Cooperatives of Distributed Energy Resources for Efficient Virtual Power Plants

Georgios Chalkiadakis[1] Valentin Robu[2] Ramachandra Kota[3] Alex Rogers[2]
Nicholas R. Jennings[2]
{gc2,vr2,rck05r,acr,nrj}@ecs.soton.ac.uk
[1] Dept. of Electronic and Computer Engineering, Technical University of Crete, Greece
[2] School of Electronics and Computer Science, University of Southampton, UK
[3] Secure Meters Ltd., Winchester, UK

ABSTRACT

The creation of Virtual Power Plants (VPPs) has been suggested in recent years as the means for achieving the cost-efficient integration of the many distributed energy resources (DERs) that are starting to emerge in the electricity network. In this work, we contribute to the development of VPPs by offering a game-theoretic perspective to the problem. Specifically, we design *cooperatives* (or "cooperative VPPs"—CVPPs) of rational autonomous DER-agents representing small-to-medium size renewable electricity producers, which coalesce to profitably sell their energy to the electricity grid. By so doing, we help to counter the fact that individual DERs are often excluded from the wholesale energy market due to their perceived inefficiency and unreliability. We discuss the issues surrounding the emergence of such cooperatives, and propose a pricing mechanism with certain desirable properties. Specifically, our mechanism guarantees that CVPPs have the incentive to truthfully report to the grid accurate estimates of their electricity production, and that larger rather than smaller CVPPs form; this promotes CVPP efficiency and reliability. In addition, we propose a scheme to allocate payments within the cooperative, and show that, given this scheme and the pricing mechanism, the allocation is in the core and, as such, no subset of members has a financial incentive to break away from the CVPP. Moreover, we develop an analytical tool for quantifying the uncertainty about DER production estimates, and distinguishing among different types of errors regarding such estimates. We then utilize this tool to devise protocols to manage CVPP membership. Finally, we demonstrate these ideas through a simulation that uses real-world data.

Categories and Subject Descriptors

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

General Terms

Economics, Experimentation

Keywords

energy and emissions, incentives for cooperation, coalition formation, simulation

Cite as: Cooperatives of Distributed Energy Resources for Efficient Virtual Power Plants, G. Chalkiadakis, V. Robu, R. Kota, A. Rogers, N. R. Jennings, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems – Innovative Applications Track (AAMAS 2011)*, Tumer, Yolum, Sonenberg and Stone (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. 787-794.

1. INTRODUCTION

The vision of a *"Smart Grid"* [12], and the resulting creation of a robust, intelligent electricity supply network which makes efficient use of energy resources and reduces carbon emissions, is a challenge that has been recently taken up by a growing number of researchers [6, 3, 7, 14, 15]. In this context, one of the main problems facing the energy supply industry is how to best achieve the utilization of the *distributed energy resources (DERs)* that, in recent years, have appeared in the electricity network. Such DERs range from electricity storage devices to small and medium capacity (2kW-2MW) renewable energy generators.

In principle, employing DERs to produce energy could reduce reliance on conventional power plants even by half [10]. Unlike conventional power plants that lie on the transmission network and are "dispatched" (i.e., called in to produce energy when needed) by the national electricity transmission network operators (termed *the Grid* herein), DERs lie in the distribution network and, due to their small size, they (and their capacity) are "invisible" to the Grid. Thus, they cannot be easily dispatched to meet demand. Moreover, due to their decentralized nature and small size, DERs are either invisible to the electricity market as well, or, lack the capacity, flexibility or controllability to participate in a cost-efficient way [10].

Now, the *reliability of supply* is a major concern of the Grid. It is essential that independent suppliers are reliable, since the failure to meet production targets could seriously compromise the smooth operation of the network as a whole. In contrast, given the unpredictability of renewable energy sources, the DERs would usually struggle to meet power generation targets when operating alone. This normally prohibits them from striking profitable deals with the Grid, and keeps them out of the electricity market for fear of suffering penalties specified in contracts (driving them to sign low-profit contracts with third-party market participants instead) [10]. To address this issue, in recent years many countries (e.g., in the EU) have enacted policies that guarantee the sale of electricity from small-scale producers to the Grid in pre-determined *feed-in tariffs* that are generally above market prices. Such policies were conceived by the need to promote the incorporation of renewable energy sources into the Grid, so that they generate appreciable percentages of total demand. However, with the number of DERs expected to rise to hundreds of thousands, and with the variable generation seen as another uncertainty to be addressed in real time through active Grid management, this is clearly unsustainable.

To counter these problems, the creation of *Virtual Power Plants (VPPs)* to aggregate DERs into the virtual equivalent of a large power station, and thus enable them to cost-efficiently integrate into the market, has been proposed in recent years. [7, 10]. A VPP is a broad term that intuitively represents the aggregated capabilities of a set of DERs. For example, it can be thought of as a portfolio of DERs, as an independent entity or agent that coordinates DERs pooling their resources together, or as an external aggregator that "hires" DERs to profit from their exploitation.

In our work here, we propose that power-producing DERs coalesce together to form *cooperatives* of agents that can profitably be integrated into the Grid, such a cooperative corresponding to a virtual power plant. Viewing the DERs as autonomous agents is natural, due to their distributed nature and inherent individual rationality, and enables them to realize their full potential as self-interested market units (as it allows for the possibility that even the smallest of DERs can carry out certain communication and intelligent decision making tasks on their own, but without imposing this as a requirement). We call these coalitions of DER-agents "cooperatives" because of *(a)* their completely decentralized nature; *(b)* their ability to sell their production without relying on any external entity that profits by using their members' resources; and *(c)* the fact that members willingly participate in a coalition, as it is in their best interests to do so. Of course, the mechanisms described in this work can also be readily used by any company that wishes to attract DERs as suppliers, aiming to resell their energy to the Grid. In the rest of the paper, we will use the terms "cooperative" and "cooperative VPP (CVPP)" interchangeably.

Given the issues discussed above, it is only natural that the Grid should encourage the emergence of cooperatives, by guaranteeing the purchase of CVPP energy at competitive rates. To this end, in this paper we incorporate ideas from mechanism design and cooperative game theory, and put forward an energy pricing mechanism to be employed by the Grid. The mechanism can be seen as an efficient alternative to feed-in tariffs, and so promotes the incorporation of the DERs (as CVPPs) in the Grid. In some detail, our mechanism promotes supply reliability, guaranteeing that CVPPs truthfully provide the Grid with estimates of their electricity production that are as accurate as possible. Further, they are rewarded for increased production, while the Grid maintains the ability to decide the flexibility of the mechanism and its degree of independence from market fluctuations. Building on that key contribution, we then propose a payment scheme to allocate payments within the cooperative, and show that, given this scheme and the pricing mechanism, a CVPP can guarantee payments to its members such that no subset of them has a financial incentive to form a CVPP of its own. Formally, we guarantee that, provided DER production estimates are accurate, the payments to CVPP members lie in the set of *core* allocations of the corresponding coalitional game [9]. We then develop a method that quantifies the uncertainty regarding DER production estimates and distinguishes between different types of errors in predicted production (i.e., those specific to individual DERs, and those common within whole DER clusters), and employ it to devise CVPP membership management protocols.

This is the first paper to discuss the formation of VPPs from a game-theoretic standpoint, extrapolating as it does mechanism design and cooperative game theory concepts and techniques to this domain. As such, this work demonstrates that multiagent research can provide the energy industry with solutions regarding the successful integration of DERs in the supply network. Note that this research has the potential of short to mid-term applicability in realistic settings, as several power trading companies that buy electricity from small scale producers to sell to the Grid already exist. Examples include *Flexitricity* in the UK and *Tata Power Trading Company Ltd.* in India (business description available online).

2. RELATED WORK

Here we briefly review existing related work that provides intelligent agents—and, more generally, AI research—solutions to energy-related problems. To begin, we note that researchers in the community have recently presented economics-inspired work to tackle such problems. Specifically, Vytelingum *et al.* [15] proposed strategies for the management of distributed micro-storage energy devices that adapt to the electricity market conditions. In separate work, they developed a market-based mechanism to automatically manage the congestion within the system by pricing the flow of electricity, and proposed strategies for traders in the Smart Grid [14].

However, ideas from *cooperative game theory* in particular—i.e., from the branch of game theory that studies the problem of forming *coalitions* of cooperating agents—have been used in the broader energy domain for more than a decade. Yeung *et al.* [16] employ coalitional game theory in a multiagent system model of the trading process between market entities that generate, transmit and distribute power. Also, Contreras *et al.* [2, 1] presented a *bilateral Shapley value* negotiation scheme to determine how to share the costs for the expansion of power transmission networks among coalescing rational agents.

Turning our attention to VPP-specific literature, Pudjianto *et al.* [10] stress the need to integrate DERs into the electricity network in an organized and controllable manner through participation in a VPP structure, and discuss the subsequent technical and commercial benefits to the electricity network as a whole. They also clearly outline the economic advantages to DERs, demonstrating as they do through specific examples that VPPs can be used to facilitate DER access to the electricity market. Dimeas and Hatziargyriou [6] also call for the emergence of VPPs, and essentially suggest an organizational structure that makes use of interacting coalitions to this purpose. Similarly, Mihailescu *et al.* [8] propose the use of coalition formation to build VPPs, but do not provide the details of the formation process or offer specific game-theoretic solutions—as they do not discuss issues of individual rationality or incentive compatibility. Though all of those papers advocate the creation of VPPs, they do not describe specific mechanisms for the market-VPP interface or the interactions among VPP members.

In contrast, the *PowerMatcher* (see [7] for an overview) is a decentralized system architecture that has been proposed as a means to balance demand and supply in clusters of DERs. It attempts to implement optimal supply and demand matching by organizing the DER-agents into a logical tree, assigning them roles and prescribing strategies to use in their interactions. The aspect of this system most relevant to us is the one proposing the aggregation of individual agents' supply offers in a cluster, serving as a VPP through the use of an *objective agent*. Such an agent has the task of implementing a "business logic" that would guide the VPP's actions. However, the authors stop short of proposing a specific business logic. Our approach can be seen as a detailed description of just such a logic, employing game-theoretic ideas and tools to this purpose.

3. AGENT COOPERATIVES

An agent cooperative (CVPP) is a collection of participating DER agents, each of which registers with the CVPP when joining. The CVPP may possess and employ any rules, tools and functionality necessary to ensure its unconstrained and profitable operation as an enterprise. We now present briefly some key CVPP characteristics and functionality most relevant to our work here.

In most countries, the day is divided into 48 half-hour electricity trading intervals, or *settlement periods*. For each of these, electricity prices are set in the market, and specific electricity production targets are specified for the various generators the day before, given predicted supply and demand. A DER i can estimate an *expected production* value $\widetilde{\text{prod}}_{i,t_j}$ for any half-hour period t_j. This is the energy it expects to be able to supply during t_j, given any known external factors (such as the prevailing meteorological conditions) and its expected technical state. Thus, the main profile parameter that describes the production of a DER i throughout each day is its *expected production vector* $\widetilde{\boldsymbol{prod}}_i = \langle \widetilde{\text{prod}}_{i,t_j} \rangle$, describing the DER's production for every half-hour period.

Note that, besides this estimated production, there is an *actual production vector* associated with each DER i: $\boldsymbol{prod}_i = \langle \text{prod}_{i,t_j} \rangle$. The value for each prod_{i,t_j}, however, becomes known only after the corresponding period elapses. We will be using the simplified notation prod_i and $\widetilde{\text{prod}}_i$ to refer to i's production when the period t_j of reference is evident or of no significance. Furthermore, we will be using prod_C and $\widetilde{\text{prod}}_C$ to denote the production and expected production of a cooperative C of DER agents. The difference between the t_j-values of the estimated and actual production vectors, gives the DER (or, similarly, CVPP) *prediction error* for the t_j period. Note that $\text{prod}_C = \sum_i \text{prod}_i$, as the total CVPP production is just the sum of the production of its DERs. Further, we assume that $\widetilde{\text{prod}}_C = \sum_i \widetilde{\text{prod}}_i$.[1]

Now, essential functionality for the CVPP operation includes rules and procedures for *(a)* the distribution of revenues, *(b)* the aggregation of individual production estimates into CVPP-wide ones, and *(c)* membership management (admitting and expelling members). That functionality might be located on some central agent responsible for "running" the CVPP, or it could be potentially distributed over several agents. The functionality localization details are unimportant to our work here. Instead, we proceed to describe the aforementioned CVPP operational activities in depth.

4. TRUTHFUL AND RELIABLE CVPPS

In this section, we present a payment mechanism that can be employed by the Grid to promote the formation of DER cooperatives. The mechanism addresses the main hurdles the Grid faces with respect to DERs' integration—namely, the *unreliability* of their production (given DERs' dependance on uncontrollable factors, like the weather), and their *large numbers* (given that it is anticipated that hundreds of thousands of DERs would be eventually embedded within a given country's distribution network).

To begin, we elucidate the main requirements of the Grid with respect to its interaction with CVPPs, and proceed to show how they translate into the features of our payment mechanism.

(a) Reliability of supply: The Grid operators are responsible for compiling production schedules to pass to the large power plants. Currently, these are based on the predicted demand for electricity. As more supply originates from smaller generators, their predicted output will also need to be incorporated into the Grid production scheduling process. Hence, the Grid requires any entity interacting with it (such as a DER or a CVPP) to provide it with reliable production estimates, and to be able to honour any agreement to supply a specific amount. Subsequently, the Grid would be willing to reward producers that are proven to be reliable suppliers.

(b) Need to minimize the number of entities the Grid interacts with: As already mentioned, widespread small-scale production will result in a huge number of DERs being connected to the Grid. However, the Grid would obviously prefer to interact with a small number of electricity producers, as this makes it easier to manage and settle accounts. This requirement mirrors the scenario on the consumption side, where the Grid interacts with only a few large utility companies, which, in turn, interact with the millions of individual consumers. Thus, it is imperative for the Grid to promote the formation of *large* CVPPs, each with a sizeable production capacity. Larger CVPPs make it possible for the Grid to interact with a smaller number of entities, and also promote supply reliability.

4.1 Payment Mechanism

With this list of requirements in mind, we now put forward a pricing mechanism that the Grid can use when making payments to the CVPPs for their contributed energy. As discussed, the CVPPs provide their estimated production for each day-ahead settlement period to the Grid authority. As stated above, $\widetilde{\text{prod}}_C$ is the estimated production declared by CVPP C, and prod_C its actual production in the given time interval. Let price be the electricity *base price* (per kWh). The "Grid-to-CVPP" payment from the Grid G to C is:

$$V_{G,C} = \frac{1}{1 + \alpha|\widetilde{\text{prod}}_C - \text{prod}_C|^\beta} \cdot log(\text{prod}_C) \cdot \text{price} \cdot \text{prod}_C \tag{1}$$

The three first factors of this payment function (or pricing mechanism) represent the *actual price* being offered by the Grid to C. Multiplying them with the actual CVPP production (the fourth factor, prod_C) gives the actual payment to C. The mechanism has specific properties that satisfy the requirements mentioned above:

(1) The first factor, $\frac{1}{1+\alpha|\widetilde{\text{prod}}_C - \text{prod}_C|^\beta}$, depends on the accuracy of the estimates provided by the CVPP. This *accuracy factor* is a bell-shaped function of $\widetilde{\text{prod}}_C$ given the actual production prod_C parameter, as the one whose graph is depicted in Fig. 1. It simplifies to 1 when $\text{prod}_C = \widetilde{\text{prod}}_C$, proportionally decreases as the difference between them increases. Importantly, this decrease is independent of whether prod_C is greater than $\widetilde{\text{prod}}_C$ or vice versa. Parameters α and β are functions of prod_C and determine the exact shape of the curve, and can be tuned so that the factor approaches zero for $\widetilde{\text{prod}}_C$ estimates that are at a distance of prod_C away from the actual prod_C production. The use of this factor guarantees that the CVPP has the incentive to truthfully provide a highly accurate estimate of its production, as acting otherwise leads to a loss of revenue (at least in expectation).

(2) The second factor, $log(\text{prod}_C)$, increases with production and thus encourages a large CVPP size. Therefore, CVPPs with more DER members generate more energy and obtain a better overall price than smaller ones. Nevertheless, being a *log* function, the factor flattens eventually at very high production amounts. This means that, though the formation of large CVPPs is encouraged, the emergence of a single CVPP containing all DERs is not a necessary consequence. Even though small CVPPs have an incentive to merge initially, they will not merge *ad infinitum*, as there is no visible benefit after some point due to the limit linearity of the function. Of course, other reasons to prevent merging, such as geographical or technical restrictions, might exist.

(3) The third factor, price, is determined by the Grid either through supply and demand in the electricity market or through other means, and will be the same for all CVPPs participating in the market.

It is evident that this pricing mechanism promotes cooperative participation in the market, and captures the aforementioned list of requirements. First, it promotes supply *reliability*, by guaranteeing that CVPPs receive higher revenues for accurate estimates. A CVPP has an incentive to provide as accurate an estimate as possible, and has no incentive to strategize about it, as the estimate is only used by the function to check how far off the actual production was from the promised supply target. As shown above, wilfully providing a wrong or biased estimate does not improve and mostly

[1] It is conceivable that CVPP-wide estimates do *not* necessarily equal $\sum_i \widetilde{\text{prod}}_i$. This would have no impact in our results.

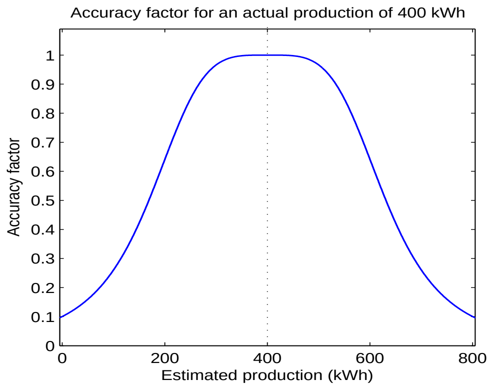

Figure 1: An accuracy factor function diagram

decreases the payment to the CVPP for the same amount of actual production. Thus, the mechanism promotes *truthfulness* on behalf of the CVPPs. Similarly, the mechanism also promotes *efficiency* at the Grid level by incentivizing the formation of large cooperatives, each attaining a substantial production aggregate.

4.2 Truthful and Reliable DER-Agents

The above mechanism incentivises the CVPP to provide accurate estimates about its production. As discussed earlier, the production of a CVPP is nothing but the aggregate of the production of all the DERs composing it. Therefore, the CVPP requires accurate estimates of production from the DERs in order to be able to calculate the production estimate to provide to the Grid. Given this, the payment from the CVPP to the DERs should encourage the DERs to truthfully provide good estimates of their production. Evidently, this mirrors the scenario between the Grid and the CVPP. Taking cue from that, we use the same principle for this "CVPP-to-DER" payment function as the Grid-to-CVPP one. Thus, the payment from CVPP C to member i for supplied energy prod_i is:

$$V_{C,i} = \frac{z}{1 + \alpha|\widetilde{\text{prod}}_i - \text{prod}_i|^{\beta}} \cdot \frac{\text{prod}_i}{\text{prod}_C} \cdot V_{G,C} \qquad (2)$$

We now describe the function in detail, demonstrating how it elicits truthful and as accurate as possible predictions from the DERs.

(1) As in Eq. 1, the first factor, $\frac{z}{1+\alpha|\widetilde{\text{prod}}_i - \text{prod}_i|^{\beta}}$, is an *accuracy factor*, encouraging the DER to provide the CVPP with its best possible production estimate. It equals z if the estimate was accurate, and drops following a bell curve otherwise. Notice that z is simply a normalization factor used to redistribute the entire $V_{G,C}$ amount back to the members. Redistribution is in proportion to the members' production and prediction accuracy—this can be easily seen with $z = \frac{\widetilde{\text{prod}}_C}{\sum_i \text{prod}_i/(1+\alpha|\widetilde{\text{prod}}_i - \text{prod}_i|^{\beta})}$. Alternatively, the CVPP can set $z = 1$ and use the residual profits to pay for maintenance costs, recruiting new members, or other such purposes.

(2) The second factor, $\frac{\text{prod}_i}{\text{prod}_C}$, gives the proportion of energy contributed by this DER w.r.t. the total CVPP production, making the payment distribution fair across all DERs.

(3) The last factor, $V_{G,C}$, denotes the total amount that is to be divided among the constituent DERs, and corresponds to the payment received by the CVPP from the Grid.[2]

To recap, by employing this payment function the CVPP promotes truthful and highly accurate predictions from its constituent DERs. A DER has an interest to truthfully and accurately report, since otherwise it does not receive the full payment corresponding to the energy it actually produced.

[2]Of course, this could be reduced by subtracting an amount if this is required to account for CVPP fees or maintenance costs.

4.3 Payment Schemes Render Stable CVPPs

Here we provide a further, game-theoretic justification for the payment scheme used by the Grid to reward CVPPs, and for that used by a CVPP to reward its members. Specifically, we show that, given the functions used to reward the cooperatives and the members payment scheme described above, and assuming that all CVPP members' stated production estimates are accurate, no members' subset has an incentive to break away and form a smaller cooperative. In addition, this result promotes the goal of large CVPP sizes.

To demonstrate this, we employ the concept of *the core* [9], the strongest of the game-theoretic solution concepts used to describe coalitional stability. Some preliminaries from cooperative game theory are in order. To begin, let N, $|N| = n$, represent a set of players; a *coalition* is a subset $S \subseteq N$. Then, a *(transferable utility) coalitional game* $G(N; v)$ is defined by its *characteristic function* $v : 2^N \mapsto \mathbb{R}$ that specifies the *value* $v(S)$ of each coalition S [13]. Intuitively, $v(S)$ represents the maximal payoff the members of S can jointly receive by cooperating, and the agents can distribute this payoff among themselves in any way. While the characteristic function describes the payoffs available to coalitions, it does not prescribe a way of distributing these payoffs. An *allocation* is a vector of payoffs $\boldsymbol{x} = (x_1, \ldots, x_n)$ assigning some payoff to each $i \in N$. Then, the *core* is the set of $\boldsymbol{x}$ payoff allocations with the property that no coalition of agents can guarantee all of its members a payoff higher to what they currently receive under $\boldsymbol{x}$. As such, no coalition would ever be motivated to break away from the grand coalition of all agents. Now, let $x(S)$ denote the payoff allocated by $\boldsymbol{x}$ to agents $S \subseteq N$, i.e., $x(S) = \sum_{i \in S} x_i$. Then, formally,

DEFINITION 1. *An allocation $\boldsymbol{x}$ is in the* core *of $G(N; v)$ iff $x(N) = v(N)$* **and** *for any $S \subseteq N$ we have $x(S) \geq v(S)$.*

That is, the value $v(N)$ of the grand coalition is efficiently distibuted by $\boldsymbol{x}$ among all agents, and the payments specified by $\boldsymbol{x}$ are such that any S already receives at least its value $v(S)$. The core of a game can be non-empty. Worse than that, it is in general **NP**-hard to determine the non-emptiness of the core (see, for example, [5]).

Returning to our setting, consider the formation of a CVPP as a coalitional game, with the characteristic function describing the value that any subset of DERs can derive by working together as a team, and the CVPP intuitively corresponding to the grand coalition of all agents. In our case, interestingly, assuming truthful and accurate DER estimates, the form of the characteristic function, $v(S) = log(\text{prod}_S) \cdot \text{price} \cdot \text{prod}_S$, allows us to prove that the payments allocated by Eq. 2 constitute a core-stable allocation, which also implies that the core of the game is always non-empty.

THEOREM 1. *Let $C = \{1, \ldots, n\}$ be a cooperative of $|C| = n$ agents, and let $G(C; v)$ be the coalitional game with characteristic function $v(S) = log(\text{prod}_S) \cdot \text{prod}_S \cdot \text{price}$ determining the value of each subset $S \subseteq C$ of agents. Consider the payoff allocation $\boldsymbol{x}$ where each agent i in C is paid according to Eq. 2—i.e., proportionally to i's contribution to the production of the CVPP (given $\widetilde{\text{prod}}_i = \text{prod}_i$). Then, $\boldsymbol{x} \in core(G)$.*

PROOF. We will show that $\boldsymbol{x}$ is in the core. We know that $\boldsymbol{x}$ distributes all payoff to the agents efficiently and therefore $x(C) = v(C)$, where $v(C) = V_{G,C}$, so the first condition of Def. 1 holds. Assume for the sake of contradiction that $\boldsymbol{x}$ is not in the core. Then, there exists some $S \subseteq C$ s.t. $v(S) > x(S)$. But $x(S) = \sum_{i \in S} x_i = \frac{\text{prod}_S}{\text{prod}_C} v(C)$ (this is easy to see by setting $\widetilde{\text{prod}}_i = \text{prod}_i$ for all i in Eq. 2). Thus: $v(S) > \frac{\text{prod}_S}{\text{prod}_C} \cdot \text{prod}_C \cdot log(\text{prod}_C) \cdot$

price $\Leftrightarrow \text{prod}_S \cdot log(\text{prod}_S) \cdot \text{price} > \frac{\text{prod}_S}{\text{prod}_C} \cdot \text{prod}_C \cdot log(\text{prod}_C) \cdot$ price $\Leftrightarrow log(\text{prod}_S) > log(\text{prod}_C)$. But, since $S \subseteq C$, this is impossible. Thus, $\boldsymbol{x}$ is in the core of $G(C; v)$. $\square$

Thus, the choice of the Grid-to-CVPP and CVPP-to-DER payment schemes described above is well justified from a game-theoretic, coalitional stability point of view also.

5. QUANTIFYING PREDICTION ERRORS

In Section 4.1 we introduced the payment function of CVPPs to their members, based partially on the accuracy of their predictions. Here we propose several methods for quantifying the uncertainty in DER predictions, and distinguishing between different types of prediction errors. This will prove helpful for devising methods to handle CVPP membership (in Section 6). To begin, consider the examples of a virtual power plant that aggregates the supply from several DER wind farms (belonging to different stakeholders) distributed in a geographical area, or from a set of solar panels installed by different houses in an extended neighbourhood. Each DER can make an error in the prediction of its future output for a given half-hour period. It is useful to distinguish between two main classes of errors:

(a) **Systematic errors:** This error type is caused by the inherent uncertainty in predicting an outside variable that is used as an input by several DERs while calculating their production estimates. For renewables, this is most likely a weather-related variable, such as wind speed or solar power. So, for example, if the meteorological office is innacurate in its prediction of wind speed at a certain time in a local area, then all the wind turbines in that area may register an error in their predicted production. We call this type of error *systematic*, as it is common to all energy resources that rely on that factor, and it is outside the control of individual DERs.

(b) **Residual errors (DER specific):** Besides the systematic errors, the predictions of an individual DER may be affected by errors caused by factors specific to itself, and (at least partially) under its control. In the example discussed above, even if a wind turbine is supplied with very accurate predictions of wind speed, its prediction of its actual output may not be that accurate (because it is an older turbine, requires maintenance work, and so on).

Against this background, we now propose a statistical method for distinguishing between the different types of prediction errors. Consider a dataset consisting of m DERs in a CVPP, which belong to the same category of energy producers (e.g., wind turbines from the same area). For each of these DERs, n half-hour data points are available within some large time period $T = \{1, \dots, n\}$ (n can be quite large as the data can span several days, weeks or months).

Formally, let $\widetilde{\text{prod}}_{i,t}$ and $\text{prod}_{i,t}$ denote the estimated and actual production of DER i in a half-hour interval t. Moreover, let $\Delta_{i,t} = \text{prod}_{i,t} - \widetilde{\text{prod}}_{i,t}, \forall i = \{1, \dots, m\}, \forall t \in T$ denote i's prediction errors in t. Given the 2-dimensional error matrix with entries $\Delta_{i,t}$ as defined above, we can define the *average* prediction error across all DERs for some $t \in T$ as: $\mu_t^{\Delta} = \frac{\sum_{i=1}^{m} \Delta_{i,t}}{m}$.

In what follows, we denote by $\boldsymbol{\Delta}_T^i$ the n-vector of errors corresponding to energy producer i for every interval $t \in T$ ($\boldsymbol{\Delta}_T^i$ is a row of $\Delta_{i,t}$ error matrix entries corresponding to i), and by $\boldsymbol{\mu}_T^{\Delta}$ the n-vector containing the average prediction errors across all DERs for all time steps $t \in T$. We can now compute the *Pearson correlation coefficient* ρ_i^{Δ} between vectors $\boldsymbol{\Delta}_T^i$ and $\boldsymbol{\mu}_T^{\Delta}$ as:

$$\rho_i^{\Delta} = \frac{cov(\boldsymbol{\Delta}_T^i, \boldsymbol{\mu}_T^{\Delta})}{\sigma(\boldsymbol{\Delta}_T^i)\sigma(\boldsymbol{\mu}_T^{\Delta})} = \frac{\sum_{i=1}^{n}(\Delta_{i,t} - \overline{\Delta_i})(\mu_t^{\Delta} - \overline{\mu_t^{\Delta}})}{\sqrt{\sum_{i=1}^{n}(\Delta_{i,t} - \overline{\Delta_i})^2}\sqrt{\sum_{i=1}^{n}(\mu_t^{\Delta} - \overline{\mu_t^{\Delta}})^2} } \quad (3)$$

where $cov(\boldsymbol{\Delta}_T^i, \boldsymbol{\mu}_T^{\Delta})$ denotes the statistical covariance between the two vectors $\boldsymbol{\Delta}_T^i$ and $\boldsymbol{\mu}_T^{\Delta}$, $\sigma(\boldsymbol{\Delta}_T^i)$ and $\sigma(\boldsymbol{\mu}_T^{\Delta})$ are their standard deviations, and $\overline{\Delta_i} = \frac{\sum_{t=1}^{n} \Delta_{i,t}}{n}$ and $\overline{\mu_t^{\Delta}} = \frac{\sum_{t=1}^{n} \mu_t^{\Delta}}{n}$ their means.

Intuitively, for each energy producer i, $\rho_i^{\Delta} \in [0, 1]$ shows how correlated its errors in predicted production were with the average errors made by the energy producers in the same category in the CVPP. In our wind turbine example, if the coefficient ρ_i^{Δ} for wind tubine i is high, it means that this turbine tends to make a prediction error when all other wind turbines in its area make a prediction error of similar proportions. Thus, its error is mostly of a "systematic" nature. If there is an uncertain, outside factor (e.g., wind speed prediction) causing an error for all these turbines, then the errors can be assigned to this factor. Conversely, if ρ_i^{Δ} is low, the errors of this wind turbine are caused by its own functioning/prediction capabilities, and appear unrelated to those of similar producers.

With this at hand, statistical theory [4] allows us to define two important measures for the error vector of each producer i: the *fraction of variance explained* by the systematic factor (also called the coefficient of determination), $FVE_i^{\Delta} = (\rho_i^{\Delta})^2$, and the *fraction of variance unexplained* (or, the fraction of residual variance) $FVU_i^{\Delta} = 1 - (\rho_i^{\Delta})^2$. In essence, these measures determine the percentage of the variance in DER i's prediction errors that can be explained by systematic factors. Thus, we can separate the variance $\sigma(\boldsymbol{\Delta}_T^i)$ in the prediction errors of each i over period T into the systematic and the residual variance, the latter defined as:

$$\sigma_{res}(\boldsymbol{\Delta}_T^i) = FVU_i^{\Delta}\sigma(\boldsymbol{\Delta}_T^i) = [1 - (\rho_i^{\Delta})^2]\sigma(\boldsymbol{\Delta}_T^i) \quad (4)$$

Thus, the residual variance provides us with a tool to determine whether the prediction error of a specific DER i is due to factors that do not affect other energy producers of the same nature and in the same area. As we shall see, this tool can be used to inform CVPP membership management decisions.

6. MANAGING CVPP MEMBERSHIP

In Section 4.3 we showed that, given the payment function described in Eq. 1, coalitions representing CVPPs are stable, in the sense that DERs do not have a financial incentive to abandon them. However, this result only holds when the DERs composing the coalition are always able to provide accurate, error-free estimates of their production. In general, cooperatives do not have an incentive to expel members, given that more members means greater expected production and thus greater expected revenues. At the same time, given Eq. 1, it is also true that, if certain DER members are consistently unreliable in their production estimates, then the additional penalty that the CVPP suffers due to increased unreliability can in the long term offset any benefits from an increased overall production. Therefore, a CVPP should perform a regular evaluation of its individual members' performance, based on which it may decide to expel some of them. In this section, we provide methods for such an evaluation.

Formally, as in Section 5 above, we consider the performance of m DERs belonging to a CVPP in a discretized time period T consisting of $t = 1, \dots, n$ half-hour periods. Furthermore, we denote by $C\backslash i$ the CVPP C if DER i was not its member. Given the Grid-to-CVPP payment of Eq. 1, we define the *marginal contribution* (or *marginal value*) of DER i to cooperative C in period t to be:

$$V_{i \to C;t}^{mg} = V_{G,C;t} - V_{G,C\backslash i;t} \quad (5)$$

Intuitively, the marginal contribution of DER i to the cooperative at any time interval is the difference between the payment that a cooperative actually receives, and the payment it would have received had i not been part of the cooperative.[3] Note that this marginal

[3]Incidentally, although perhaps intuitively appealing, using the

value is influenced by both the estimated and actual productions, $\widetilde{\text{prod}}_{i,t}$ and $\text{prod}_{i,t}$, of DER i (and, implicitly, by its errors $\Delta_{i,t}$).

Given this, we now propose a method to assess the long-term performance of i within a time frame of interest T. The same mechanism could be applied to the process of deciding whether to accept a new member in the CVPP, if historical data regarding its predictions' reliability were available.

Note that a first, simple solution would be to assess the members' performance by ranking them according to their marginal contribution during a time period T consisting of intervals $t = 1, \ldots, n$. That is, we can simply add the marginal contributions of DER i for the intervals $t \in T$: $V^{mg}_{i \to C;T} = \sum_{t=1,\ldots,n} V^{mg}_{i \to C;t}$. Then, each producer can be ranked by its marginal contribution to the revenues of the CVPP, as described by $V^{mg}_{i \to C;T}$ across the period T of interest. This method captures the exact contributions of members, but does not account for systematic errors. So, for example, a DER situated in an area with poor wind/solar power prediction for a given period, would be penalized for elements outside its control.

A fairer method would be to use the residual variance specific to each DER. Such a method involves ranking the producers according to their residual variances, as computed in Eq. 4, over a period T. The least accurate producers could then possibly be expelled from the CVPP, as a high residual variance shows their prediction accuracy underperforms that of others in the same area for a considerable period of time. However, that would have the disadvantage that it completely disregards the contributions of individual DERs to the CVPP revenues. Indeed, a CVPP could be reluctant to expel a member that, though consistently inaccurate, still contributes significantly to the CVPP production and, therefore, revenues.

Thus, here we propose a method that actually *weighs* the marginal contribution of a DER by its residual variance (normalized to $[0, 1]$ through division by the sum of residual variances across all m agents). Specifically, C calculates, for each i over T, the following:

$$\text{score}^i_T = \Big(1 - \frac{\sigma_{res}(\mathbf{\Delta}^i_T)}{\sum_{j=1}^{m} \sigma_{res}(\mathbf{\Delta}^j_T)}\Big) V^{mg}_{i \to C;T} \tag{6}$$

Intuitively, DERs with higher residual variance have their marginal contribution disregarded more, while still taking some credit for it. The CVPP then ranks the DERs in terms of their score, and has the option to expel members with low performance. The advantage of this method is that it avoids punishing individual DERs for systematic errors, while taking into account their marginal contributions at the same time.

7. EMPIRICAL EVALUATION

We tested our payment mechanisms by examining the incentives of a set of individual DERs to form a cooperative, in the context of a renewables generation scenario. The data used in our analysis comes from the *Sotavento* experimental wind farm, in Galicia, Spain, and is made freely available for research purposes from their website (http://www.sotaventogalicia.com/). The farm produces roughly the energy required to serve 12,000 homes. In what follows, we first discuss how we constructed individual wind turbine profiles from the available data, and describe our prediction scenarios. We then apply our mechanisms to this setting, demonstrate the benefits to individual turbines from forming a cooperative, and evaluate our method for ranking DERs according to prediction performance.

To begin, the main characteristic of a wind turbine is its *power curve*, describing, for a given level of wind speed, its electrical output (in MW). The generic power curve of wind turbines is typically a *sigmoid function*, with a threshold level, beyond which the power output increases more sharply. A turbine's *nominal capacity* describes its maximum power (and, subsequently, energy produced per hour) output for "optimal" wind speed.

The Sotavento farm contains 24 wind turbines, with an installed nominal capacity of 17.5 MW which jointly produce an average of 38,500 MWh yearly. The available dataset we used in our experiments contains, for each *hourly* slot for the entire year from 2 September 2009 to 2 September 2010, both the actual wind speeds recorded, as well as the farm production (in kWh) for that time slot. There are 8600 records/year provided in total, due to some records being corrupt. Fig. 2 shows a scatter plot of all the yearly data points from Sotavento, as well as the power curve (the sigmoid function) we derived based on this data.

Next, we divided the derived curve for the entire farm (i.e., in our terminology, the CVPP) into 24 identical power curves, one for each individual wind turbine (DER). Note that, while no detailed data was available about individual turbines, considering them equal in nominal production capacity is realistic, and sufficient to illustrate the functioning of our model. Therefore, based on the real data, each of the 24 turbines has a nominal capacity of ∼700 kW (or, it can supply ∼500 homes). If w_t is the wind speed at an hourly timepoint t (in m/s), the *generic* power curve of each turbine is:

$$prod^{generic}_{i,t}(w_t) = \frac{700}{1 + e^{0.66*(9-w_t)}} \tag{7}$$

The shape of this function for each individual turbine is the same as that in Fig. 2, but with a maximal capacity of 700 kW.

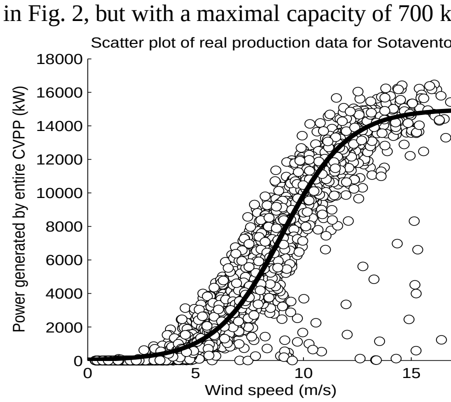

Figure 2: Scatter plot of the yearly data points from Sotavento.

7.1 Forming CVPPs of Wind Turbines

Although the Sotavento site provides real data about production and wind speeds, it does not provide us with any long-term data about the *predictions* of individual turbines. Furthermore, all wind turbines in Sotavento are owned by the same entity (a government agency). By contrast, our goal is to examine more decentralized settings, with these turbines belonging to individual stakeholders. Specifically, our aim here is to verify experimentally that, given our payment mechanisms, "self-interested" turbines (DERs) with different abilities have an incentive to coalesce into a CVPP.

To this end, we consider experimental scenarios in which the main parameter varied is the prediction ability of individual turbines regarding future production. Formally, given a wind speed

marginal contributions to distribute the CVPP revenues to the DERs is problematic as an approach, because it compromises DER truthfulness. Specifically, it provides agents with a reason to strategize and base their reports on those of others, since their payment would be based on whether they can accurately predict and "correct" the reports of others, so that they are awarded the marginal gains resulting from improved CVPP performance. Though the study of such collective "auto-correction" mechanisms is perhaps interesting, it is out of the scope of this work.

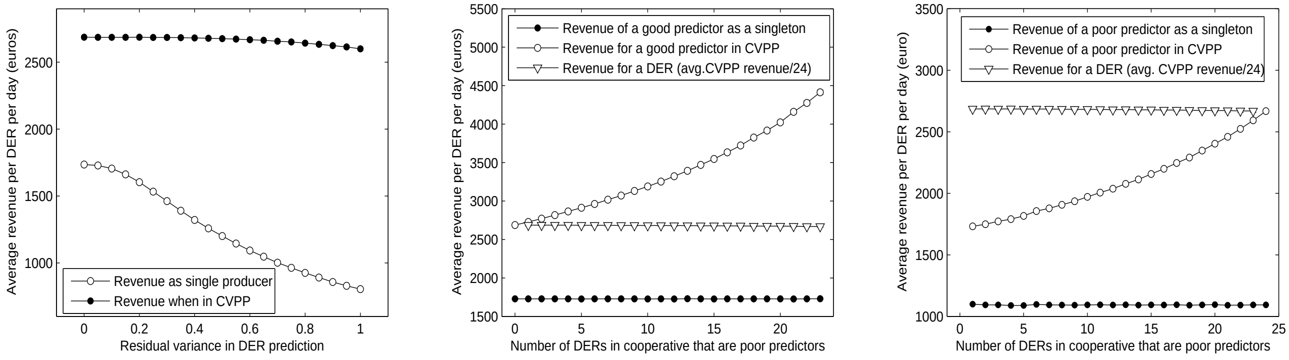

(a) (b) (c)

Figure 3: Benefits from joining a CVPP vs. selling to the Grid as an independent producer (singleton): (a) symmetric case; (b) asymmetric case: performance of *good predictors*; and (c) asymmetric case: performance of *poor predictors*. Averages are over 86,000 steps (8600 hourly time points available in a year & yearly simulation run 10 times). Error bars too small to be visible.

prediction w_t, we first compute a *generic* (idealized) production $prod^{generic}_{i,t}$ of each wind turbine i at time t using Eq. 7.[4] Then, the *actual* production for each DER $i = 1 \dots 24$ is given by

$$prod_{i,t} = prod^{generic}_{i,t} \cdot \mathcal{N}(1, \sigma_{syst})$$

where the variance factor σ_{syst} captures the systematic error that is common to all turbines (i.e., the actual wind speed is not the same as predicted). While the actual productions are drawn independently for each DER i, the deviation σ_{syst} of the normal perturbation distribution is the same for all, reflecting the fact that they are all subject to the same uncertain, outside factor (wind speed).

Now, the DERs can have rather different capabilities w.r.t. deriving future production estimates. This is captured by a DER-specific (or residual) error factor σ^i_{res}. Then, the estimated production reported by each DER $i = 1 \dots 24$ is:

$$\widetilde{prod}_{i,t} = prod^{generic}_{i,t} \cdot \mathcal{N}(1, \sigma^i_{res})$$

Against this background, we use two simulation settings to explore the benefits to individual DERs from being in a CVPP. In both settings, the number of DERs is fixed at 24 (as in Sotavento), each with generic production functions as in Eq. 7, and with the systematic error variance set to $\sigma_{syst} = 0.1$. We set $\mathrm{price} = 0.05$; this is combined with the first two factors of Eq. 1 to give the *actual price* in euros/kWh. We consider the following cases:

(a) **The symmetric case:** All DER-agents are *equally good* or *equally bad* in predicting their own production. In other words, the residual deviation σ^i_{res} is the same accross all agents i.

(b) **The asymmetric case:** The agents in the cooperative are divided into 2 classes: one of *good predictors*, having a low residual deviation $\sigma_{res(low)} = 0.05$ regarding their production estimates, and a second class of *poor predictors*, having a high residual deviation of $\sigma_{res(high)} = 0.6$. The relative proportion of the two class sizes varies from 0/24 to 24/24 (out of the 24 agents in the CVPP).

For both scenarios, we ran a series of experiments where the real wind data for all hourly intervals for an entire year was used. The simulation of the hourly wind speeds over the entire year was repeated 10 times[5] to reduce the outcomes' variance, resulting to $86,000$ tests for each data point shown in the results of Fig. 3.

[4] As already discussed, our simulation uses the real wind speeds for each hour for the 365 days in the year.

[5] The simulation parameters were chosen with the computational requirements of the various experimental settings in mind, but in all cases our results are statistically significant.

Joining a CVPP is beneficial in the symmetric case.

Turning our attention to Fig. 3(a) which depicts the results for the symmetric predictions scenario, we can see that, whatever the residual uncertainty in prediction is, individual DERs have an incentive to join together to form a CVPP. For small values of the deviation in prediction error σ^i_{res} (i.e., when all agents predictions are almost entirely accurate), this effect is due to the superadditive structure of the reward function of Eq. 1. This was not surprising, given the result of Theorem 1. Interestingly, however, the impact of our payment schemes is even more profound when highly inaccurate DERs (i.e., those with high values of residual variance) are considered. In this case, the revenue for singleton DERs more than halves when compared to their average gains when participating in a CVPP (from 1700 to 800 euro/day), as the agents are punished by the Grid for their inability to predict their production accurately.

As expected, when agents interact with the Grid as a CVPP, the cooperative's revenue also drops when its members become less accurate in prediction. However, the drop is much smaller, from 2700 to about 2600 euro/day for each of the 24 members. This is mainly because, if added over the entire cooperative, residual prediction errors cancel each other. Thus, quite interestingly, even a virtual power plant consisting of 24 DERs with poor prediction ability is able to issue a reasonably accurate estimate to the Grid.

Results for the asymmetric case.

We now examine a setting in which DERs can be separated into two distinct classes, one of *good* and one of *poor* predictors (with a residual variance of $\sigma_{res(low)} = 0.05$ and $\sigma_{res(high)} = 0.6$ respectively). The main experimental parameter varied here is the number of agents of each type that make up the CVPP; these are varied from 0 to 24.

Simulation results appear in Fig. 3(b) and 3(c). We observe that, in general, both types are better off being in a CVPP than interacting with the Grid as singletons. This is regardless of whether the other participants are good predictors or poor. However, there are some additional interesting observations to be made in this setting.

Somewhat surprisingly, *good predictors* actually do much better if the rest of the cooperative members are poor. The reason for this is the way the CVPP-to-DER payment redistribution function works. If an agent is the only accurate one (or among the few accurate ones) in the cooperative, it gets a large proportion of the joint payments, as it is among the few with a low error factor, and thus enjoys high returns following the (normalized to reward accuracy, as explained in Section 4.2) redistribution of CVPP's revenues.

In general, *poor predictors* also have a strong incentive to join the CVPP, as the results in Fig. 3(c) show. An interesting point to note is that it would appear from these results that both poor and good predictors prefer the other agents in the cooperative to be poor predictors (unless their errors are all biased towards the same direction and thus do not cancel out—an improbable scenario for large CVPPs). However, as shown in our figures, a *random* member of the cooperative would on average expect to do slightly better if the number of good predictors is high, as the cooperative as a whole gains more revenue on average in that case.

7.2 Ranking DERs by Prediction Performance

For the last set of results, we use a similar setting as the asymmetric case described above. We divide the DER-agents into two categories: *good predictors* (with $\sigma_{res(low)} = 0.05$) and *poor predictors* (with $\sigma_{res(high)} = 0.3$). The number of each agent type in the cooperative was varied from 1 to 23 (out of 24 agents in total). Recall that in Section 6, two methods for assessing the contribution of a DER i to the CVPP were discussed: one based on only its marginal contribution to the cooperative, and the other taking into account both i's marginal contribution and the residual error variance $\sigma_{res}(\Delta_T^i)$. In our experiments, we compare these two methods, taking T to be one year of hourly data, as before.

The graph in Fig. 4 shows, for settings with $k = 1 \ldots 24$ poor predictors, the number of *real* poor predictors detected by each method (i.e., how many actual poor predictors are among the k lowest scoring agents returned by each method used). Note that the ranking shown is actually an average over 25 runs, sufficient to reduce the results' variance to very low levels (since, in fact, each data point represents the results from 25 years of real hourly data). As we can see, the method that weighs marginal values by residual variance (Eq. 6), making use of the techniques of Section 5, is clearly better in distinguishing poor predictors than ranking by marginal contributions alone; in fact, it rarely identifies a predictor wrongly in this setting. In contrast, the strategy of ranking solely by marginal values does degrade, especially when the number of good predictors roughly equals that of poor ones. In any case, the results in this setting show that both methods manage to distinguish poor predictors from good ones with a very high degree of accuracy.

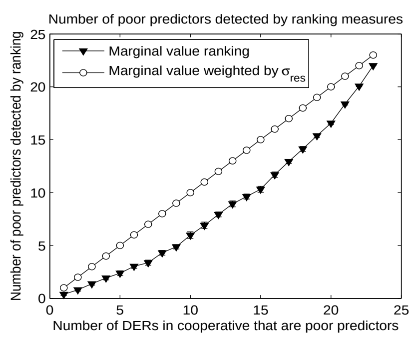

Figure 4: Results for the efficiency of ranking measures.

8. CONCLUSIONS

In this paper we applied several game-theoretic ideas in the energy domain. We presented a pricing mechanism that can be used as an alternative to feed-in tariffs, in order to promote the creation and cost-efficient operation of DER cooperatives. We also proposed a method to allocate CVPP revenues to its members, and showed that this method promotes CVPP stability (assigning payoffs that are core-stable, under the condition of DER accuracy). We also showed that the payment functions incentivize truth-telling when CVPPs interact with the Grid and when DERs interact with the CVPP; and that our methods promote supply reliability and production efficiency. Moreover, we provided a generic method for CVPP membership management, which was experimentally shown to be successful in ranking DERs w.r.t. predictions' accuracy. Crucially, our ideas were evaluated on scenarios using data from a real-world wind-farm. Our results confirm that joining CVPPs which make use of our proposed payment schemes is almost always beneficial to any individual DER.

In future work, we intend to study alternative pricing schemes to the one proposed here. For instance, residual errors-related information could perhaps be incorporated in the payment function. Doing so optimally and in a fair manner is not straightforward, since determining the residual part of the error requires the study of an agent's performance over an extended period, while the payment function rewards the agent for its immediate performance. We also intend to examine alternative ways to distribute rewards among CVPP members, perhaps by utilizing their *Shapley value* [9]. Although its exact calculation is an intractable problem, the use of *bilateral* Shapley value approximation schemes could be an option. Furthermore, assuming DERs could provide production estimates in the form of a full distribution (rather than just an expected value), it would be interesting to devise *scoring rules* [11] to elicit those estimates, and to reward both estimates that turn out to be accurate, and those provided with high precision (low variance). Moreover, we would be interested in implementing a web service to accommodate CVPP formation and member management activities.

9. REFERENCES

[1] J. Contreras, M. Klusch, and J. Yen. Multi-agent coalition formation in power transmission planning. In *Proc. of the 4th Int. Conf. on Artificial Intelligence in Planning Systems*, 1998.

[2] J. Contreras, F. F. Wu, M. Klusch, and O. Shehory. Coalition formation in a power transmission planning environment. In *Proc. of Conf. on Practical Applications of Intelligent Agents and Multiagent Systems*, 1997.

[3] E. M. Davidson, M. J. Dolan, S. D. J. McArthur, and G. W. Ault. The use of constraint programming for the autonomous management of power flows. In *Proc. of the 15th Int. Conf. on Intelligent System Applications to Power Systems*, pages 1–7, 2009.

[4] M. DeGroot and J. Schervish. *Probability and Statistics*. Addison-Wesley, 2002.

[5] X. Deng and C. H. Papadimitriou. On the complexity of cooperative solution concepts. *Math. Oper. Res.*, 19(2):257–266, 1994.

[6] A. Dimeas and N. Hatziargyriou. Agent based control of virtual power plants. In *Proc. of the 14th Int. Conf. on Intelligent System Applications to Power Systems*, pages 1–6, 2007.

[7] K. Kok, M. Scheepers, and R. Kamphuis. Intelligence in electricity networks for embedding renewables and distributed generation. In *Intelligent Infrastructures*. Springer, 2009.

[8] R. C. Mihailescu, M. Vasirani, and S. Ossowski. Towards agent-based virtual power stations via multi-level coalition formation. In *Proc. of the 1st Int. Workshop on Agent Technologies for Energy Systems*, pages 107–108, 2010.

[9] R. Myerson. *Game Theory: Analysis of Conflict*. Harvard Univ. Press, 1991.

[10] D. Pudjianto, C. Ramsay, and G. Strbac. Virtual power plant and system integration of distributed energy resources. *IET Renewable Power Generation*, 1(1):10–16, 2007.

[11] L. J. Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.

[12] U.S. Department of Energy. Grid 2030: A national vision for electricity's second 100 years, 2003.

[13] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, 1944.

[14] P. Vytelingum, S. D. Ramchurn, T. D. Voice, A. Rogers, and N. R. Jennings. Trading agents for the smart electricity grid. In *Proc. of AAMAS 2010*, pages 897–904, 2010.

[15] P. Vytelingum, T. D. Voice, S. D. Ramchurn, A. Rogers, and N. R. Jennings. Agent-based micro-storage management for the smart grid. In *Proc. of AAMAS 2010*, pages 39–46, 2010.

[16] C. S. K. Yeung, A. S. Y. Poon, and F. F. Wu. Game theoretical multi-agent modeling of coalition formation for multirateral trade. *IEEE Transactions on Power Systems*, 14(3):929–934, 1999.