the auctioneer given a bid $B_{i,j,k}$ for having a task T_i made, is defined as follows:

$$u(T_i, B_{i,j,k}) = v(T_i) - f(B_{i,j,k}) \tag{3}$$

where $v(T_i)$ is the value considered by the auctioneer for having task T_i completed and $f(B_{i,j,k})$ is the cost of bid $B_{i,j,k}$ for the auctioneer considering all the dimensions involved in the allocation (economic cost, ending time and energy consumption). Note that, given T_i, maximising $u(T_i, B_{i,j,k})$ is equivalent to minimising $f(B_{i,j,k})$. Thus, the winner determination problem (WDP) is defined as:

$$argmin_{j,k} \sum_{i,j,k} x_{i,j,k} * f(B_{i,j,k}) \tag{4}$$

Subject to

- $x_{i,j,k} = 1$ if bid $B_{i,j,k}$ is selected; otherwise $x_{i,j,k} = 0$
- Each task is assigned to and executed by a single bidder/bid $\sum_{j,k} x_{i,j,k} = 1, \forall i$
- All tasks constraints are satisfied

However, this minimisation problem is not trivial due to the set-up times regarding $M_{i,j,k}$, $E_{i,j,k}$ and $\Delta_{i,j,k}$. One possible way to simplify the problem is to use auxiliary variables to express the final price $b_{i,j,k}$ of a bid, the final end time $t_{i,j,k}$, and the final energy consumption $e_{i,j,k}$:

$$b_{i,j,k} = \mu_{i,j,k} + \sum_{l=1}^{N_j} M_{i,j,k}(l) \cdot x_{i,j,l} \tag{5}$$

$$t_{i,j,k} = s_{i,j,k} + \delta_{i,j,k} + \sum_{l=1}^{N_j} \Delta_{i,j,k}(l) \cdot x_{i,j,l} \tag{6}$$

$$e_{i,j,k} = \epsilon_{i,j,k} + \sum_{l=1}^{N_j} E_{i,j,k}(l) \cdot x_{i,j,l} \tag{7}$$

where N_j is the number of bids sent by the jth bidder.

The winner determination problem can be then reformulated as follows:

$$argmin_{j,k} \sum_{i,j,k} x_{i,j,k} * f(b_{i,j,k}, t_{i,j,k}, e_{i,j,k}) \tag{8}$$

Subject to the same constraints as above. Note that the minimisation problem considers all tasks ($\forall i$).

Therefore the problem of the determination of the auction winner(s) can be solved by minimising f, which combines the different attributes of bids (price, time and energy), becoming by definition a key issue for the winner determination problem. For the mechanism to be feasible, we consider V (a particular case of f) as an aggregation function which must be a real-valued monotonic bijective function [17]. In particular, in this paper we use the weighted sum but other functions could be considered as well (see [17] for alternative evaluation functions):

$$V(b_{i,j,k}, t_{i,j,k}, e_{i,j,k}) = w_0 \cdot b_{i.j,k} + w_1 \cdot t_{i,j,k} + w_2 \cdot e_{i,j,k} \tag{9}$$

$$\sum_k w_k = 1 \tag{10}$$

The complexity of solving the problem is exponential [23,3], and complete methods cannot provide a solution in a realistic amount of time when the number of tasks and bids increases. Therefore, the use of meta-heuristic methods is a good alternative to obtain near optimal solutions. We decided to use Genetic Algorithm (GA) [12,8], because its use does not involve many mathematical assumptions about the problem (they can handle any kind of objective function and constraint). Also, it is a very effective tool for global search (there is no need for convexity in the objective function).

The GA used represents solutions as strings of bids (chromosomes) where each slot of the string corresponds to a particular task. To create new chromosomes it uses a selection operator, a crossover operator and a mutation operator. The selection operator is 3 tournament selection, [8], which consists in selecting 3 random chromosomes and choosing the best as the 1st parent. The process is repeated for choosing the 2nd parent. Once the parents are selected, it uses the 2 cross-point crossover operator, [8], to create 2 new chromosomes exchanging the genetic information of the parents. Finally, the new chromosomes mutate changing each bid (gene) for another randomly selected with a probability of 5%. To maintain the population, an elitism operator is used. It consists in removing all the chromosomes except the best. The algorithm is explained in Algorithm 1, where N_g is the number of generations and N_p is the size of the population.

Algorithm 1 Genetic Algorithm

Require: $N_g = 2000$, $N_p = 300$

Create N_p random chromosomes
for $g \leftarrow 1$ **to** N_g **do**
  for $i \leftarrow 1$**to** $N_p/2$ **do**
    Select 2 parents using the 3 tournament selection
    Breed two new chromosomes using 2 cross-point crossover
    Apply mutation operator over the new chromosomes
    Compute fitness of the new chromosomes using V_0
  end for
  Elitism: remove all chromosomes except the N_p best
end for
select the best chromosome as solution

3.4 Payment

A payment rule is used to establish the economic amount that auctioneers must pay to the auction winner(s) for performing any task. Given the multi-dimensional nature of the allocation problem we are dealing with, payment is not only conditioned by the bidded economic amounts but also by other attributes. For instance, delivering a task later than agreed may involve receiving less money than the initial bid amount. Moreover, the auctioneer cannot assume that bidders will follow a truthful bidding strategy. In order to encourage bidders to bid truthfully regarding their economic costs the VCG payment mechanism can be used [25]. This payment considers that the payment $p_{i,j,k}$ for bidder j for performing task i according to bid $B_{i,j,k}$ will correspond to the difference of the welfare all bidders would have obtained if the winning bid had not been sent to the auction and the welfare they receive with the chosen allocation excluding the welfare for bid $B_{i,j,k}$. However, such a mechanism considers a single attribute, price, and does not guarantee that bidders deliver tasks to the terms agreed during the bidding process (i.e. due to estimation errors [18]). So we modify the VCG payment mechanism in order to reduce the auctioneer's utility loss when bidders do not deliver tasks to the agreed attributes.

The payment rule proposed is a two case method: on the one hand, when winning bidders are successful (delivering the task as agreed) they receive a payment $p_{i,j,k}$ according to a classical VCG auction schema. On the other hand, if the bidder delivers a task in worst conditions than the agreed (i.e. $t^{'}_{i,j,k}, e^{'}_{i,j,k}$ instead of $t_{i,j,k}, e_{i,j,k}$), it will receive a smaller payment in such a way that the valuation of the obtained payment $p_{i,j,k}$ and the delivered attributes matches the valuation of the initially presented bid, as follows:

$$V(p_{i,j,k}, t'_{i,j,k}, e'_{i,j,k}) = V(b_{i,j,k}, t_{i,j,k}, e_{i,j,k}) \tag{11}$$

where $t'i, j, k$ and $e'_{i,j,k}$ are the true delivery time and the final energy consumption. Therefore the payment is defined as follows:

$$p_{i,j,k} = \begin{cases} V^{-1}\left(\Phi_{i,j,k}, t^{'}_{i,j,k}, e^{'}_{i,j,k}\right) & \text{if } t^{'}_{i,j,k} \prec t_{i,j,k}, e^{'}_{i,j,k} \prec e_{i,j,k} \\ V^{-1}\left(V\left(b_{i,j,k}, t_{i,j,k}, e_{i,j,k}\right), t^{'}_{i,j,k}, e^{'}_{i,j,k}\right) & \text{otherwise} \end{cases} \tag{12}$$

where

$$\begin{aligned} \Phi_{i,j,k} = & \sum_{(l,m,n)\in G_{-(i,j,k)}} V\left(b_{l,m,n}, t_{l,m,n}, e_{l,m,n}\right) - \\ & \sum_{(x,y,z)\in G\setminus(x,y,z)\neq(i,j,k)} V\left(b_{x,y,z}, t_{x,y,z}, e_{x,y,z}\right) \end{aligned} \tag{13}$$

where $\prec$ means *worse than*, G is the set of winning bids, $G_{-(i,j,k)}$ is the set of bids that would have won the auction if bid $B_{i,j,k}$ had not been sent, $G\setminus(x, y, z) \neq (i, j, k)$ indicates the set of winning bids different to $B_{i,j,k}$ and where

$$V^{-1}\left(\Phi_{i,j,k}, t^{'}_{i,j,k}, e^{'}_{i,j,k}\right)$$

is the reverse function of $V(b_{i,j,k}, t_{i,j,k}, e_{i,j,k}) = x$ which given $x, t_{i,j,k}, e_{i,j,k}$ returns $b_{i,j,k}$. Note that for achieving the set $G_{-(i,j,k)}$, we need to resolve the WDP but removing the bid $B_{i,j,k}$.

	T_1				T_2				T_3			
	b	e	t	V_0	b	e	t	V_0	b	e	t	V_0
Bidder 1	**20**	**5**	**5**	**10**	10	7	7	8	**6**	**5**	**3**	**5**
Bidder 2	20	10	6	12	**7** (5)	**5** (4)	**3** (3)	**5** (4)	10	7	7	8
Bidder 3	25	10	10	15	8	8	5	7	15	10	5	10

Table 1. Example of 3 bidders bidding for 3 different tasks. It shows the values of the attributes and the global value of the bids considering the weighted sum (with all weights equal to $\frac{1}{3}$ in Equation 9). Winning bids are in bold face. Numbers in brackets correspond to the bid values (considering set-up costs) if tasks T_1 and T_2 are assigned to bidder 2.

In this way bidders are encouraged to bid truthfully: on the one hand, if they underbid regarding any attribute, bidders do not increase their utility (and they could lose utility, because if they win they are forced to work under the bid conditions). On the other hand, overbidding will reduce their chances of winning the auction. Finally, underdelivering confers a payment reduction which will reduce the bidder's utility (encouraging it to improve its attribute estimation) whilst avoiding a loss of utility from the auctioneer's side (for instance, paying less to the winning bidder will allow the auctioneer to hire better resources in future).

Example 1 (Payment rule example).
Consider the example of Table 1 where three different bidders have sent three bids each (for three tasks), and where the evaluation function V of the auctioneer is a weighted sum with all the weights set to $w = \frac{1}{3}$. According to the values of Table 1, bidder 1 is the winner for performing tasks T_1 and T_3, and bidder 2 is the winner for T_2.

When a task is delivered to the agreed conditions, the payment to the bidder for performing a task according to a particular bid is computed according to Equations (12) and (13). First, we compute the payment of bidder 1: $\Phi_{1,1,1}$ is the difference between the valuations of bids $B_{2,1,1}$, $B_{2,2,2}$ and $B_{3,1,3}$ (winning bids if bid $B_{1,1,1}$ had not been sent) and the valuations of bids $B_{2,2,2}$ and $B_{3,1,3}$ (winning bids except $B_{1,1,1}$). Thus, considering Table 1,

$$\Phi_{1,1,1} = (12 + 4 + 5) - (5 + 5) = 11 \tag{14}$$

Note that when we consider that bid $B_{1,1,1}$ is not sent, we have to consider set-up costs of bid $B_{2,2,2}$ because T_1 would have been assigned to bidder 2. Then, $b_{2,2,2} = 5$, $t_{2,2,2} = 3$, $e_{2,2,2} = 4$ and $V(5, 3, 4) = 4$.

Then, the payment $p_{1,1,1}$ corresponding to bidder 1 for doing task 1 according to $B_{1,1,1}$ is calculated according to Equation (12) as follows:

$$p_{1,1,1} = \frac{\Phi_{1,1,1}}{w} - \left(t'_{1,1,1} + e'_{1,1,1}\right) = \frac{11}{0.33} - (5 + 5) = 23 \tag{15}$$

Similarly, payments corresponding to bids $B_{2,2,2}$ and $B_{3,1,3}$ are 13 and 16 respectively if the tasks are delivered to the agreed conditions.

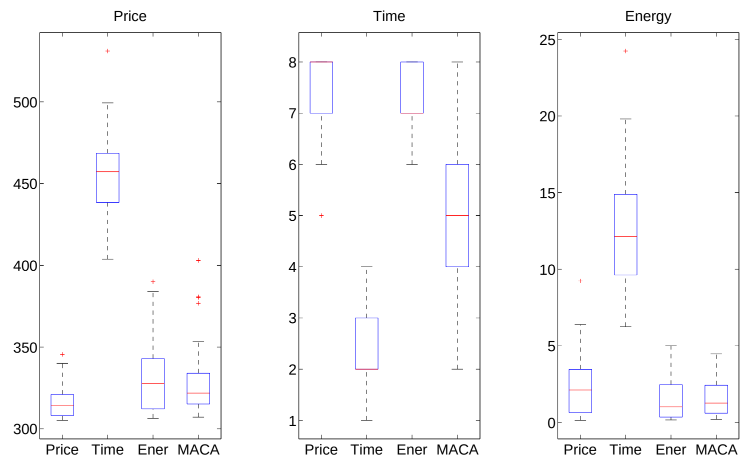

Fig. 3. Values of the attributes of the winning bids when we optimise a single attribute (monetary cost, time or energy) or the aggregation of all of them (horizontal axis). Y axis: (left) price, (center) time, (right) energy

However, if we assume that bidder 2 does task T_2 with an energy consumption of $e'_{2,2} = 8$ instead of 5, the corresponding payment is calculated according the second branch of Equation (12). Thus,

$$p_{2,2,2} = \frac{V\left(b_{1,1,1}, t_{1,1,1}, e_{1,1,1}\right)}{w} - \left(t'_{1,1,1} + e'_{1,1,1}\right) = \frac{5}{0.33} - (3 + 8) = 4 \qquad (16)$$

So, bidder 2 would receive a payment of 4 instead of 13 for not fulfilling the agreed energy consumption.

4 Experimentation

In this section we analyse the performance of the presented methodology using a multi-agent simulator based on real data [16]. First, we analyse the results when we perform uni-criteria allocations and when we perform a multi-criteria allocation. Second, we analyse the allocation results obtained with the multi-attribute combinatorial auction mechanism proposed in this paper and we compare them with the results obtained with VMA2, a multi-attribute (but) sequential auction (one task auctioned after the other). Finally we discuss the results achieved regarding our research objectives.

4.1 Experimentation set-up

The data over which we conducted experimentation is based on a real industry process[1]. Tasks are managed by an agent (auctioneer) that outsources some of its tasks to 7 other

[1] Data available at `http://eia.udg.es/~apla/fac_data/`

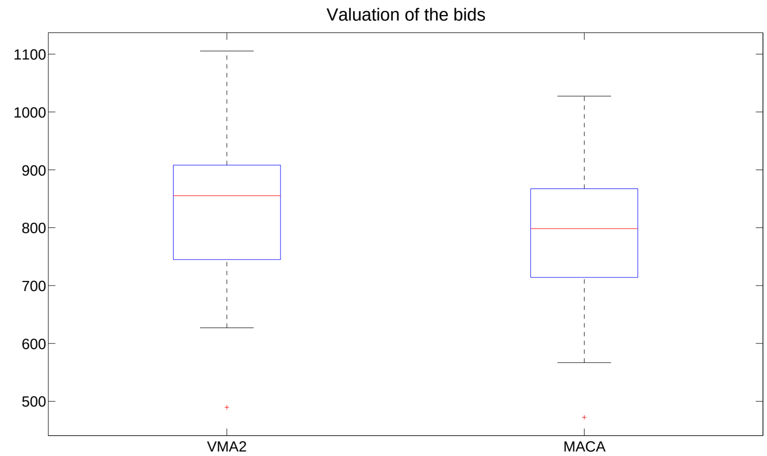

Fig. 4. Comparison of the average aggregated cost of the winning bids when using MACA or VMA2.

agents (bidders) with different skills. Each task requires a particular skill and conveys an economic cost, a particular execution time and an energy consumption. The tasks and attribute values are obtained from probability distribution functions which were modelled using data from real business processes[1]. Each bidder is assigned a particular energy tariff which conveys variable energy prices. Agents' behaviour is modelled as competitive and greedy.

4.2 Experiment 1: uni-attribute versus multi-attribute combinatorial auctions

The goal of the experiment is to point out the importance of aggregating all the objectives that an organisation needs to consider, especially, when they cannot be optimised simultaneously. Then the use of aggregation functions provides solutions with a trade-off between the objectives.

For that purpose, in this experiment we compare the allocation of the tasks of a single day considering uni-criteria (combinatorial auctions) and multi-criteria valuation functions (MACA). We computed the resulting task allocation using our auction mechanism considering a uni-attribute approach (considering only the price, or the delivery time, or the energy consumption of the bids in Equations (8) and (12)) and using a multi-criteria approach (determining the auction winner using aggregation function V). The experiment was conducted over 50 sets of tasks.

Figure 3 shows the box plot of the attributes of the winning bids when the auctioneer wants to optimise a single attribute (price, or time, or energy) or the aggregation of all of them. It points out that in this experiment, it is impossible to optimise all the attributes, i.e. the optimisation of time greatly increases the price and energy consumption. However, when we aggregate all attributes, the obtained allocation is a trade-off

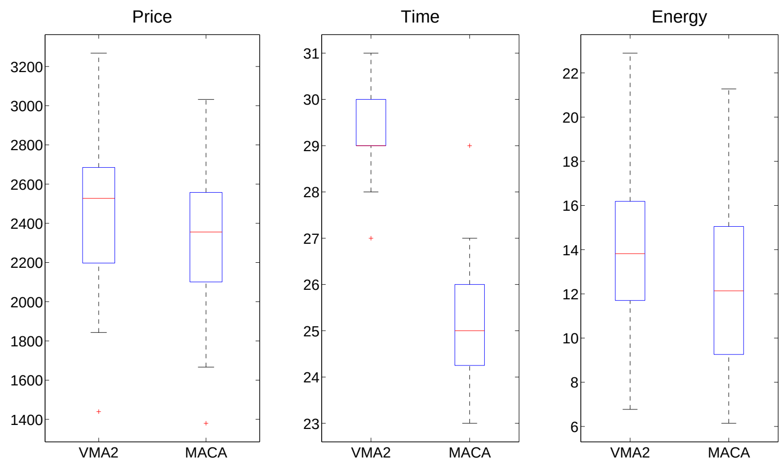

Fig. 5. Comparison of the values of the attributes of the the winner bids when using MACA or VMA2. The values for VMA2 consist of the aggregation of the results of auctioning one after the other the same tasks than in MACA.

between the objectives. Such trade-off is determined by the aggregation function. For example, we see in Figure 3 that optimising the aggregation of all attributes, reports a solution in terms of price and energy very close to the optimal; and in terms of time the solution is between the optimal and the solutions obtained when we optimise only either price or energy (which are far worse than when we only optimise time).

ANOVA analysis over the values of price, time and energy of Figure 3 shows that the results obtained optimising different attributes can be considered that come from different distributions with p-values lower than 10^{-73}. Even the results from optimising either price, energy, or the aggregation (MACA) are different, with p-values lower than 10^{-2}. Even paired-response tests tell us that with a significance value of 0.05 we can consider that the values of either price or time or energy, obtained when we optimise the either price or time or energy respectively, are better than we optimise another objective.

4.3 Experiment 2: multi-attribute sequential versus multi-attribute combinatorial auctions

Experiment 2 compared the performance of VMA2 and MACA. VMA2 [18] auctions one task at a time (sequential auction). Therefore, the order in which tasks are auctioned could affect the results. On the other hand, with the approach presented in this paper, MACA, all the tasks are auctioned at the same time (combinatorial). Although the benefits of combinatorial auctions as compared to sequential ones are very well known, in this paper we are considering multi-attribute auctions, both in VMA2 and MACA. In particular, experiment 2 is used to point to the cost differences when using each method.

For that purpose, we computed and compared the allocation for the tasks of a single day using a multi-attribute combinatorial auction approach (MACA) and a multi-attribute sequential auction approach (VMA2). Experiments were also repeated 50 times to obtain meaningful results. To compare VMA2 and MACA we auctioned the same tasks, but VMA2 auctioned them sequentially and MACA auctioned them concurrently. To compare the results we aggregated the results of VMA2. We calculated the makespan, for VMA2, as the difference between the ending of the last task and the auction time of the first time. MACA computes the makespan as the difference between the end time of the last task performed and the auction time.

Figures 4 and 5 show the results obtained in this scenario. As we expected, MACA outperforms VMA2 in terms of aggregated cost (price, time and energy) because it is able to consider bundles of tasks and is therefore, able to provide better allocation. We also tested the results with pair-response tests, which showed that we can assume that the aggregated cost of the winning bids when using MACA is lower than when using VMA2 at the significance level of 0.05. ANOVA analysis also discards that both collection of values come from a population with the same mean with a p-value of 0.0472.

Regarding the values of the attributes, pair-response tests also show that with a significance level of 0.05 we can assume that the values of the attributes using MACA and GA are better than VMA2. ANOVA analysis also discard that results of MACA and VMA2 regarding the values of the attributes come from populations with the same mean with p-values of 0.0417 (for price), $1.25 \cdot 10^{-35}$ (for makespan) and 0.012 (for energy consumption).

4.4 Discussion

Results obtained in the first experiments show that our auction mechanism, MACA, is able to deal with several objectives at a time, meaning that we are able to take account of energy issues when allocating resources to tasks. On the other hand, the second experiment corroborates the benefits of using MACA to deal with bundles of tasks, handling set-up constraints regarding different attributes (time, money, energy), which results in a better outcome than assigning tasks in a sequential way. Therefore, our new mechanism MACA improves upon previous mechanisms (such as VMA2), whilst meeting the multi-criteria requirements that enables it to deal with energy issues. Energy consumption is handled by the auctioneer, while the details of variable prices of energy usage are handled in private by the bidders. The auctioneer does not take into account energy usage constraints, it just seeks the best allocation of resources taking into account the attributes of the bids.

In future work it would be productive to include task precedence constraints within the auction model. To address this issue it is important to study how this precedence can be modeled within the call for proposals and the bidding steps, but also how delays might affect future tasks, how they might provoke bid withdrawals, and how these issues should be considered in the payment mechanism.

5 Conclusions

This paper has presented a multi-attribute combinatorial auction mechanism for allocating bundles of tasks under demand in environments where the production schedule is unknown in advance, and the cost of the performance is variable. In particular, we have applied the presented approach to allocate tasks based on economic cost, delivery time and energy consumption under variable energy prices.

To deal with the problem of allocating tasks based on more than one attribute we have used a multi-attribute auction mechanism which uses a multi-criteria function to establish the winner of the auction. The combinatorial dimension of the auction mechanism proposed enables bidders to handle variable costs for pairs of tasks, whilst the auctioneer focuses on selecting the best bids.

We tested the mechanism in a simulated environment based on real data, and the results show that the presented mechanism is suitable for allocating tasks and resources under demand while considering different attributes and taking advantage of combinatorial bids to reduce cost in terms of price, time and energy. Moreover, they show that the MACA mechanism outperforms VMA2.

The work remains open to further interesting lines of research, for instance, considering precedences between tasks and robustness issues in the winner determination problem and determining payment amounts.

References

1. Odile Bellenguez-Morineau and Emmanuel Néron. A branch-and-bound method for solving multi-skill project scheduling problem. *RAIRO-Operations Research*, 41(02):155–170, 2007.
2. YK Che. Design competition through multidimensional auctions. *rand Journal of Economics*, 1993.
3. John Collins, Güleser Demir, and Maria Gini. Bidtree ordering in IDA combinatorial auction winner-determination with side constraints. *Agent-Mediated Electronic Commerce IV. Designing Mechanisms and Systems*, 2531:17–33, 2002.
4. Laureano F. Escudero, Mercedes Landete, and Alfredo Marín. A branch-and-cut algorithm for the Winner Determination Problem. *Decision Support Systems*, 46(3):649–659, February 2009.
5. European Co-Directorate for Research Co-operation Energy. European technology platform smartgrids - strategic research agenda for europe's electricity networks of the future, 2007.
6. European Co-Directorate General for Research Sustainable Energy Systems. European smartgrids technology platform - vision and strategy for europe's electricity networks of the future, 2006.
7. Sebastian Gottwalt, Wolfgang Ketter, Carsten Block, John Collins, and Christof Weinhardt. Demand side management - A simulation of household behavior under variable prices. *Energy Policy*, 39(12):8163–8174, 2011.
8. RL Haupt and SE Haupt. *Practical genetic algorithms*. John Wiley & Sons, Ltd, 2004.
9. Simonis Helmut and Tarik Hadzic. *A Resource Cost Aware Cumulative*, pages 76–89. ecent Advances in Constraints. Springer Berlin, Heidelberg, lecture no edition, 2011.
10. Wenzhao Jia, Chongqing Kang, and Qixin Chen. Analysis on demand-side interactive response capability for power system dispatch in a smart grid framework. *Electric Power Systems Research*, 90(0):11–17, 2012.
11. Beatriz Lopez, Aditya Ghose, Bastin Tony Roy Savarimuthu, Mariusz Nowostawski, Michael Winikoff, and Stephen Cranefield. Energy-aware optimisation of business processes. In *Proceedings of the 3rd International Workshop on Agent Technologies for Energy Systems (ATES 2012)*, Valencia, December 2012. University of Otago.
12. Sean Luke. Essentials of metaheuristics. *Lecture notes, George Mason University.*, 2013.

13. Jeffrey K. MacKie-Mason and Hal R. Varian. Generalized Vickrey Auctions, July 1994.
14. Amir-Hamed Mohsenian-Rad, Vincent W. S. Wong, Juri Jatskevich, Robert Schober, and Alberto Leon-Garcia. Autonomous Demand-Side Management Based on Game-Theoretic Energy Consumption Scheduling for the Future Smart Grid. *IEEE Transactions on Smart Grid*, 1(3):320–331, December 2010.
15. David C. Parkes and Jayant Kalagnanam. Models for Iterative Multiattribute Procurement Auctions. *Management Science*, 51(3):435–451, March 2005.
16. Albert Pla, Pablo Gay, Joaquim Melndez, and Beatriz Lpez. Petri net-based process monitoring: a workflow management system for process modelling and monitoring. *Journal of Intelligent Manufacturing*, pages 1–16, 2012.
17. Albert Pla, Beatriz Lopez, and Javier Murillo. Multi criteria operators for multi-attribute auctions. *Modeling Decisions for Artificial Intelligence*, 2012.
18. Albert Pla, Beatriz López, Javier Murillo, and Nicolas Maudet. Multi-attribute auctions with different types of attributes: Enacting properties in multi-attribute auctions. *Expert Systems with Applications*, 41(10):4829–4843, August 2014.
19. SD Ramchurn and C Mezzetti. Trust-based mechanisms for robust and efficient task allocation in the presence of execution uncertainty. *Journal of Artificial* . . . , 2009.
20. Red Eléctrica de España. Boletín mensual diciembre 2012. Technical report, Red Eléctrica de España, 2013.
21. Tuomas Sandholm, Subhash Suri, Andrew Gilpin, and David Levine. CABOB: A fast optimal algorithm for winner determination in combinatorial auctions. *Management Science*, 51(3):374–390, 2005.
22. Tuomas W Sandholm. Limitations of the vickrey auction in computational multiagent systems. In *Proceedings of the Second International Conference on Multiagent Systems (ICMAS-96)*, pages 299–306, 1996.
23. Ferran Torrent-Fontbona and Beatriz López. Comparison of workflow scheduling using constraint programming or auctions research report 13-01-rr. Technical Report March, Institute of Informatics and Applications, University of Girona, Girona, 2013.
24. Z A Vale, H Morais, and H Khodr. Intelligent multi-player smart grid management considering distributed energy resources and demand response. In *Power and Energy Society General Meeting, 2010 IEEE*, pages 1–7, 2010.
25. WE Walsh and MP Wellman. A market protocol for decentralized task allocation. *Multi Agent Systems, 1998.* . . . , 1998.
26. Yang Yu, Maolin Pan, Xuguang Li, and Huan Jiang. Tabu search heuristics for workflow resource allocation simulation optimization. *Concurrency and Computation: Practice and Experience*, 23(16):2020–2033, 2011.