

Optimal Combinatorial Electricity Markets

Yoseba K. Peña

Department of Information Systems
Vienna University of Economics and
Business Administration, Austria.
yoseba.pena@wu-wien.ac.at

Nicholas R. Jennings

School of Electronics and Computer Science
University of Southampton
Southampton, UK.
nrj@ecs.soton.ac.uk

Abstract—The deregulation of the electricity industry in many countries has created a number of marketplaces in which producers and consumers can operate in order to more effectively manage and meet their energy needs. To this end, this paper develops a new model for electricity retail where end-use customers choose their supplier from competing electricity retailers. The model is based on simultaneous reverse combinatorial auctions, designed as a second-price sealed-bid multi-item auction with supply function bidding. This model prevents strategic bidding and allows the auctioneer to maximise its payoff. Furthermore, we develop optimal single-item and multi-item algorithms for winner determination in such auctions that are significantly less complex than those currently available in the literature.

KEYWORDS

Multi-Agent Systems, Distributed Artificial Intelligence, Combinatorial Auctions

I. INTRODUCTION

The deregulation of electricity markets began in the early nineties when the UK Government privatised the electricity supply industry in England and Wales. This process has been subsequently followed in many other countries. In most cases, the restructuring involves separating the electricity generation and retail from the natural monopoly functions of transmission and distribution. This, in turn, leads to the establishment of a *wholesale electricity market* for electricity generation and a *retail electricity market* for electricity retailing. In the former case, competing generators offer their electricity output to retailers and in the latter case end-use customers choose their supplier from competing electricity retailers.

Here we focus on retail markets, which differ from their more traditional counterparts because energy cannot be stored or held in stock (as tangible goods can). Consequently, retailers are forced to work with consumption prognoses [1], which, in turn, creates a number of risks. First, producing more than is consumed is not economical. Moreover, the price of the energy mainly depends on the production cost and this typically rises with the amount of energy produced. Second, if the demand exceeds the prediction, suppliers must find additional energy to avoid a blackout. Finally, there are non-negligible costs stemming from the variation in the electricity production volume that most of the traditional types of energy generators (e.g. hydroelectric, thermoelectric, nuclear) have to face.

In this way, the desideratum is to achieve a market model where retailers have the most accurate possible prognosis and

the capability of influencing or guiding customers' consumption. To this end, there have been a number of initiatives, grouped under the general banner of *Demand-Side Management* (DSM), which has been promoted by Utility Companies (UCs) as an alternative to building new power plants. Its (consciously infeasible) objective consists of smoothing demand so that ideally it is a flat constant energy consumption 24 hours a day, 365 days a year. This profile embodies the ideal circumstances for energy producers and Transmission System Operators (TSOs) since the former may employ cheap and stable production methods, harmless to the environment, and the latter get to control a transmission grid with constant load [2].

Now, the easiest way to accomplish this goal is by setting the price of the energy depending on the actual demand load. Thus, the higher the demand, the more expensive the price, and vice versa. Based upon these premises, many utility companies already present a basic form of DSM by offering a cheaper night tariff. In this way, our aim in this work is to improve and extend this simple market model to permit UCs to express more complex aims and, thus, increase their influence on customers. For instance, in order to lighten the peak-time load, the supplier can offer a discount for consuming a small amount of energy at 8 am (peak-time) and a larger amount at midnight (off peak). This incentivises the customer to reschedule some tasks to midnight (e.g. the dishwasher or the washing machine). If many clients accept this compromise offer, the UC will have achieved a double goal. It will have a more accurate prognosis for 8 am and midnight and it will also have shifted some of the peak-time consumption to off peak.

In e-commerce terms, this process can be seen as a *reverse combinatorial auction*. It is "reverse" because the customers pick one of the available companies and tariffs to supply their future consumption. And it is "combinatorial" because bidding for a bundle of items is typically valued differently from bidding separately for each of the constituent items (e.g. the *combination* of consuming at 8am and midnight is more appreciated, and thus rewarded, than, for instance, the combination of consuming at 10am and 11am).

While combinatorial auctions provide very efficient allocations that can maximise the revenue for the auctioneer, their main drawback is the complexity of the *clearing* process in which buyers and sellers are matched and the quantities of items traded between them are determined. Specifically, clear-

ing combinatorial auctions is non-deterministic polynomial-time hard (NP-hard) [3], [4], which can make combinatorial auctions intractable in practice [5]. Thus, a significant amount of work has been concentrated on developing strategies to overcome this shortcoming. Unfortunately, most work in this area deals with clearing combinatorial auctions with *atomic propositions* [6], [7], [8], [9]. Thus, bids are either accepted or rejected in their entirety, which may limit the profit for the auctioneer. A more efficient solution is to allow bidding with demand/supply functions [10], [5], in which bidders submit a function¹ to calculate the cost of the units to be bought or sold. This allows the customer to accept parts of different bids and constitutes a powerful way of expressing complex pricing policies. In our case, production costs can be easily reflected in the supply function and if bids are accepted partially, there may be more than one winner for the same auction and item. Therefore, customers may accept different parts of bids from different bidders so they can get energy simultaneously from several suppliers. This kind of function was chosen since it can readily approximate any curve.

In our problem domain, such a setting enables customers to accept different parts of bids from different bidders so they can get energy simultaneously from several suppliers. Since the transmission and distribution grids are shared and the path followed by the electricity cannot be tracked down, it is impossible to determine the producer of the energy being consumed. Therefore, the hypothesis of customers being simultaneously supplied by several UCs does not pose any technical problems.

Against this background, this paper advances the state of the art in two main ways. First, we present, for the first time, an energy retail market designed as a system of reverse combinatorial auctions with supply function bidding. This novel market allows customers to increase their profit and provides UCs with a mechanism to influence customers' behaviour. Second, we develop new optimal clearing algorithms tailored to electricity supply functions that perform better than the existing more general clearing algorithms.

The remainder of the paper is organised as follows. Section II details the overall market design. Section III presents the single-item and the multi-item clearing algorithms, analyses their complexity and optimality, examines the results of comparing the multi-item algorithm with the only other optimal multi-item one defined in the literature and presents a strategy for keeping the multi-item within tractable ranges. Section IV discusses related work. Finally, Section V concludes and outlines the avenues of future work.

II. ELECTRICITY RETAIL MARKETS

This section discusses the nature of current electricity retail markets and outlines the design of our solution.

¹In case of a forward auction, bidders submit a demand function. In case of a reverse auction, bidders submit a supply function.

A. Requirements

Currently, most customers only partially enjoy the benefits of a deregulated market. They typically sign mid-term contracts with a single supplier and the tariffs do not reflect the pressure of competition. Moreover, whereas classical capitalist pricing policies encourage demand by applying discounts on quantity (the more you buy, the cheaper the unit price becomes), actual electricity contracts often include a threshold above which the consumption becomes more expensive. To move to a more dynamic environment where the benefits of competition can be more fully realised, we put forward the following requirements for our market design. The arrangement of customers' electricity supply from multiple UCs should be achieved by having contracts that specify the provision of an amount of energy for a certain period of time (say one hour). These contracts should not necessarily be exclusive and, thus, customers may have agreements with different companies for the same hour if this is the best thing to do. Finally, we assume customers auction, on a daily basis, their next 24 hours consumption divided into 24 items (representing one hour each). They subsequently receive bids from the UCs and make their decision for the next 24 hours, which is a trade-off between the very static situations of today and the possibility of auctioning on a per minute basis for the coming minute. Finally, in order to make the whole process automatic, both UCs and customers are represented by software agents. Thus, hereafter when we write "UCs" or "customer", we mean "UCs' software agent" and "customer's software agent".

B. Market Design

The requirements detailed above can be best met by structuring the market as a reverse auction. Furthermore, we assume customers don't issue any bids, but simply choose among those offered by the UCs. An exchange (in which multiple buyers and sellers submit their bids and offers to an independent auctioneer that decides the winners [11]) was rejected because it scales poorly. In practice, the number of customers may be up to tens of thousands, each of which is selling 24 items, and with combinatorial bidding, clearing such an exchange becomes intractable very fast. Unlike exchanges, reverse auctions have the advantage that they may be performed in parallel. This means the complexity can be divided between the number of customers because instead of one big auction, many *smaller* ones are carried out at the same time. For these reasons, we have designed our system as a series of simultaneous reverse auctions despite the risk of *overbooking*.

This is a problem that cannot be underestimated (see the consequences in, for instance, the 2000-2001 California electricity crisis [12], [13], [14]). It refers to the impossibility of UCs of controlling exactly how many customers will accept their bids (and, therefore, they cannot predict their total demand). It can be approximated or foreseen with the help of statistical means, but depending on the quality of these measures, the threat of a blackout will always be present. Such a dilemma exists in every market in which a single producer cannot supply all the demand itself. Therefore, a

producer that is too successful could get to a point above which, paradoxically, it is not economical to sell (because it cannot produce so much and, therefore, must buy the difference somewhere else).

Think for instance of a UC starting a price war². In case it suddenly decreases the rates and more customers than expected opt to change to this UC, it will have problems to meet the new forthcoming demand. Less dramatically, a normal UC can also submit an offer that by sheer chance is accepted by more clients than it expected. In other words, when UCs prepare their tariffs, they must also take into account the number of clients they may supply, in addition to the factors mentioned before. Therefore, again, an extremely successful tariff, ironically, may not be highly desirable.

From another point of view, carrying out separate auctions in parallel, where duplicated resources are available from each auction, implies the risk of selling the same resource simultaneously in separate auctions at the same time. This is a problem in our overbooking approach and stems from the fact, as stated before, that resource suppliers have uncertainty about the outcome of each single reverse auction. In other words, suppliers place bids *in all auctions simultaneously* as if the auctions were separate entities and their results do not influence each other in any way (which is clearly not the case).

In this way, in order to prevent this blackout (worst-case) scenario, we see only two feasible solutions left, detailed as follows.

1) *Airline-like Model*

In order to minimise the number of unoccupied seats (and thus, the revenue of the flight) within a passenger aeroplane, airlines are allowed to accept a number of bookings that is higher than the actual number of seats available on the plane [16]. Then, passengers denied boarding due to overbooking are compensated. Similarly, we may allow UCs to offer to sell more electricity than they can actually supply. In case of overbooking, we envisage two approaches: Either UCs withdraw (part of) their bids and compensate affected customers³ or UCs buy the energy they need in a second market. Nevertheless, both pose new game-theoretic issues (such as the influence of second markets and bid-withdrawal possibilities in strategies) that require further research. Indeed, the second market, together with a non-realistic ceiling price, were the main factors which brought about the aforementioned 2000-2001 Californian electricity crisis. Therefore, we recommend the solution detailed next.

2) *Regions Model*

A different approach to this problem is addressed in

²For an interesting simulation and discussion of price wars with software agents see [15].

³This is not an uncommon practice. The (USA) Federal Communications Commission (FCC, <http://www.fcc.gov/>) [17], [18] allows bid withdrawal at a penalty consisting of the maximum of zero or the difference between the value of the withdrawn bid on a license and the highest bid after the withdrawal on that license. See [19], [20] for more information on this subject.

[21], [22] and [23]. Here, Haque, Jennings and Moreau present a distributed resource allocation protocol that allocates end-to-end network bandwidth by means of using market-based agents that are deployed in a communications network. The agents compete to buy and sell bandwidth resources from auction servers that use combinatorial reverse auctions. Their approach consists of dividing their whole communications network into distinct local regions, in which resources are auctioned exclusively (i.e. only within the region they belong to). This setting provides much needed benefits since communication messages do not have to be broadcast to all auction servers from where resources are needed, but only to the desired ones.

We can borrow this idea and apply this strategy to our framework in two ways. In the first possibility, customers are concentrated randomly or geographically in *heterogeneous* groups, so suppliers can estimate the maximum demand they will offer to each of the groups. In the second modality, customers are concentrated in *homogeneous* groups, according to the priority of their supply. For instance, one group might include all public-service clients (e.g. hospitals, police stations, etc.), the next most important group might be composed by strategic private companies (e.g. petrol stations or big industries), and so on. In both alternatives, auctions are carried out in steps, in a random order, for the heterogeneous groups case, or in a priority order for the homogeneous groups case and, therefore, the risk of overbooking is minimised.

Coming back to our market design, as combinatorial bidding is permitted, UCs submit their *special* discounts together with the usual hour tariffs. In this case, having 24 hours (or items) means that there may be up to 2^{24} different combinations of discounts. This is obviously a worst-case scenario because, in practice, our experience in the domain indicates that UCs are highly unlikely to issue a different discount for each possible combination. Moreover, we decided that the auctions should be sealed (to reveal the least possible information) and single-round (to minimise communication and other delays). The auctions also need to be both multi-item and multi-unit. As each item is the supply of electricity in one hour, there are 24 items to allocate in an auction. In addition, each bidder may not allocate the whole consumption within an hour, but rather just a portion of it (i.e. some units). Figure 1 illustrates this new market setting.

Another important component to set is the price paid by the winner. We do not want to have a first price auction because it offers incentives for strategic behaviour (i.e. the participants act according to beliefs formed about others' values and types, which does not assure them of maximising their payoff). To circumvent this, we choose a uniform second price for combinatorial auctions (Vickrey-Clarke-Groves) since this has the dominant strategy of bidders bidding their true valuations of the goods [24], [25]. The price paid by the winner is not directly specified in the bid because bidders submit a supply

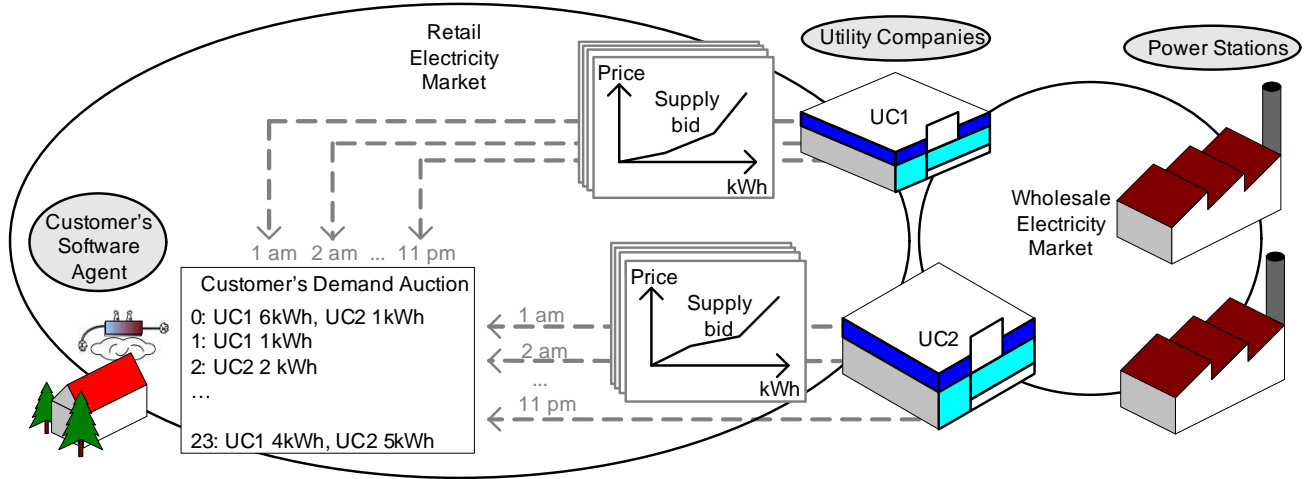


Fig. 1. Novel market format: simultaneous reverse auctions of demand.

function. Thus, the customer must calculate the energy it wants to consume within a time slot (i.e. the units of that item to be auctioned) and then decide the cheapest combination with the supply functions submitted (i.e. number of units to be allocated with each bidder). Therefore, the bids are accepted partially. To this end, we use the compact notation introduced in [26], where bidders submit for a certain item a piece-wise linear supply function P composed of n linear segments. Each segment l , $1 \leq l \leq n$, is described by a starting quantity s_l , an ending quantity e_l , a unit price π_l , and a fixed price C_l . Thus, if a customer wants to buy q units of that item from the supplier, it will pay $P_l = \pi_l \cdot q + C_l$ if $s_l \leq q \leq e_l$. Additionally, bidders submit a correlation function, ω , which shows the reward or penalty of buying a number of items together (it is this that makes the bidding truly combinatorial). For instance, $\omega_1(A, B) = 0.95$ would mean that if buying x units of items A and y units of B (i.e. consuming x Kw at time A and y Kw at time B), the price paid will have a 5% discount. Thus, if the unit price of item A is p_a and the unit price of item B is p_b , the total price would be $0.95 \cdot ((x \cdot p_a) + (y \cdot p_b))$.

Currently, there is only one optimal algorithm to solve this problem. Specifically, the one presented by Dang and Jennings in [26] (described in more detail in section IV). However, we believe this is inapplicable in our scenario because it scales poorly (as we show in section III-C). Therefore, with the market described above in place, the next step is to design a clearing algorithm that solves the winner determination problem more efficiently and allows it to be actually applied in realistic contexts.

III. OPTIMAL CLEARING ALGORITHMS

This section details the optimal single-item (sPJ) and the optimal multi-item clearing algorithm (mPJ) that we have developed for the electricity retail market described in section II. Furthermore, we analyse their complexity, prove their optimality, and analyse strategies to keep them tractable. First

of all, let us introduce some basic definitions that will be used thereafter:

Definition 1: A **single allocation** is a set \langle time-slot t , supplier s , amount q , price p \rangle meaning that s wants to pay p to buy q units of energy to be consumed at time t .

Definition 2: An **allocation** is a list containing a number (between one and the number of suppliers) of single allocations that detail the supply of electricity to be provided to the customer at a given time-slot.

Definition 3: A **more profitable allocation** from two alternatives is the one that for a given total demand q , has the lower total price p .

Definition 4: An **optimal allocation** is one in which the demand constraint is satisfied and there is no more profitable allocation.

Definition 5: An **optimal day allocation** is a set of 24 optimal allocations, each of which corresponds to a different item (i.e. there is an optimal allocation for each hour).

The clearing algorithms we present in this section are related in that the multi-item one is a consecutive and iterative processing of the single-item one (i.e. the result of the multi-item algorithm is obtained by executing the single-item one with different values). Specifically, clearing a single-item case implies finding the optimal allocation for that item, so this enterprise deals only with the supply functions submitted to one item. The multi-item case has a broader remit (an optimal day allocation) and, thus, it also takes into account the relationships between the different items of the optimal allocations (i.e. the correlation functions). Let us first start with the explanation of the single-item case.

A. The Optimal Single-Item Clearing Algorithm

Clearing a single-item algorithm with piece-wise supply function bids involves determining the amount to be allocated to each submitted bid function. In essence, in each loop the algorithm selects one segment of each supply function (the one corresponding to the already allocated demand) and allocates

k units to the segment with the best price (i.e. the lowest price for k units after applying any relevant discount on the amount). The loop is repeated until the demand is satisfied. Note that the value of k is dynamically assigned in each loop to guarantee the optimality of the algorithm. Specifically, it always has the ending quantity value (e_i) of the shortest segment being evaluated at that moment.

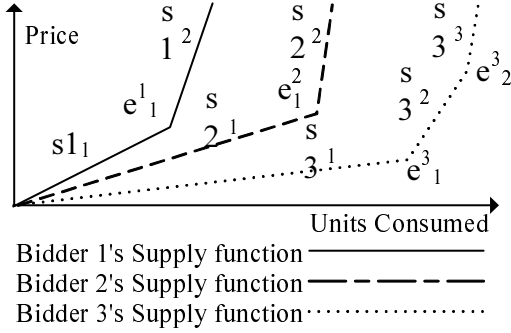


Fig. 2. Linear piece-wise supply functions submitted to a single item.

Let us now illustrate this procedure with the example of Figure 2. Assume there are three potential buyers 1, 2, and 3 that submit their supply functions s_1 , s_2 , and s_3 for a certain item (i.e. the consumption in one hour). In the first loop, the algorithm processes the segments s_{11} , s_{21} , and s_{31} . Since the shortest of the three is segment s_{11} (i.e. $e_1^1 < e_1^2 < e_1^3$), $k = e_1^1$ and the algorithm compares $s_{11}(e_1^1)$, $s_{21}(e_1^1)$, and $s_{31}(e_1^1)$. Suppose the price of $s_{31}(e_1^1)$ is less than the price of $s_{11}(e_1^1)$ and $s_{21}(e_1^1)$; then, it selects s_{31} to supply these first e_1^1 units. In the second loop, the algorithm processes the segments s_{11} , s_{21} , and s_{31} (but starting from e_1^1) and gives k the value of $e_1^3 - e_1^1$ because it is less than e_1^1 and e_1^2 . Then, it compares $s_{11}(e_1^3 - e_1^1)$, $s_{21}(e_1^3 - e_1^1)$, and $s_{31}(e_1^1)$, and so on. The algorithm continues until the amount of allocated units is equal to the demand.

As we can see, the algorithm evaluates one function per bidder in each step so it has a complexity $O(m)$ per loop, where m is the number of bidders. As the loop is repeated k times, where k is the number of segments of the function with the highest number of them, the overall complexity is $O(km)$. A safe way to reach an optimal allocation is to select for each unit the segment that offers the best price (i.e. $k = 1$). However, it is not necessary to repeat the process for each single unit since price and discount are constant in each segment. So, as long as the segments evaluated in each loop are the same (unit price and fixed price remain unchanged), the winner will also be the same. Thus, in each loop it is only necessary to compare the price of allocating the lowest ending quantity of the segments being processed, repeating this process until the demand is satisfied. Therefore, sPJ (detailed in Figure 3) always finds the most profitable optimal allocation.

Input: m supply functions f and *demand*.

- Pre-loop: initialise needed variables: *allocated* to keep the total allocated demand, the list *allocation* showing the allocated demand per bidder, and the temporal storage variable k .

- Loop: in each loop, until the demand is satisfied, select the segment with the lowest gradient and allocate the minimum ending quantity units.

While (*allocated* < *demand*) do

/* Select the the segment with the lowest gradient */

$k = \text{select the minimum ending quantity of the current segments}$

if (*demand* - *allocated* < k) then

winner = select the minimum $f_m(k)$

/* Allocate the minimum ending quantity units */

allocated += k

allocation[*winner*] += k

else

winner = select the minimum $f_m(\text{demand})$

/* Allocate the minimum ending quantity units */

allocated += *demand* - *allocated*

allocation[*winner*] += *demand* - *allocated*

Output: *allocation*, the variable detailing the amount allocated to each bidder.

Fig. 3. The sPJ clearing algorithm.

B. The Optimal Multi-Item Clearing Algorithm

This algorithm, detailed in Figure 4, is more complex since it cannot simply be generalised from the single-item one. If there were no correlations, it would be sufficient to run the sPJ case once for each item. However, the existence of correlations poses the problem of the inconsistent application of discounts. First, if a supplier bids for two items and offers a reduction if both bids get accepted, no reduction should be applied if only one of them succeeds. Second, functions become different after applying a discount. For example, assume P_l is a piece-wise supply function for the item l and it is included in the correlation $\omega(l, \dots) = x$. Then, P'_l is the new supply function with the value $P'_l = xP_l$. Thus, the optimal allocation of a set of functions in which P_l is included may not be the same as the one in which everything else is the same but with P'_l instead of P_l .

In this way, mPJ must process all possible combinations of discounted and non-discounted functions and check that discounts are applied consistently. To this end, we use a brute-force strategy [27], [28], [29] for identifying all the possibilities. Here, all possible bids from each bidder are combined with all possible bids from the rest of the bidders. However, it is not necessary to evaluate all the combinations since some of them are repeated. For instance, Table I shows an auction with two suppliers (1 and 2) and two items (a and b). In this case, there is one possible correlation for each bidder, $\omega^1(a, b) = x$ and $\omega^2(a, b) = y$. Thus, clearing the multi-item case implies evaluating the combinations where supplier 1 and 2 bid normally for item a (so the single-item clearing algorithm is run with supply functions $P_a^1 - P_a^2$);

Input: j supply functions s_j , j correlation functions ω_j and demand q_i for each item i .

- Pre-loop: Initialise variable *day-set* to keep the optimal allocation for each item, *item-set* to keep a group of supply functions to be evaluated by the single-item clearing algorithm, *all – item – sets* to keep already processed sets of supply functions (*item – sets*), and a boolean variable *ok*.

- Loop: For each item calculate the optimal allocation of a possible set of supply functions and then check whether the selected discounts are applicable.

```

Do
  for each item  $i$ 
    for each supplier  $s_j$ 
      /* Generate a new combination of segments*/
      add next  $s_j^i$  to item-set
      if item-set not in all-item-set then
        /* Calculate the optimal allocation of a single item*/
        optimal-allocation = single_item_algorithm(item-set)
        store item-set in all-item-sets
        add optimal-allocation to day-set
      /* Check whether the allocation is consistent*/
      ok = check_constraints (day-set,  $\omega_j$ ).
      If ok then compare day-set with best so far
      and keep the best of both
    until all the combinations are explored
  
```

Output: *day-set*, a set of i optimal allocations (one for each item) with the lowest total price.

Fig. 4. The mPJ clearing algorithm.

supplier 1 bids for item a and b with discount and supplier 2 bids normally for item b (so the single-item algorithm clears item a with supply function xP_a^1 and item b with xP_b^1 and P_b^2), and so on.

TABLE I
SINGLE-ITEM EVALUATIONS WITH TWO ITEMS AND TWO BIDDERS,
REPEATED IN BOLD.

	P_a^2	yP_a^2 yP_b^2	P_b^2	-
P_a^1	$P_a^1 - P_a^2$	$P_a^1 - yP_a^2$ yP_b^2	P_a^1 P_b^2	P_a^1
xP_a^1 xP_b^1	$xP_a^1 - P_a^2$ xP_b^1	$xP_a^1 - yP_a^2$ $xP_b^1 - yP_b^2$	xP_a^1 $xP_b^1 - P_b^2$	xP_a^1 xP_b^1
P_b^1	P_a^2 P_b^1	yP_a^2 $P_b^1 - yP_b^2$	$P_b^1 - P_b^2$	P_b^1
-	P_a^2	yP_a^2 yP_b^2	P_b^2	-

This brute-force strategy evaluates all possible bid combinations (without repeating some of them) and, therefore, it always finds the most profitable optimal day allocation. However, it also scales poorly. First, the number of possible combinations depends on the number of items. In our case, with 24 items, there are 2^{24} different combinations. Second, it also rises exponentially as the number of bidders grows: with n items, and two bidders, 2^{2n} ; with three bidders 2^{3n} , and

so on. In the extreme situation with two bidders submitting a different supply function for each one of the 24 items and 2^{24} correlations, there are $2 \cdot 2^{48}$ possible combinations. This is $n \cdot (2^n)^m$ where n is the number of items and m the number of bidders. For instance, in the example of Table 1, there are $2 \cdot (2^2)^2 = 32$ possible combinations, but half of the combinations do not need to be re-calculated (in bold format in Table 1). Thus, if bidders bid for all items and submit all possible correlations, the number of times that the multi-item algorithm clears the single-item one is $n \cdot (2^n - 2^{n-1})^m = n \cdot (2^{n-1})^m$. Therefore, the complexity is $O(kmn \cdot 2^{(n-1) \cdot m})$, where n is the number of items, m the number of suppliers and k the number of segments of the supply function with more segments. Note, however, that this is a pathological worst-case scenario, which is highly unlikely to happen in practice. Furthermore, as we discuss afterwards, it can be mitigated against by constraining the agent's bidding behaviours.

C. Evaluation

In this section we present the results of comparing the performance of our mPJ algorithm with the only other optimal algorithm for this class of problem. Specifically, our benchmark is the multi-item algorithm mDJ presented by Dang and Jennings in [5] (described in more detail in section IV).

The comparison shown in Figure 5 details how the complexity (defined in terms of X, the number of bids) scales when the number of items n increases for a constant number of bidders m . As can be seen, mDJ soon becomes intractable (i.e. prohibitively high complexity), and mPJ scales better.

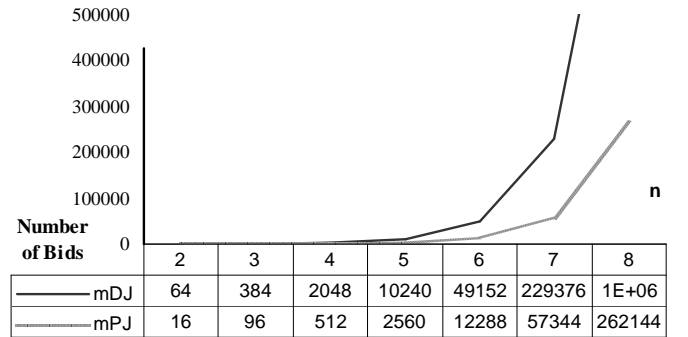


Fig. 5. Complexity evolution with n increasing and m steady ($m = 2$)

Figure 6 tests how the algorithms react to the increment of m (bidders) when n (items) remains steady. Again, mDJ becomes intractable as soon as it did in Fig. 5, whereas mPJ presents a significantly better performance profile. The main reason for this behaviour is the sensitivity of mDJ to the increment of both n and m (while mPJ is only sensitive to the increment of n , as seen in Fig. 5). For mDJ, a larger number of items and clients means a larger number of single-allocations to form the set from which the allocations will be formed. Whereas for mPJ, more clients means more correlations to clear, but half of which need not be processed since they are repeated.

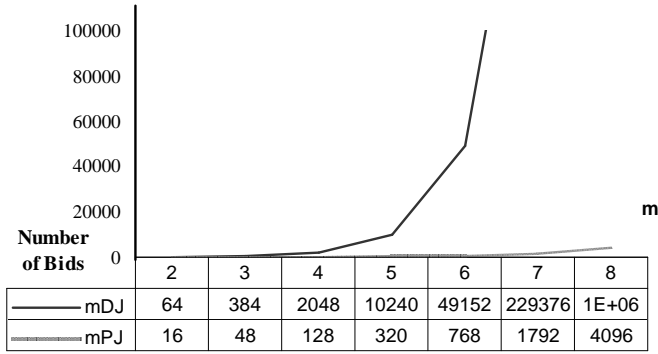


Fig. 6. Complexity evolution with m increasing and n steady ($n = 2$).

Similarly, Figure 7 illustrates the behaviour of the algorithms when both n (items) and m (bidders) increase. Again, mDJ performs worse than the others. Its $n = 2$ series is almost equivalent to the $n = 3$ of our multi-item algorithm.

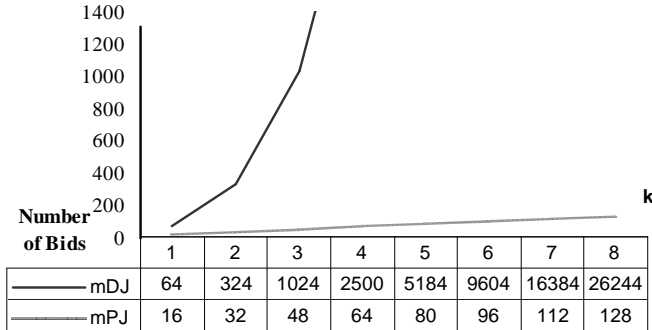


Fig. 7. Complexity evolution with n and m steady and k increasing ($n, m = 2$)

Finally, Figure 8 depicts the dependence of each algorithm on k , the number of units allocated in each iteration of the single-item algorithm. In this dimension, mPJ again performs well. For mDJ, increasing k implies increasing the number of single allocations that may be combined with each other (therefore the algorithm grows exponentially with k as the base). In contrast, for mPJ increasing k just implies that the single-item algorithm is going to process more steps (therefore the algorithm grows linearly with k as the factor).

D. Constrained Bidding

In order to prevent such pathological scenarios from occurring, it is possible to constrain the choice of possible discounts so bidders can submit only a certain number of correlations. This type of restriction has already been successfully applied to atomic propositions bidding [30], where when limiting the allowable combinations to tree structures or sequential combinations, the NP-hard winner determination problem can be solved in polynomial time [9]. In a similar vein, mPJ can also take advantage of such an approach. Specifically, we can constrain the number of correlations to a value c . Thus, bidder 1 can issue, for instance, the following: $\omega_1^1(n_1, n_2 \dots n_i)$, $\omega_2^1(n_1, n_2 \dots n_i) \dots \omega_c^1(n_1, n_2 \dots n_i)$ where i is the number

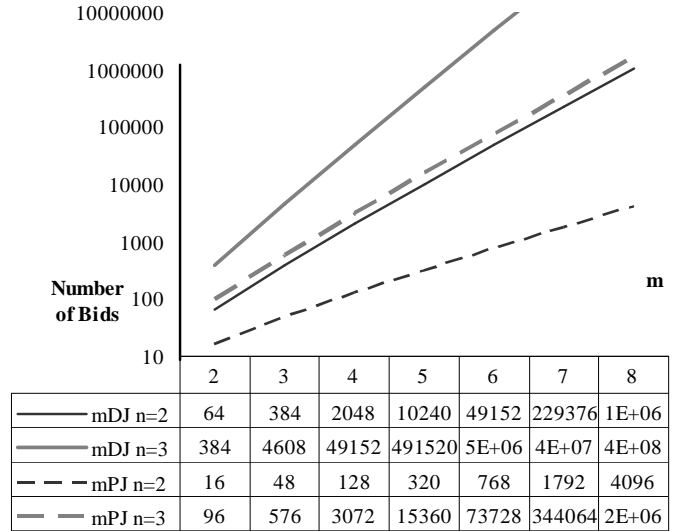


Fig. 8. Complexity evolution with n and m increasing.

of items included in each discount (for the sake of simplicity, let us suppose it is a fixed number less than n , the number of items, but big enough to allow the bidder to be sufficiently flexible in its offering).

In this way, the single-item algorithm sPJ will be executed, with m bidders, $i \cdot c^m$ times (again, supposing that i is fixed) and the complexity of the mPJ algorithm will drop to $O(ki \cdot c^m)$. Unfortunately, in this case, the mPJ algorithm cannot skip evaluating half of the combinations (as in section 4). With this constrained discount choice, the reduction depends much more on the specific discount combinations chosen. For instance, if the combinations include many items (i.e. i is bigger), the single-item algorithm will be executed more often than if the combinations only include two items each. In short, there is no way to accurately determine it *a priori*. Similarly, restricting the available amounts assigned to the discount increases the number of repeated combinations. Thus, if a supplier offers the same reduction for accepting two different items (e.g. $\omega(a, b) = \omega(c, d)$), the number of repeated combinations would increase further and the complexity would continue decreasing.

For the comparison shown next, we have set the maximum number of bids to be issued as half of the maximum possible ($c = 2^{n-1}$) and the maximum number of items included in a correlation as the number of items ($i = n$). For instance, in figure 9, we present the reaction of the three algorithms to the increment of the number of items n for a constant number of bidders m .

Here, the constrained variant presents the best profile for our purposes. This would have been even clearer if we had not set the value of c and i depending on the number of items n (as detailed above). With a fixed c and i , the constrained variant would have presented a flat line, whereas mDJ and mPJ would had grown exponentially because, in contrast to mDJ and mPJ, the constrained variant does not depend directly on the number of items being auctioned.

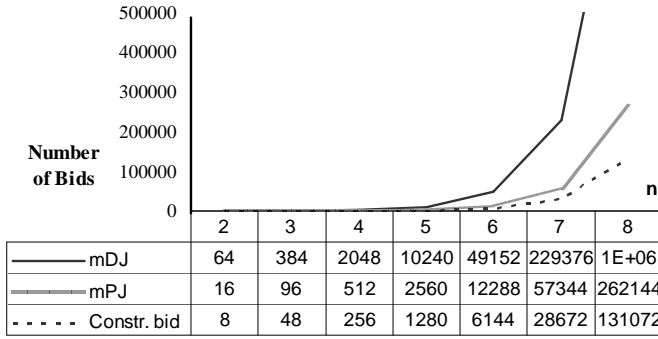


Fig. 9. Complexity evolution with n increasing and m steady ($m = 2$).

Figure 10 tests how the algorithms react to the increment of m (bidders) when n (items) remains steady. The results of the constrained variant are better than those from mPJ. Again, mDJ becomes intractable as soon as it did in Fig. 9.

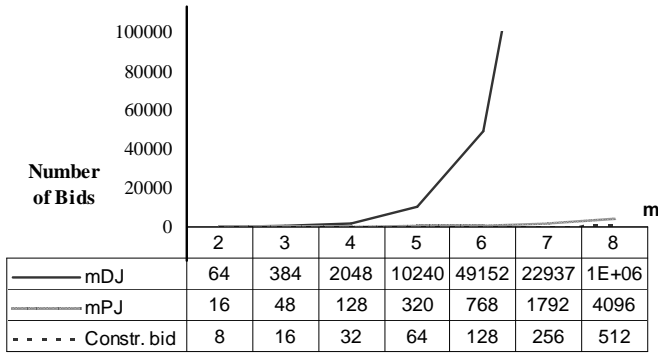


Fig. 10. Complexity evolution with m increasing and n steady ($n = 2$).

Figure 11 presents the reaction of the algorithms to the increment of both n (items) and m (bidders). As we would expect, the best results are again achieved by the constrained variant.

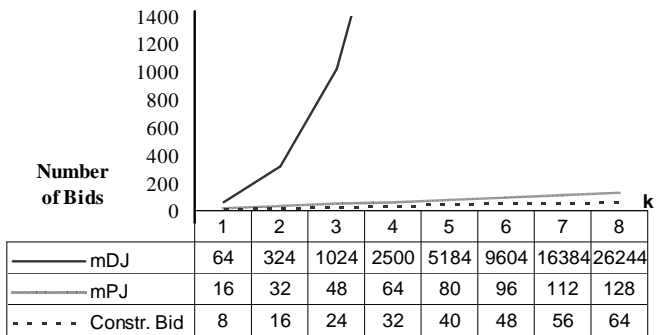


Fig. 11. Complexity evolution with n and m steady and k increasing ($n, m = 2$).

Finally, Figure 12 shows the dependence of each algorithm on k , the number of units allocated in each iteration of the single-item algorithm. Again, mDJ is the worst of the three and the constrained variant by far the best. This behaviour is due to the direct dependency of mDJ on k .

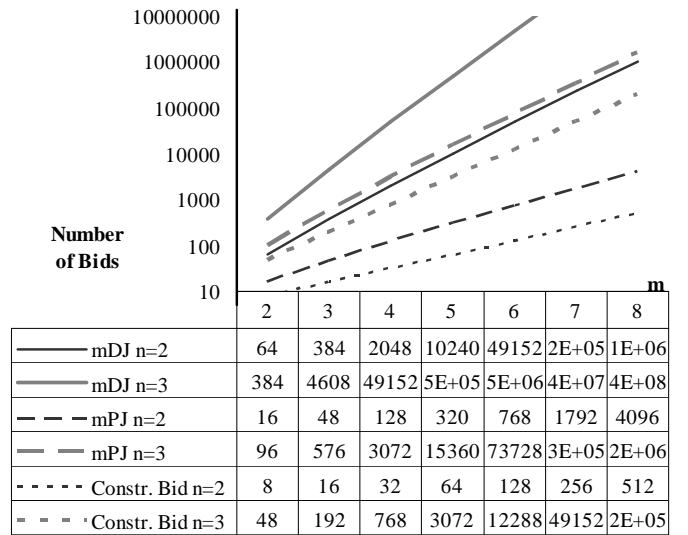


Fig. 12. Complexity evolution with n and m increasing.

Note that the complexity of the constrained variant can be further reduced depending on the values of i and c . With the values we assigned to i and c for these comparisons, it is only m times less complex than mPJ (since $c = 2^{n-1}$, $i = n$ and $O(ki \cdot c^m)$, then the complexity after substitution of c and i is $O(kn \cdot 2^{(n-1)m})$). The genuine advantage of the constrained variant can be found when there are higher values of n and m . Thus, based on our beliefs about the likely operation of the retail energy market some “typical” values might be to have 24 items (e.g. 24 hours) and around 20 bidders (e.g. 20 UCs trying to sell their energy). Therefore, if we set $k = 1$ and restrict the number of possible correlations to 10, each one with 5 items (which experience indicates will provide UCs with enough *persuasive* power), the results are clear: mDJ presents a complexity of $1,498E+147$, our mPJ $1,429E+141$ and the constrained variant $5E+20$. In our opinion, this means the constrained variant is sufficiently close to the optimal to be useful, but is still sufficiently tractable to be practicable.

IV. RELATED WORK

There has been comparatively little previous work in combinatorial energy markets, but there is a much larger literature on clearing algorithms for combinatorial auctions. However, these two strands of work have not been brought together before.

In more detail, there has been a surge of research in this area lately, most of which has concentrated on properties of the underlying multi-agent system and is thus slightly out of focus to our needs. Moreover, only the work of Palensky et al. [31], [2], [32], [33] is explicitly directed to remove the shortcomings posed by deregulation (or seen in a positive way, to exploit the new opportunities it brings). Specially interesting is his PhD thesis, where he addresses a genetic algorithm as an optimal method to find the proper demand schedule. Nevertheless, the problem target is different to ours since he only intends to offer a solution to the consumer side.

In this way, by ignoring the supplier side, aspects such as the electrical market are disregarded and, thus, so is the possibility of influencing customers that we offer to the suppliers.

Worthy of mention is also the work of Ygge et al. [34], [35], [36], [37], [38], [39], [40], [41], [42], which is seminal in the area of agents and energy management. Specifically, they combine power load management with market-oriented programming. However, there are two reasons that make their approach inapplicable to our domain. First, market-oriented programming is based on the Theory of Competitive Equilibrium (as opposed to Game Theory). This model studies equilibrium conditions in which participants deal only with parameters such as the price, but not with possible actions of the others. That is, agents are cooperative and not competitive, a feature that we require to design electricity markets in which bidder agents compete to be chosen to supply an auctioneer’s energy demand. Second, they introduce a hierarchical structure of *HomeBots*, intelligent agents that represent every load in the system and buy the energy in a system of forward non-combinatorial auctions. With only one energy supplier, this approach places all the initiative on the *HomeBots* so the UCs cannot express their preferences for having more or less demand at a certain time. We address this shortcoming by allowing combinatorial bidding.

There has been an enormous amount of research in combinatorial auctions [11] recently, but most of this has focused on atomic propositions that may limit the choice (and hence the profit) to the auctioneer, as already stated. Addressing this limitation, several authors have developed algorithms that deal with demand/supply bidding. For instance, Sandholm and Suri [10] considered this possibility but only for the single-item case (which is not sufficient since our problems demand a multi-item auction). Moreover, working with a single item implies that the auction is not truly combinatorial [43], [5]. In this way, there exist only two algorithms dealing with this auction setting.

In the first one, Dang and Jennings [43] develop the first single-item and multi-item algorithm for multi-unit combinatorial reverse auctions with demand/supply functions. They sacrifice optimality at the cost of running in polynomial time, which is a common practice in AI (see a compilation of different polynomial algorithms for diverse problems in [44] or examples in [45], [46], [47]). The algorithms are not guaranteed to find the optimal solution, but do produce solutions that are shown to be within a finite bound of the optimal, which sometimes is an acceptable trade-off. Still, we would like to have an optimal algorithm for our problem.

In the second one, again Dang and Jennings [5] present another pair of single and multi-item algorithms (referred to here as *sDJ* and *mDJ*) for the same environment, but this time they are optimal. The strategy they use consists of defining a dominant set containing an increasingly sorted group of single allocations, so they search within this dominant set for the combinations that form the most profitable day allocation. Given m bidders, n items and k , the upper bound on the number of segments of the dominant set, the complexity in a

worst case scenario is $O(n \cdot (k + 1)^n)$ in the single-item case and $O(mn \cdot (k + 1)^{mn})$ in the multi-item.

In comparison to this work, the single-item algorithm presented here, *sPJ*, is less general than *sDJ* (because it only clears continuous piece-wise supply functions), but both our algorithms present significantly lower computational complexity even in a worst-case scenario ($O(km)$ in the single-item case and $O(kmn \cdot 2^{(n-1) \cdot m})$ in the multi-item). That is even if, for instance $k = 1$, our *mPJ* algorithm is still 2^{n-1} times less complex than *mDJ*. Moreover, in the constrained bidding variant of *mPJ*, this difference is even higher. If we again set $k = 1$ and restrict the number of possible correlations to 10, each one with 5 items (which experience indicates will provide UCs with enough *persuasive* power), the results are clear: *mDJ* presents a complexity of $1,498E+147$, our *mPJ* $1,429E+141$ and the constrained variant $5E+20$. In our opinion, this means the constrained variant is sufficiently close to the optimal to be useful, but is still sufficiently tractable to be practicable.

V. CONCLUSIONS AND FUTURE WORK

The deregulation of the electricity industry offers new opportunities for providers and consumers. In this environment, customers can choose their suppliers to get cheaper energy and suppliers can compete to increase the number of their customers and, subsequently, their profits. To make this happen in practice, however, efficient electricity markets need to be developed. To this end, traditionally, energy management techniques have presented the two different sides with their own purposes and measures. On the one hand, suppliers and retailers aim to smooth the overall energy consumption to avoid sudden peak loads. On the other hand, customers intend to reduce their energy bills without giving up freedom (meaning they can use energy at any time). Our system addresses both needs. It helps to reduce peak loads and to distribute them amongst less-loaded time slots. Specifically, by including off-peak hours in the discounts, UCs reward customers that consume electricity off-peak. Thus, they have an additional tool for energy management besides setting off-peak prices lower than peak ones. Moreover, the use of combinatorial auctions helps to produce efficient allocations of goods because combinatorial bidding allows the expression of more complex synergies between auctioned items [48]. Together with the use of supply functions and non-atomic propositions, consumers are able to accept energy from diverse UCs simultaneously, which, in turn, helps them to maximise their benefits.

Against this background, this paper presents, for the first time, an electricity retail market as a system of simultaneous reverse combinatorial auctions with supply-function bidding. Furthermore, we have developed the novel single and multi-item clearing algorithms *sPJ* and *mPJ* that are optimal, as well as a strategy to keep the multi-item algorithm within tractable ranges for the real-world problem we face. Still, it is not enough to find the optimal allocation of demand of a single customer if this demand is not consumed according to the profile. Hence, we need a method that coordinates the customers’

devices to find the best consumption schedule with capacity to analyse the submitted bids and find the best consumption profile according to these bids. Thus, we have two interrelated combinatorial problems: finding the best schedule among the possible consumption alternatives of devices and finding the best demand allocation among the submitted supply bids and bundle discounts. The solution to this second problem has been presented in [49] and [50].

Future work will focus on evaluating the whole electricity market system and on reducing the complexity of the multi-item clearing algorithm with additional restrictions on combinatorial bidding. Moreover, a number of game-theoretic aspects should be reconsidered. These include the suitability of the Vickrey-Clarke-Groves mechanism (given the criticism about its applicability to real-world environments by Ausubel and Milgrom in [51]) and the possibility of replacing it (for instance with the Ausubel auction [51]). In this vein, the work of Kahneman and Tversky [52], [53], [54] suggests that human's motivations and decision making do not always follow mathematical (say game theoretical) models and, therefore, it could be worth studying this issue and its influence on both the pricing strategies of suppliers and their real behaviour in our system of simultaneous reverse combinatorial auctions. In the same way, adopting the regions model as tool for reducing the possibility of overbooking poses new questions regarding UC's pricing strategies because both with the heterogeneous and with the homogeneous groups the auction would be multi-round. Finally, further work should also focus on developing a simulator to test both the electricity market and the electricity consuming environment with real values. This would lead us to an accurate assessment of the whole system under concrete and realistic conditions.

ACKNOWLEDGMENT

An early version of this paper appeared at the IEEE/WIC/ACM International Conference on Intelligent Agent Technology, Compiègne, France, Sep. 2005, pp. 626-632. The authors would like to thank Viet Dang, Raj Dash, Alex Rogers and Nadim Haque from the IAM Group at the University of Southampton for their support and assistance.

REFERENCES

- [1] D. Mell, "Demand forecasting," Report, Cap Gemini Ernst & Young. Available <http://www.cgey.com/energy>, 2001.
- [2] P. Palensky, "Distributed reactive energy management," Ph.D. dissertation, Institute of Computer Technology, Vienna University of Technology, 2001.
- [3] Y. Fujishima, K. Leyton-Brown, and Y. Shoham, "Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches," in *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, 1999, pp. 548-553.
- [4] T. Sandholm, "Making markets and democracy work: A story of incentives and computing," in *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, Acapulco Mexico, 2003, pp. 1649-1671.
- [5] V. Dang and N. Jennings, "Optimal clearing algorithms for multi-unit single item and multi-unit combinatorial auctions with demand/supply function bidding," in *Proceedings of the International Workshop on Internet Computing and E-Commerce*, Pittsburgh PA, 2003, pp. 25-30.
- [6] N. Nisan, "Bidding and allocation in combinatorial auctions," in *ACM Conference on Electronic Commerce*, 2000, pp. 1-12.

- [7] K. Leyton-Brown, Y. Shoham, and M. Tennenholtz, "An algorithm for multi-unit combinatorial auctions," in *Proceedings of the 17th National Conference on Artificial Intelligence*, Minneapolis MN, 2000, pp. 13-20.
- [8] M. Tennenholtz, "Some tractable combinatorial auctions," in *Proceedings of the 17th National Conference on Artificial Intelligence*, 2000, pp. 98-103.
- [9] T. Sandholm, "Algorithm for optimal winner determination in combinatorial auctions," *Artificial Intelligence*, vol. 135, no. 1-2, pp. 1-54, 2002.
- [10] T. Sandholm and S. Suri, "Market clearability," in *Proceedings of the IJCAI'01*, Seattle WA, 2001, pp. 1145-1151.
- [11] T. Sandholm, "Algorithm for optimal winner determination in combinatorial auctions," *Artificial Intelligence*, vol. 135, pp. 1-54, 2002.
- [12] V. Authors, "Manifesto on the california electricity crisis." Available at http://www.haas.berkeley.edu/news/california_electricity_crisis.html, 2001.
- [13] S. Oren and P. Spiller, "High electricity prices in the West: What can be done about it?" *Public Utilities Fortnightly*, vol. 138, no. 20, 2000.
- [14] R. Haas and H. Auer, "Die kalifornische stromkrise," *Energiewirtschaftliche Tagesfragen*, vol. 51, no. 5, pp. 280 - 287, 2001.
- [15] J. Kephart, J. Hanson, and J. Sairamesh, "Price-war dynamics in a free-market economy of software agents," in *6th Conference on Artificial Life*, University of California at Los Angeles, California, 1998, pp. 53-62.
- [16] "Regulation (EEC) No 295/91 of 4 February," Official Journal of the European Communities No L 36, 8.2.1991, p. 5, 1991.
- [17] M. Bykowsky, R. Cull, and J. Ledyard, "Mutually destructive bidding: The FCC auction design problem, Social Science Working Paper 916. Pasadena: California Institute of Technology, Division of Humanities and Social Sciences," 1995.
- [18] J. McMillan, "Selling spectrum rights," *Journal of Economic Perspectives*, vol. 8, pp. 145-162, 1994.
- [19] D. Porter, "The effect of bid withdrawal in a multiobject auction," *The Review of Economic Design*, vol. 4, pp. 73-97, 1999.
- [20] D. Parkes, "iBundle: An efficient ascending price bundle auction," in *Proceedings of the 1st ACM Conf. on Electronic Commerce*, 1999, pp. 148-157.
- [21] N. Haque, N. Jennings, and L. Moreau, "Resource allocation in communication networks using market-based agents," in *Proceedings of the 24th SGAI International Conference on Innovative Techniques and Applications of Artificial Intelligence*, M. A. Bramer, Ed., 2004, pp. 187-200.
- [22] —, "Resource allocation in communication networks using market-based agents," *International Journal of Knowledge-Based Systems*, vol. 18, no. 4-5, pp. 163-170, 2005.
- [23] —, "Resource allocation in communication networks using market-based agents," in *Proceedings of the 1st International Workshop on Incentive-Based Computing*, 2005, pp. 30-39.
- [24] T. Groves, "Incentives in teams," *Econometrica*, vol. 41, pp. 617-31, 1973.
- [25] R. K. Dash, D. C. Parkes, and N. R. Jennings, "Computational mechanism design: A call to arms," *IEEE Intelligent Systems*, vol. 18(6), pp. 20-47, 2003.
- [26] V. D. Dang and N. R. Jennings, "Optimal clearing algorithms for multi-unit single item and multi-unit combinatorial auctions with demand/supply function bidding," in *Proceedings of ICEC'03*, Pittsburgh PA, 2003, pp. 25-30.
- [27] S. Russell and P. Norvig, *Artificial Intelligence: A Modern approach (2nd Edition)*. Prentice Hall, 2002.
- [28] A. Levitin, *Introduction to the Design and Analysis of Algorithms (2nd Edition)*. Addison Wesley, 2006.
- [29] M. Weiss, *Data Structures and Problem Solving Using Java (3rd Edition)*. Addison Wesley, 2006.
- [30] A. Babanov, J. Collins, and M. Gini, "Scheduling tasks with precedence constraints to solicit desirable bid combinations," in *Proceedings of the 2nd International Conference on Autonomous Agents and Multi-Agent Systems*, 2003, pp. 345-352.
- [31] P. Palensky, D. Dietrich, R. Posta, and H. Reiter, "Demand side management by using lonworks," in *Proceedings of the 2nd IEEE international Workshop on Factory Communication System*, 1997, pp. 341 - 347.
- [32] P. Palensky, P. Roessler, and D. Dietrich, "Heim- und gebaudeautomatisierung zur effizienzsteigerung in gebauden," *Elektrotechnik und Informationstechnik*, vol. 4, pp. 104-110, 2003.

- [33] M. Stadler, P. Palensky, B. Lorenz, M. Weihs, and C. Roesener, "Integral resource optimization networks and their techno-economic constraints," *International Journal on Distributed Energy Systems*, vol. 11, no. 4, pp. 199–319, 2005.
- [34] F. Ygge and H. Akkermans, "Power load management as a computational market," in *Proceedings of the 2nd International Conference on Multi-Agent Systems*, 1996, pp. 393–400.
- [35] H. Akkermans, F. Ygge, and R. Gustavsson, "Homebots: Intelligent decentralized services for energy management," in *Knowledge Management: Organization, Competence and Methodology*, J. Schreinemakers, Ed. Ergon Verlag, 1996.
- [36] H. Akkermans and F. Ygge, "Smart software as customer assistant in large-scale distributed load management," in *Proceedings of the International Conference and Exhibition on Distribution Automation and Demand Side Management*, 1997.
- [37] F. Ygge and H. Akkermans, "Duality in multi-commodity market computations," in *Proceedings of the 3rd Australian Workshop on Distributed Artificial Intelligence*, C. Zhang and D. Lukose, Eds., 1997, pp. 65–78.
- [38] F. Ygge, "Market-oriented programming and its application to power load management," Ph.D. dissertation, Department of Computer Science, Lund University, 1998.
- [39] —, "A multi-commodity market approach to power load management, Blekinge University of Technology Research Report 1997:15," 1998.
- [40] F. Ygge, H. Akkermans, A. Andersson, M. Krejic, and E. Bortjes, "The homebots system and field tests: A multi-commodity market for predictive load management," in *Proceedings of the 4th International Conference on Practical Application of Intelligent Agents and Multi-Agent Technology*, London UK, 1999, pp. 363–382.
- [41] R. Gustavsson, "Agents with power," *Communications of the ACM*, vol. 42, no. 3, pp. 41–47, 1999.
- [42] F. Ygge and H. Akkermans, "Resource-oriented multi-commodity market algorithms," *Autonomous Agents and Multi-Agent Systems Journal*, vol. 3, no. 1, pp. 53–72, 2000.
- [43] V. Dang and N. Jennings, "Polynomial algorithms for clearing multi-unit single item and multi-unit combinatorial reverse auctions," in *Proceedings of the 15th European Conference on Artificial Intelligence*, Lyon France, 2002, pp. 23–27.
- [44] D. Hochbaum, Ed., *Approximation Algorithms for NP-Hard Problems*. PWS Publishing Company, 1997.
- [45] G. Nemhauser, L. Wolsey, and M. Fisher, "An analysis of the approximations for maximizing submodular set functions," *Mathematical Programming*, vol. 14, pp. 265–294, 1978.
- [46] R. Hassin, "Approximation schemes for the restricted shortest path problem," *Mathematics of Operations Research*, vol. 17, no. 1, pp. 36–42, 1992.
- [47] R. Mettu, "Approximation algorithms for np-hard clustering problems," Ph.D. dissertation, Department of Computer Science, University of Texas at Austin, 2002.
- [48] Y. Fujishima, K. Leyton-Brown, and Y. Shoham, "Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches," in *Proceedings of the IJCAI'99*, Stockholm Sweden, 1999, pp. 548–553.
- [49] Y. Peña, N. Jennings, and G. Neumann, "An optimal distributed constraint optimisation algorithm for efficient energy management," in *IEEE Proceedings of the 2005 International Conference on Intelligent Agents, Web Technology and Internet Commerce*, 2005.
- [50] Y. Peña, "Optimal allocation and scheduling of demand in deregulated energy markets," Ph.D. dissertation, Technical University of Vienna, 2006.
- [51] L. Ausubel and P. Milgrom, "The lovely but lonely vickrey auction," in *Combinatorial Auctions*, P. Cramton, Y. Shoham, and R. Steinberg, Eds. MIT Press, 2006.
- [52] D. Kahneman and A. Tversky, "On the psychology of prediction," *Psychological Review*, no. 80, pp. 237–251, 1973.
- [53] —, "Prospect theory: An analysis of decisions under risk," *Econometrica*, no. 47, pp. 313–327, 1979.
- [54] —, "Choices, values and frames," *American Psychologist*, no. 39, pp. 341–350, 1984.