Stable coalition formation among energy consumers in the smart grid

Meritxell Vinyals,
Filippo Bistaffa,
Alessandro Farinelli
Department of Computer Science
University of Verona, Italy
{meritxell.vinalssalgado,alessandro.farinelli}@univr.it

Alex Rogers
Agents, Interaction and Complexity Group
School of Electronics and Computer Science
University of Southampton,UK
acr@ecs.soton.ac.uk

ABSTRACT

The vision of the Smart Grid includes demand-side peak shaving strategies, such as real-time pricing or profile's based tariffs, to encourage consumption such that the peaks on demand are flattened. Up to date, most works along this line focused on optimising via scheduling of home appliances or micro-storage the individual user consumption. Alternatively, in this paper we propose to exploit the consumer social side by organising them into coalitions of energy consumers with complementary needs. To this end, we propose the concept of virtual energy consumer (VEC) to capture the notion of a number of energy consumers, coming together to buy electricity, as an aggregate. To create such VEC's we consider that each consumer looks for potential partners for its coalitions through its contacts in a social network. In more detail, we propose a network-restricted coalitional game, where: (i) each feasible VEC is evaluated with a metric that estimates the expected joint payment of the coalition of consumers within the electricity markets; (ii) the set of most efficient VEC's are identified (by solving the corresponding Coalition Generation Problem); and (iii) the joint payment of each VEC is divided among its members in such a way that any consumer can not be better off by deviating and forming a new VEC (i.e., we compute a core-stable payoff distribution if this exists or alternatively detect core emptiness). Moreover, we evaluate our approach on consumption data for a set of households located in UK. Our analysis provides interesting insights into the relationship between structure and stability of VEC's and prices within the electricity markets.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Multiagent systems

General Terms

Economics

Keywords

Coalition formation, virtual energy consumer, stability, smart grid

1. INTRODUCTION

Since energy cannot be stored efficiently on a large scale, the electricity grid must perfectly balance the demand of all customers at any instant with supply. In all current electricity grids this balance is achieved by varying the supply-side to continuously match demand. The amount of demand required on a continuous basis is usually carried by the baseload stations owing to low cost generation, efficiency and safety. However, these stations are slow to fire up and cool down, so they are not able to match the peakload periods that exceed this baseload that, in contrast, require expensive, carbon-intensive, peaking plants generators. Although only running when there is high demand, these peaking plants generators are responsible for most part of consumers electricity bill.

Along this line, the vision of the Smart Grid includes demand-side peak-shaving strategies such as real-time pricing or profile's based tariffs to encourage consumption such that the peaks on demand are *flattened* [1]. A flattened demand results in a more efficient grid not only with lower carbon emissions but also with lower prices for consumers. Hence, some works [13, 15] focused on techniques that flatten individual consumer demand by automatically controlling home domestic or micro-storage devices. Unluckily, since each consumer independently optimises its own consumption, the effectiveness of this approach has a clear limit on the consumer's restrictions and comfort (e.g. it will be unavoidable to get a consumption peak in the non-working hours of consumers). Even more worrying, demand-side management technologies based on individual price reaction have shown tendency to reduce the natural diversity of consumers' peak demands leading to the shifting of current peaks to new specific periods [14].

Against this background, in this paper we aim at improving the grid efficiency from a social perspective by promoting the formation of coalitions among energy consumers with near-complementary consumption restrictions. Then, a coalition of consumers can act in the market as a single virtual energy consumer (VEC) with a flattened demand for which it gets much better prices. As we analyse in this work, several challenges arise in the formation and management of these energy coalitions. On the one hand, from the Grid's perspective, it is important to ensure that negotiations among individuals energy consumers converge in such a way that the most efficient VECs are formed (i.e., achieving the maximum social welfare). On the other hand, consumers are rational utility maximizers, and convergence is only achieved when all the members of the VEC agree on their share of the coalition's payment. Moreover, consumers may not want to join coalitions with unknown consumers for which they do not have any source of trust regarding their reported profiles or their capacity to meet their payment obligations.

In this paper we address the above-identified requirements by proposing a game-theoretical model for VEC formation that finds the most efficient VEC's to form and splits each total VEC payment among its members. Our solution is based on modeling the energy consumer coalition formation problem as a coalitional game [10], where: (i) the set of coalitions with maximum collective value, that is, an optimal coalition structure, has to be identified; and (ii) each coalition's value has to be distributed among its members in such a way that coalition members have no incentive to break away from the identified optimal structure. Moreover, we restrict the coalitions membership using consumer's acquaintances in a social network to provide some form of trust among coalition members. Thus, in more detail, this paper makes the following contributions:

- We formally define the concept of VEC along with a metric that, given an estimation of the aggregated coalition consumption, computes the total payment optimizing the buying strategy within the electricity markets.
- We propose an algorithm that allows consumers in a social network to distributedly generate and evaluate the network-feasible energy coalitions.
- We use linear programming approaches to identify the most efficient VEC's and to subsequently allocate core-stable payments to individual consumers (if the core is not empty). When such core-stable allocation exists the total payment of each VEC is split among its members such that the formed VEC are stable (i.e. consumers can not be better off leaving their current VEC and forming a new one).
- We evaluated our model on a real dataset based on the electric consumption of households in the UK. The results are analyzed in terms of the structure and stability of the formed coalitions as well as the gain obtained by consumers as a result of engaging in the coalition formation process. Results show that stability of most efficient VEC's is significantly affected by the density of the social network and the market prices conditions. Moreover, the structure of the VEC's formed is highly depending on the ratio between different market prices: close prices between electricity markets does not incentivize the formation of coalitions whereas as the distance between prices is increased larger coalitions appear in the market (until after some limit parameter value the grand coalition always emerge).

This paper is structured as follows. In Section 2, we review the literature and in Section 3, we describe our model for coalition formation among energy consumers as a coalitional game. Section 4 presents our empirical evaluation. Finally, Section 5 concludes and outlines some paths for future research.

2. BACKGROUND

2.1 Today's electricity markets

In most European countries, the current operation of the exchange electricity market is composed of multiples markets available for trading electricity, each with different operation and purpose [3, 2] . In particular, most countries define and distinguishes between, at least, two different kinds of markets: the *spot electricity markets* and the *forward electricity markets*. The main goal of spot markets lies in the facilitation of the trading of short-term energy delivery. Thus, in a spot market, energy is traded independently for each time slot and hence, each time slot may have a different price (e.g. the day-ahead market is a spot market where hourly blocks of electricity are negotiated for the next day). In contrast, forward markets are the venue where forward electricity contracts for long periods (e.g. month, quarter or year) with delivery and withdrawal obligation are negotiated. Thus, the contract in a forward market specifies a single quantity that will be delivered at constant rate for the contract period and a single price. Finally, any real-time excess or shortfall in supply and demand (with respect to the contracted volume) is reconciled in the balancing market. The balancing market is cleared just before the actual power is delivered by producers.

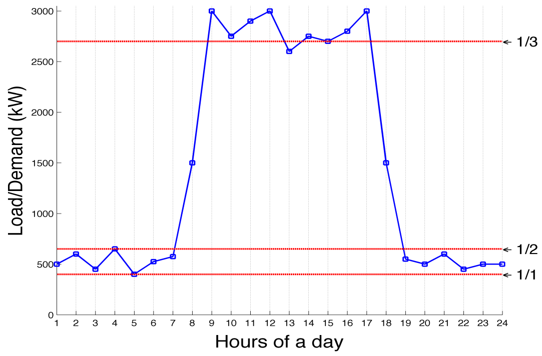

Figure 1: A sample of a hourly load energy profile and the different quantities to buy in the forward market given different forward over day-ahead market prices ratios.

To date, although in many countries (e.g. the US, UK and most of European countries) the electricity market is deregulated, market operation and conditions restricts explicitly or implicitly the participation to wholesale companies who subsequently sell the electricity to final consumers in form of standard products (i.e. fixed contracts and tariffs). However, with the advent of the smart grid this is expected to change. The vision of the smart grid involves significant changes in the way energy is bought and sold including a two-way communication between the grid and consumers that allow a more active role of the latter. For example, as part of the smart grid community, electricity consumers have already access to smart meters that allow them to monitor its (load) energy profile in an hour-day basis. Figure 1 shows an example of an energy profile as a graph that plots the variation in the electrical load (measured in kW) versus time (measured in hours). Formally, we define the energy profile of a consumer a_i as a vector $E_i = \{e_i^1, \ldots, e_i^N\}$ where e_i^t is the amount of energy consumed at time slot t.

2.2 Coalitional games

A coalitional game is traditionally defined as follows. Let $A = \{a_1, \ldots, a_n\}$ be a set of agents. A subset $S \subseteq A$ is termed a coalition. However, depending on the domain not all coalitions may be feasible. In particular, here we are interested on restricting coalitions by a *graph* G: (i) each node of the graph represents an agent; and (ii) a coalition S is allowed to form iff every two agents in S are connected by some path in the subgraph induced by S. We denote the set of graph feasible coalitions as $F(G)$. Then, a coalitional game CG is completely defined by its *characteristic function* $v : F(G) \rightarrow \Re$, which assigns a real value representing (transferable) utility to every feasible coalition [10]. Agents in a coalition are then permitted to freely distribute coalitional utility among themselves. Given a game CG, a *coalition structure* $CS = \{S_1, \ldots, S_k\}$ is an exhaustive disjoint partition of the space of agents into feasible coalitions. We refer to the coalition composed of all agents as grand coalition and to the coalition composed by a single individual agent as singleton coalition. We overload notation by denoting by $v(CS)$ the (intuitive) worth of a coalition structure: $v(CS) = \sum_{S \in CS} v(S)$.

Then, the coalition formation process can generally be considered to include three differentiated activities: *Coalitional Value Calculation*, *Coalition Structure Generation* and *Payoff Distribution*. First, on *coalitional value calculation*, agents enumerate and evaluate all possible feasible coalitions that can be formed. Next, given the values of feasible coalitions, the key challenge addressed in *coalition structure generation* is to identify the coalition structure CS^* that maximizes *social welfare* - i.e. the coalition structure with maximal value. Finally, *Payoff Distribution* determines the utility that each agent in a coalition should obtain as a result of the actions taken by the coalition as a whole. A vector $\rho = \{\rho_1, \ldots, \rho_n\}$ assigning some payoff to each agent $a_i \in A$ is called an *allocation*. We denote $\sum_{i \in S} \rho_i$ by $\rho(S)$. An allocation ρ is an *imputation* for a given CS, if it is efficient ($\rho(S) = v(S)$ for all $S \in CS$), and individually rational (that is, $\rho_i \geq v(\{i\})$ for all a_i). Note that if ρ is an imputation for CS, then $\rho(A) = v(CS)$. A game outcome is a (CS, ρ) pair, assigning agents to coalitions and allocating payoffs to agents efficiently. However, in a selfish environment, agents are only concerned with maximizing their own payoffs. Thus, with the presence of selfish agents we need to determine *stable* allocations. Here, stability requires agents to have no incentive to deviate from the coalitions to which they belong. Cooperative game theory provides several stability concepts, here we focus on the core which is arguably the most well-studied. The core is composed of all coalition structure-imputation tuples (CS, ρ) such that no feasible coalition has any incentive to deviate. Formally:

$$Core(CG) = \{(CS, \rho) : \rho(A) = v(CS) \;\&\; \rho(S) \geq v(S) \; \forall S \in F(G)\}$$

The core is a strong solution concept, as it is empty in a plethora of games. Moreover, notice that *only optimal coalition structures might admit an element in the core*. Intuitively, if the current structure is suboptimal then a subset of agents can be made strictly better off by moving to an optimal coalition structure. Hence, for any core-pair allocation CS is an optimal coalition structure, CS^*, and the allocation ρ is efficient with respect of the value of the optimal coalition structure ($\rho(A) = v(CS^*)$). In this work, we are interested on the question of how to compute a core member: this includes to solve the CSG problem to get the optimal coalition structure CS^* and compute the stable payoff allocation over CS^* or detect the emptiness of the core.

3. THE MODEL

In this section, we model the problem of demand-side coalition formation among energy consumers as a coalitional game. Let $A = \{a_1, \ldots, a_n\}$ be the set of agents, each one representing an energy consumer with its associated energy profile E_i. Agents can form energy coalitions $S \subseteq A$, where an energy coalition S stands for the set of consumers S acting as a VEC in the market along with their joint consumption.

The first issue to be addressed is which coalitions consumers are going to consider and which is the metric they are going to use to evaluate them. In particular, we propose that consumers use social networking tools as free available technologies to support the discover, formation and restriction of their energy coalitions. From a game point of view, the metric simply represents the characteristic function of the coalitional game whereas the social network constrains the set of feasible coalitions (as defined in Section 2.2).

Now, the process of forming VECs at a technical level require of mechanisms and strategies that allow energy consumers to come with the most efficient coalitions (i.e. if some consumers only consume at specific times of the day, they will want to choose those partners they can complement better at those times) and to an economical agreement (i.e. how they share the payments generated by the total consumption the VEC). From a game theoretic point of view, this involves to solve the CSG problem and find a core-stable payoff distribution (as defined in Section 2.2). In our model, we use a linear programming approach to solve both the CSG problem and to find the core-stable payments that divides the payments of optimal coalitions among its members (or alternatively, detects the inexistence of such payments).

In the next sections, we specify in more details how we solve the three main activities that underline the coalition formation process for this particular domain.

3.1 Coalitional Value Calculation

In this section we formalize coalitional value calculation, namely the generation and evaluation of feasible coalitions, for the VEC formation. First, in Section 3.1.1, we define a metric to evaluate coalitions that computes the total payment that coalition of consumers S will need to carry out to get their aggregated demand. Next, we address the problem of enumerating and evaluating all energy coalitions in the social network in Section 3.2.

3.1.1 Coalition value metric

To determine the value of a coalition, we define a metric that computes, for each coalition, the total payment estimated for the coalition. In more detail, this metric optimises the buying strategy across energy markets taken, to meet the expected VEC aggregate consumption.

The first issue that arises in this context is how a coalition of agents predict their aggregated consumption over time. Although predicting the joint demand of a coalition is a topic of relevance itself, in this work we do not tackle this problem. Thus, for the sake of clarity, we simply take the joint average energy profile of the coalition as a predictor of the daily coalition consumption. Similarly to singleton consumer coalitions, the (expected) demand of any coalition of consumers S is represented by their joint energy profile $E_S = \{e_S^1, \ldots, e_S^N\}$ where $e_S^t = \sum_{i \in S} e_i^t$.

Now, following the operation of the current grid, we consider that consumers buy directly their electricity in two different markets: the day-ahead market (that forms part of the spot electricity markets) and the forward electricity market. Let p_F be the unit energy price in the forward market and p_D the average unit energy price among daily hours in the day-ahead market (prices are negative values to denote the direction of payment). The value of the expected payment for the coalition S is given by:

$$v(S) = \sum_{t=1}^{N} q_D^t(S) \cdot p_D + N \cdot q_F(S) \cdot p_F \tag{1}$$

where $q_F(S)$ stands for the time unit amount of energy to buy in the forward market and $q_D^t(S)$ for the amount of energy to buy in the day-ahead market at time slot t. Notice that whereas the amount of energy bought in the day-ahead market can vary at each time slot, the quantity to buy in the forward market has to be continuous for all the period (i.e. the same quantity for all N time slots). Also, to guarantee that the demand for each time slot is covered, these quantities must satisfy the following constraints:

$$q_D^t(S) + q_F(S) \geq e_S^t \quad \forall t = 1 \ldots N \tag{2}$$

Hence, to compute the value of a coalition in this domain agents face the decision problem of determining the quantities to buy in the forward and the day-ahead market such that Equation 1 is maximised (i.e. the payment regarding their joint consumption is minimized) whereas satisfying constraints in Equation 2 that guarantee that these quantities meet the coalition energy needs.

Next, we describe a procedure (outlined in Algorithm 1) that allows agents to optimally solve the above-defined optimization problem. This procedure takes as input the coalition energy profile, E_S, and the ratio between prices among the two available markets, $\frac{p_F}{p_D}$. Intuitively, in order for agents in a coalition to be advantageous to buy a certain continuous quantity q in the forward market

Algorithm 1 computeCoalitionBuyingStrategy(E_S, $\frac{p_F}{p_D}$)

1: Sort E_S in descending order;
2: $q_F(S) \leftarrow E_S[round(\frac{p_F}{p_D} \cdot N + 0.5)]$; /*The amount of energy that is covered at least the $\frac{p_F}{p_D}$ of the time interval is the continuous quantity that coalition should buy in the forward market*/
3: **for** $t = 1 \ldots N$ **do**
4: $\quad q_D^t(S) = max(e_S^t - q_F(S), 0)$;
5: **end for**
6: **return** $\langle q_F(S), q_D^1(S), \ldots, q_D^N(S) \rangle$

(instead of buying the individual q^i quantities necessary for each time slot), this continuous amount should be used at least $\frac{p_F}{p_D}$ of the time interval. Figure 1 shows how different ratios between the forward/day-ahead market lead to different buying strategies within electricity markets for the same coalition profile. Horizontal lines stand for different quantities to be bought in the forward market given the market ratios as labeled at the end of the line. To compute such quantity given the discretization of the profile in N elements, we first order the coalition profile values in descending order (line 1). Then the quantity to buy in the forward market is simply the energy consumption value at position $\frac{p_F}{p_D} \cdot N$ (assuming the profile array values starting at 1 to get the value on such position we need to round adding one half, line 2). Observe that in the particular case when there is no economical incentive to buy in the forward market ($p_F = p_D$), the forward quantity represents the safer base load corresponding to the minimum quantity that is expected to be continuously consumed along hours. Thus, in the load profile of Figure 1 if the ratio between markets is 1 over 1 the amount to buy in the forward market is exactly the minimum among the hours consumption (corresponding to time slot 3). Finally, given the forward energy quantity to be bought in the forward market, the day-ahead quantity for a time slot t is simply computed as the amount of demanded energy that exceeds the forward quantity (line 3). Notice that as the incentive to buy in the forward market increases, agents increases the quantity bought in the forward market, in detriment of this bought in the day-ahead, by buying continuous amount even when they are not expect to used it all hours of the day. Thus, in Figure 1 if the ratio between markets is 1 over 2 the amount to buy in the forward market is exactly the 12th measure in magnitude among all the hours consumption (corresponding to time slot 4) although for half of the day the consumption is expected to be less than this amount.

It is worth noting that the computation of coalition's values has the primary objective of maximizing agent's profits, since $p_F < p_D$ it indirectly encourages the formation of flattened profiles, capturing the synergies that exist between consumers to improve the efficiency of the grid.

3.2 Network-based coalitions

As discussed above, in our model we consider that each consumer looks for potential partners for its coalitions through its contacts in a social network. In this way, coalition membership is restricted to coalitions composed of *friends of friends*, being always somebody in the coalition responsible for the introduction of a new member. From the game perspective this restriction implies that feasible coalitions are restricted by a graph. More formally, a coalition among agents is feasible if its members form a vertex-connected induced subgraph. Here, we observe that this problem can be cast to the problem of connected induced subgraph enumeration, for which several algorithms have been proposed on the literature [4, 7]. Exploiting similar procedure we propose a distributed algorithm that allows agents organised into a network to list all network-restricted coalitions and compute their values.

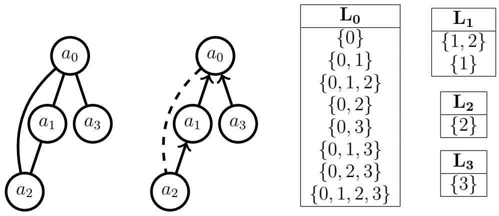

(a) G (b) PT (c) G-restricted coalitions

Figure 2: Example of (a) a network with a cycle (G); (b) a pseudotree PT of G and (c) the set of G-restricted coalitions partitioned in leading coalitions per agents.

Let G be a connected (undirected) graph with vertex set $A(G) = \{a_1, a_2, \ldots, a_n\}$ and let E be the set of edges among agents. An example of a 4-agents network with a cycle that defines a graph G is given in Figure 2 (a). For an agent $a_i \in A$ let $N(\cdot)$ be a function that returns the set of neighbours of a_i in G, that is, $N(i) = \{j \in A(G), j \neq i, (i,j) \in E(G)\}$. Thus, in Figure 2 (a), the neighbours of agent a_1 is the set $N(1) = \{0, 2\}$. Tables in Figure 2 (b) list the set of feasible coalitions restricted by the graph in Figure 2 (a). Thus, a_3 can form an energy coalition with agents a_0 and a_1 ($S = \{013\}$) but not a coalition with a_1 without a_0 ($S = \{13\} \notin F(G)$).

Instead of using a linear ordering as in [7], we propose to use the partial ordering that defines a pseudotree arrangement of the agent's graph [5]. A pseudotree PT of G is a rooted tree with agents $A(G)$ as nodes and the property that any two agents that share an edge in G are on the same branch in PT. Pseudotrees are a common structure used in search and inference procedures given their ability to exploit independencies between nodes in a graph, allowing parallel processing of independent branches. Figure 2(b) shows a pseudotree, rooted at agent a_0, of the cyclic graph G in Figure 2(a). A PT has two kinds of edges: tree-edges (bold lines) that link parent with children (e.g. a_2 is child of a_1); and pseudoedges (dashed lines) that link pseudoparents with pseudochildren (e.g. a_2 is pseudochild of a_0). Let's denote $A(PT)$ the set of agent' nodes in PT and PT_i the subtree of PT rooted at a_i. Thus, in Figure 2(b), PT_1 is a tree rooted at a_1 composed of agents a_1, a_2. Finally, given an agent $a_i \in A(PT)$ we will denote as Ch_i its children, An_i its ancestors (the set composed of its parent and its pseudoparents), and D_i its descendants (the set composed of its children and pseudochildren) in PT. Then, in Figure 2(b), $Ch_2 = D_2 = \emptyset$ and $An_2 = \{a_0, a_1\}$. Then, given a game on a graph $CG = \langle A(G), v, F(G) \rangle$ and a pseudotree PT over G, the partial ordering that PT defines among agents allows us to partition the set of feasible coalitions into $|A|$ disjoint sets $\{\mathbf{L}_i | a_i \in A\}$, one per agent. The set of (leading) coalitions $\mathbf{L}_i$ contains all the feasible coalitions in which a_i is the leader (precedence position in the ordering), that is all coalitions that include agent a_i but no agent up a_i in PT, $\mathbf{L}_i = \{S \in F(G) | i \in S, \forall j \in S : \ level(i) \leq level(j)\}$). Figure 2(c) shows the different sets of agents leading coalitions for the pseudotree in Figure 2(b).

Next, we describe the main steps of a distributed procedure that allow agents to compute the set of leading coalitions on a graph. Thus, at the end of this process, each agent will know its set of leading coalitions $\mathbf{L_i}$.

Each agent a_i uses the PT partial ordering to order its descendants $D_i = \{d^1, \ldots, d^m\}$ from higher to lower. For example, in Figure 2(b), a_0 can order its descendants as a_1,a_2,a_3 (as far as a_1 is placed before a_2 the order is valid). Then, a_i will proceed to generate its set of leading coalitions $\mathbf{L}_i$ into two steps: first, generating a set of basic coalitions, and second, generating a set of composed coalitions, that result from combination of basic ones.

Step 1. (Basic coalitions). In this step, each agent a_i will generate the set of *basic* coalitions. For each descendant $d^{j=1\ldots m}$, a_i generates all coalitions S such that $\{d^j\} \subseteq S \subseteq A(PT_i) \setminus \{E\}$ where $E = \{d_k | k < j\}$ stands for all descendants placed before d^j in the ordering. To generate these sets of coalitions we use a distributed version of a recursive connected induced subgraph enumeration algorithm proposed in [7] (for further details see [7], section 5). For example, in Figure 2(b), agent a_0 will generate the set of coalitions that include: (i) a_1 and other agents reachable from a_1 in PT_1 ($\{1\}$,$\{12\}$); (ii) a_2 and other agents reachable from a_2 in PT_1 excluding a_1 ($\{2\}$); and (iii) a_3 and other agents reachable from a_3 in PT_1 excluding a_1 and a_2 ($\{3\}$). Each agent a_i records for each coalition S a set of frontiers nodes, F, these are nodes that are reachable from S but not included in S. For example, in Figure 2(b) the set of frontiers for $\{1\}$ is $\{2\}$ since a_2 is reachable from a_1 but not included. Finally, a_i stores each generated coalition S, adding a_i, $(S \cup \{i\})$, as well as a_i's singleton coalition ($\{i\}$), as part of its leading coalitions $\mathbf{L}_i$.

Step 2. (Composed coalitions). In this step, each agent a_i will generate the set of *composed* coalitions. For each descendant $d^{j=1\ldots m}$ a_i will combine all coalitions reachable from d^j with all compatible coalitions reachable from d^{j+1}, from d^{j+2}, $\ldots$, and until d^m, storing at each step the new coalitions as reachable from d^j. Thus, in Figure 2(b), a_0 will combine coalitions reached from a_1 with coalitions from a_2 (storing any new coalition as reachable from a_1) and coalitions reachable from a_1 with coalitions from a_3. Two coalitions, S reachable from d^j, and S' reachable from $d^{k>j}$, are compatible if S' does not contain any agent in S or in its frontiers ($S' \cap S \cap F_S \neq \emptyset$). Thus, in Figure 2(b), a_1 will not combine $\{1\}$ from a_1 and $\{2\}$ from a_2 because a_2 is a frontier for $\{1\}$. A composed coalition is generated as $\{S \cup S'\}$ with frontier agents $\{F_S \cup F_{S'}\}$. Thus, in Figure 2(b), a_0 the result of combining $\{2\}$ from a_2 with $\{3\}$ from a_3 is a composed coalition $\{2, 3\}$ reachable from a_2 with $F = \{2\} \cup \emptyset$.

3.3 Coalition Structure Generation

To solve the CSG problem, which is known to be NP-Hard, we use an integer programming (IP) approach (see [11], pages 38-39). Compared with with other state-of-the-art CSG algorithms [8, 12], this approach has the important advantage that can be applied given any set of feasible coalitions as an input and hence, it can directly model network-based coalitions as the ones we are interested in. The CSG problem is formulated as a binary integer programming problem containing a set of binary decision variables $x_S \in \{0, 1\}$, one per feasible coalition $S \in F(G)$. Then, solving the CSG amounts to solving the following IP:

$$max \sum_{S \in F(G)} v(S) \cdot x_S$$

Subject to:

(1) Each energy customer can join at most one coalition:

$$\forall a_i \in A : \sum_{S \in F(G) | S \ni i} x_S = 1$$

where having a variable $x_S = 1$ corresponds to coalition S being selected in the optimal coalition structure CS^*.

3.4 Core-Stable Payoff Distribution

Given the optimal coalition structure, we can compute a core element (or alternatively, detect that no core allocation exist) by solving a linear program (LP). Linear programs can be solved in polynomial time in the number of variables and constraints. Our aim is to find a set of negative real values that stand for agent's payments ρ, one $\rho_i \in \rho$ for each agent $a_i \in A$. Finding such stable payments, once CS^* has been found, amounts to solving the following LP:

$$min \;\; \rho(A)$$

Subject to:

(1) There are no deviating coalitions for these payments:

$$\forall S \in F(G) : \; \rho(S) \geq v(S)$$

(2) Agents payments are equal or lower than 0 (i.e. agents do not make a positive profit exploiting other agents)

$$\forall a_i \in A : \; \rho_i \leq 0$$

Then, if the value of the objective function of this LP yields to the value of the optimal coalition structure, $\rho(A) = v(CS^*)$, then the problem has a non-empty core and the values ρ define an allocation in the core. Otherwise, the problem has an empty core. It should be emphasized that the optimal coalition structure CS^* is a given parameter which means that although this program can be solved in polynomial time it needs as input the outcome of the IP program defined in Section 3.3 which is NP-Hard. The LP defined above has a number of variables equal to the number of agents, $|A|$, and a number of constraints linear to the number of feasible coalitions (again in the worst case represented by a complete graph this number is exponential to the number of agents).

It is worth noting that several techniques have been developed to solve IP and LP problems such as the ones defined in Section 3.3 and 3.4 (e.g. the dual simplex method, and the interior-point algorithm, linear relaxation coupled with branch-and-bound). Thus, next in the experimental section, we use standard, off-the-shelf-software such as CPLEX to solve them.

4. EMPIRICAL EVALUATION

In this section we provide an empirical evaluation of the coalition formation model among energy consumers introduced in Section 3. The IP and the LP problems defined in Sections 3.3 and 3.4 for computing the optimal coalition structure and the core-stable payments are solved using implementations on CPLEX 12.3. First, we explain the details of our experimental setup in Section 4.1. Next, we analyse our empirical results in Section 4.2.

4.1 Empirical settings

4.1.1 Problem generation

To analyse the sensitivity of the coalition formation process with respect to the underlying network topology, we evaluate our model on three different network models with different density levels. Formally, the density of a graph is defined as the ratio between the number of links and the number of agents in the graph ($\frac{|E|}{|A|}$). In more detail, in our experiments we test our model on the following network configurations:

Random Networks. Graphs are created by *randomly* adding a number of links d for each agent. Densities used in this case are: $d = 1$ (low), $d = 2$ (medium) and $d = 3$ (high).

Market	$\mathbf{p}_F$	$\mathbf{p}_D$	$\mathbf{r}_M$
M1	70	80	1
M2	70	80	70/80
M3	1	2	1/2

Table 1: Different market conditions explored in the experiments varying the prices of forward ($\mathbf{p}_F$), day-ahead ($\mathbf{p}_D$) and market ratios ($\mathbf{r}_M$).

Scale Free Networks. Graphs are created by using an implementation of the Barabasi-Albert model. At each step, a node is added and attached to d neighbours using a biased random selection giving more chance to a node if it has a high degree. Graphs are generated using three different densities: ($d = 0.92$, low), ($d = 1.75$, medium) and ($d = 3.17$, high).

Small-World Networks. Graphs are created by following the Watts and Strogatz model. This model generates a ring of graph where each node is connected to its k nearest neighbours in the ring (k/2 on each side, which means k must be even). Then it process each node on the ring "rewiring" each of their edges toward the not yet processed nodes. The rewiring process chooses a node randomly among the ones not yet processed and takes place according to a rewiring probability of 0.1. Graphs are generated using three different values for parameter k: $k = 2$ ($d = 1$, low), $k = 4$ ($d = 2$, medium) and $k = 6$ ($d = 3$, high).

Notice that whereas scale free and small-world networks are known to capture some characteristics of social networks [9], random networks constitute a more synthetic model for our domain. All experiments are run using networks of 12 nodes. For each instance, the energy profile of each node is randomly selected from a real dataset composed of energy profiles characterizing the real domestic electricity consumption of 5000 households in the United Kingdom. Each consumer has been monitored for a time period of a month (December 2009), recording the power consumption every half an hour, for a total of 48 daily time slots[1].

4.1.2 Market's parameters

As described in Section 3.1.1, the value of a coalition in our model depends on two market parameters: the price of the electricity in the forward, p_F, and the day-ahead market, p_D (although the price of electricity in the day-ahead market varies on each time slot, we consider here that p_D is calculated by averaging the hourly price of a day). In our experiments, we explore three different markets conditions, denoted as M1, M2 and M3 and detailed in Table 1, to evaluate how the coalition structure formation process will respond to price signals.

Notice that whereas in M1 agents follow a naive buying strategy in which the market price ratio is set to buy in the forward market the minimum continuous consumption, in M2 and M3 the market ratio price ($r_M = \frac{p_F}{p_D}$) is the one corresponding to market prices and thus the buying strategy minimizes the amount to be paid by the coalition. Regarding market prices, in $M1$ and $M2$ prices used are those of current electricity markets in Italy[2] whereas $M3$ explores a different scenario in which buying in the forward market is more incentivized with better prices.

[1]The initial data contained some corrupted entries due to a problem with a sensor, so before running any experiment, the data has been filtered keeping only valid entries.

[2]Available at: `http://www.mercatoelettrico.org/En/Default.aspx`

4.2 Results

We evaluate our model by performing repeated simulations (50 instances) for each possible configuration detailed in section above. Next sections provide an analysis of our results in terms of individual consumer gain and the structure of the formed coalitions.

4.2.1 Consumer's social gain

In this section we analyse the consumers' effective gain obtained by adopting the coalitional approach proposed with respect to the non-coalitional one, composed of singleton coalitions. Figure 3 show the results for 12 agents on a random, scalefree and small-world networks in the three different market scenarios respectively. Only instances for which empty core was not detected are considered in these results. Let ρ_i be the payment of agent $a_i \in A$ in a coalition and $v(\{i\})$ the payment of the agent in its singleton coalition. Then, the average percent consumer gain is assessed as $\frac{\sum_{a_i \in A} \rho_i - v(\{i\})}{\sum_{a_i \in A} v(\{i\})}$[3]. We also plotted the standard error of the mean as a measure of the variance in each graph. Results show that in all configurations, although as expected the average percent consumer gain is increased with density (more links among agents lead to more feasible energy coalitions among them) this increment is not significant. Regarding different market conditions, the average percent consumer gain is much higher (around 10%) in M3 that in M1 and M2 (around 1%). Thus, the economical incentive to join coalitions is directly proportional to the economical incentive to enroll in the forward market.

Table 2 shows the percentage of instances under each configuration for which the core was detected as empty. Notice that in all network topologies, the number of instances for which the core is empty increases with the density of the network. These results are coherent with the well-known results that any acyclic network (which has by definition the lowest density) is guarantee to have a non-empty core [6]. As we increase the density the number of cycles also increase and results show that the probability of core emptiness is higher (i.e. a higher number of instances show the inexistence of an stable economical agreement among consumers). Regarding different network topologies, we observe that the number of instances with core emptiness is higher in scale free networks, where the links are concentrated on hubs, than not on random and small-world networks, where each node in average have the same degree. Finally, we also observe that the number of instances with core empty is much higher on M1 and M2 than not in M3. Although we need to perform a deeper analysis on these results, they lead to the hypothesis that the larger the distance of prices in the market the less the probability of having an empty core in the coalitional game.

4.2.2 Structure of energy coalitions

In this section we analyse the structure of the energy coalitions obtained in the experiments. For each configuration, we plot the mean of the minimum, average and maximum size of coalitions formed. Figure 4 plots the results for networks of 12 agents on a random, scale free and small-world networks in two different market scenarios. We also plotted the standard error of the mean as a measure of the variance in each graph. Market scenario M3 is omitted because we detected that the grand coalition was formed in all tested instances. In contrast we observe that for markets M1 and M2, the market conditions lead to coalitions of middle size in all

[3]Notice that since consumers payments are an imputation of the optimal coalition structure this is equivalent to $\frac{v(CS^*) - \sum_{a_i \in A} - v(\{i\})}{\sum_{a_i \in A} v(\{i\})}$

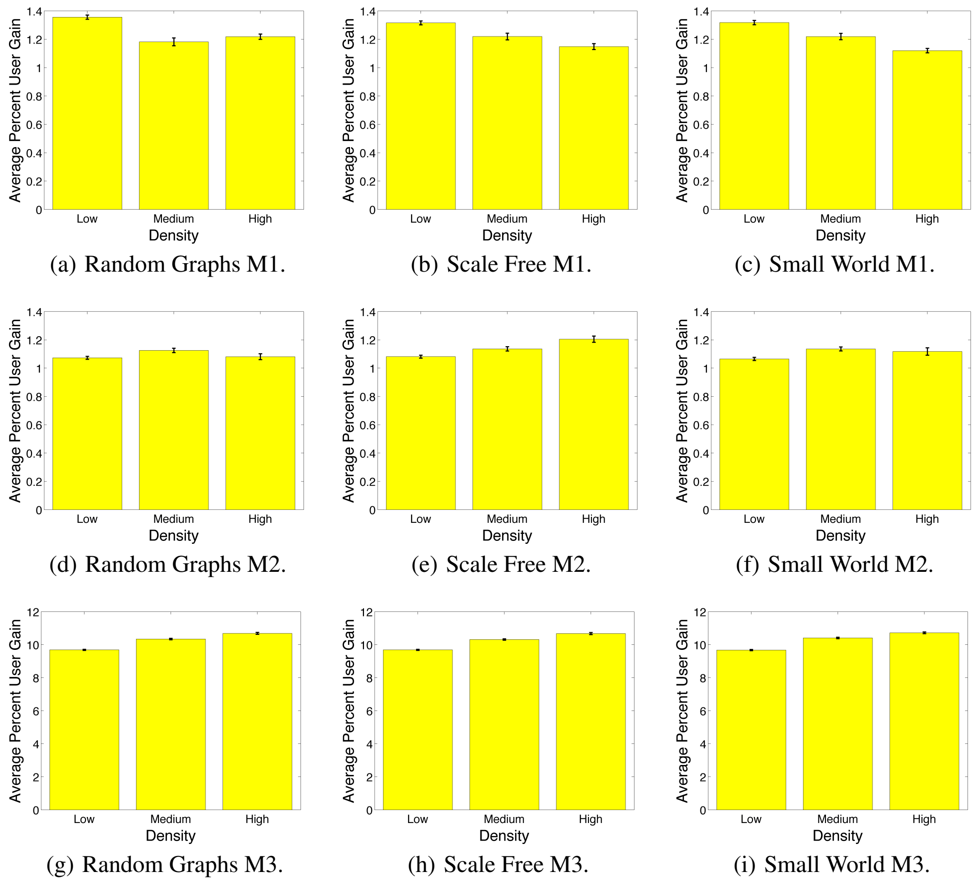

(a) Random Graphs M1. (b) Scale Free M1. (c) Small World M1.
(d) Random Graphs M2. (e) Scale Free M2. (f) Small World M2.
(g) Random Graphs M3. (h) Scale Free M3. (i) Small World M3.

Figure 3: Graphs showing the average percent gain of consumers on different topologies and densities under market conditions M1 (a)-(c), M2 (d)-(f) and M3 (g)-(i).

network structures. Therefore, our results show that larger differences between prices in the two markets, leads to larger coalitions sizes and that the structure of the coalitions formed is very sensitive to these market conditions. Finally, we also observe that as we increase the density of the network, more coalitions of middle size are formed since the size of the maximum coalition decreases with density whereas the average size increases. In contrast, low density networks tend to lead to larger coalitions.

Topology	**Density**	% **Empty Core**		
		M1	M2	M3
Random	Low	8%	0%	0%
	Medium	50%	26%	6%
	High	56%	44%	10%
ScaleFree	Low	0%	0%	0%
	Medium	52%	22%	2%
	High	46%	38%	12%
SmallWorld	Low	8%	6%	2%
	Medium	46%	18%	8%
	High	46%	48%	6%

Table 2: Percentage of instances with empty core under different configurations.

5. CONCLUSIONS AND FUTURE WORK

In this work we proposed a novel demand-side peak strategy that promotes the formation of coalitions among energy consumers with complementary energy needs. On so doing, we addressed the challenges that arise in the formation and management of these energy coalitions, the so-called virtual energy consumers (VEC's), by modeling the VEC formation process as a coalitional game. We defined a metric to evaluate coalitions taking into account that members of a VEC are typically motivated to minimize their joint payment within the electricity markets, while capturing the synergies that exist between consumers to improve the efficiency of the grid. Our model uses social networking as a tool for consumers to provide member engagement and trust on energy coalitions. Thus, we defined an algorithm that allows agents organised into a network to list all network-feasible coalitions and compute coalition values in a distributed fashion. We used linear programming techniques to efficiently identify the most efficient VEC's and to allocate the payments to individual members of VECs while taking into account that each consumer is typically motivated to maximize its own profit (as defined by core-stable solution concepts).

Secondly, we tested our model on a real dataset varying the topology and density of the social network and the market conditions. Our results show that whereas the density of the social network does not affect significantly consumer's coalitional gain, it affects the stability of the economical agreement among consumers (in many dense networks such stability simply does not exist). We also show that, as the distance between the price of energy in the forward market and the day-ahead market decreases, not only the coalition formation process yields to higher gains, but coalitions are more likely to be stable (this effect is observed even in very dense networks). Finally, depending on the economical incentive to buy in the forward market, we show that the game process converges

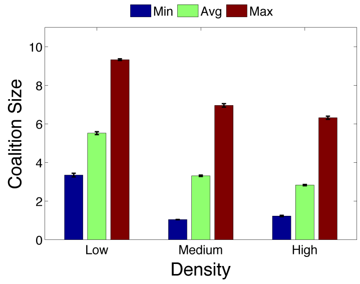

(a) Random Graphs M1.

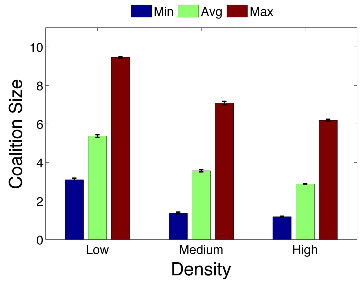

(b) Scale Free M1.

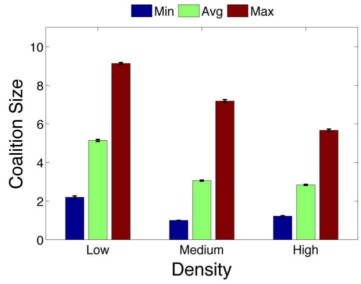

(c) Small World M1.

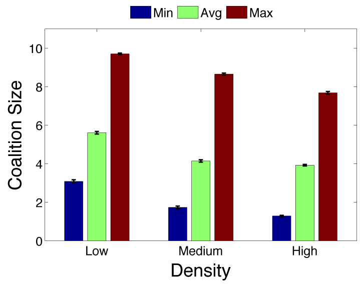

(d) Random Graphs M2.

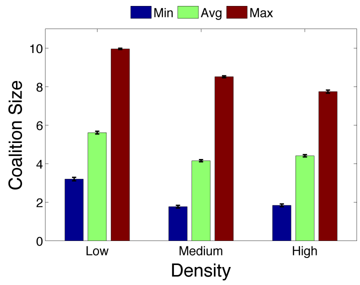

(e) Scale Free M2.

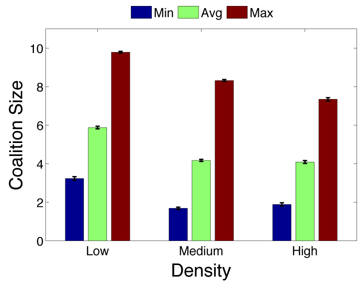

(f) Small World M2.

Figure 4: Graphs showing the minimum, average and maximum size of coalitions formed on different topologies and densities under market conditions M1 (a)-(c), M2 (d)-(f).

from singleton coalitions to the absorbing state of the grand coalition, through a wide variety of middle size coalitions (i.e. as the distance between prices in the forward and the day-ahead market increases, the sizes of the coalitions formed also increases).

As a future work, we plan to explore multiple lines. First, given the decentralisation nature of this domain, it would be desirable to provide a decentralised solutions for the coalition structure generation and payoff distribution activities instead of the centralised linear programming one used in this paper. Second, considering the scale and dynamism of the optimisation problem, it would be important to explore approximate solutions to the proposed model that provide more scalability. Finally, it is left as future work to test the proposed model with more sophisticated metrics to predict coalition consumption that explicitly model the risk associated with the decisions made by the agents.

6. REFERENCES

[1] Grid 2030. a national vision for electricity's second 100 years. Washington DOE, 2003. http://www.oe.energy.gov/DocumentsandMedia/Electric_Vision_Document.pdf.

[2] Tradewind, further developing europe's power market for large scale integration of wind power. report d. 4.1., 2007. `http://www.trade-wind.eu/fileadmin/documents/publications/D4.1_Summary_report_of_market_rules.pdf`.

[3] Vademecum of the italian power exchange. Gestore dei mercati energetici S. p. A., 2009. `http://www.mercatoelettrico.org/En/MenuBiblioteca/documenti/20091112VademecumofIpex.pdf`.

[4] D. Avis and K. Fukuda. Reverse search for enumeration. *Discrete Applied Mathematics*, 65(1-3):21–46, 1996.

[5] R. Dechter. *Constraint processing*. Elsevier Morgan Kaufmann, 2003.

[6] G. Demange. On Group Stability in Hierarchies and Networks. *Journal of Political Economy*, 112(4):754–778, August 2004.

[7] G. Gutin, A. Johnstone, J. Reddington, E. Scott, A. Soleimanfallah, and A. Yeo. An algorithm for finding connected convex subgraphs of an acyclic digraph. In *Algorithms and Complexity*. College Publications, 2008.

[8] T. Michalak, J. Sroka, T. Rahwan, M. Wooldridge, P. Mcburney, and N. Jennings. A distributed algorithm for anytime coalition structure generation. In *AAMAS*, pages 1007–1014, May 2010.

[9] M. Newman. *Networks: An Introduction*. Oxford University Press, 2010.

[10] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge MA, USA, 1994.

[11] T. Rahwan. *Algorithms for Coalition Formation in Multi-Agent Systems*. PhD thesis, University of Southampton, August 2007.

[12] T. Rahwan and N. R. Jennings. An improved dynamic programming algorithm for coalition structure generation. In *AAMAS*, pages 1417–1420, 2008.

[13] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings. Agent-based control for decentralised demand side management in the smart grid. In *AAMAS*, pages 5–12, 2011.

[14] G. Strbac. Demand side management: Benefits and challenges. *Energy Policy*, 36(12):4419–4426, December 2008.

[15] S. L. Tompros, N. P. Mouratidis, M. Draaijer, A. Foglar, and H. Hrasnica. Enabling applicability of energy saving applications on the appliances of the home environment. *IEEE Network*, 23(6):8–16, 2009.