Interdependent Multi-Issue Negotiation for Energy Exchange in Remote Communities

Muddasser Alam, Alex Rogers and **Sarvapali D. Ramchurn**

Agents, Interaction and Complexity Research Group,
Electronics and Computer Science,
University of Southampton, U.K.
{moody,acr,sdr}@ecs.soton.ac.uk

Abstract

We present a novel negotiation protocol to facilitate energy exchange between off-grid homes that are equipped with renewable energy generation and electricity storage. Our protocol imposes restrictions over negotiation such that it reduces the complex interdependent multi-issue negotiation to one where agents have a strategy profile in subgame perfect Nash equilibrium. We show that our negotiation protocol is tractable, concurrent, scalable and leads to Pareto-optimal outcomes in a decentralised manner. We empirically evaluate our protocol and show that, in this instance, a society of agents can (i) improve the overall utilities by 14% and (ii) reduce their overall use of the batteries by 37%.

1 Introduction

Lack of access to electricity is a serious hindrance to economic and social development in the developing world (UNDP 2012, p.14), and currently affects 1.4 billion people in small communities in Sub-Saharan Africa and Asia (IEA 2010, p.239). Recent initiatives have sought to provide these remote communities with off-grid renewable microgeneration such as solar panels and electric batteries (Alam et al. 2013). At present, these resources (i.e., microgeneration and storage) are operated in isolation, however, we envision that their interconnection and autonomous coordination could result in their more efficient use. As a step towards this vision, we explore the possibility of energy exchange between homes in such communities. We represent an individual home as a software agent that acts on the household's behalf. The whole community can be perceived as a multi-agent system composed of self-interested agents that negotiate with each other to reach energy exchange agreements while maximising their own utility. However, negotiation in this context poses many issues that come from the very nature of communities and realities of life in developing countries, e.g., lack of banking systems, low-processing power at hand, absence of a centralised infrastructure. Furthermore, negotiation over energy exchange involves multiple issues, specifically, the amount of energy exchange and also, how this amount is scheduled across the day. These issues are interdependent as the recipient's utility for any period may depend on the energy received in earlier periods. This interdependent multi-issue negotiation, along with the socio-economic limitations of remote communities, make negotiation over energy exchange a very challenging task for agents.

To address this challenge, Alam et al. (2011) presented a bi-lateral negotiation protocol to facilitate negotiation over energy exchange between two agents. Their attempt is inspired by more general work of Rosenschein and Zlotkin (1994) which shows that careful design of a negotiation protocol can reduce the complexity in negotiation. Their protocol restricts the type and number of offers that agents can make such that each agent has a weakly dominant strategy to reveal its true energy needs, resulting in a Pareto-optimal outcome. Their empirical evaluation demonstrates that agents can use a smaller battery capacity (40% less) without losing their utilities when they exchange energy. However, their protocol is applicable to two agents only and not scalable to larger communities.

More general work on interdependent multi-issue negotiation is focused on two tracks. The first focuses on settings where interdependence between issues is removable or reducible. For example, Fujita et al. (2010) and Hindriks et al. (2006) remove dependencies by approximating the utility space. However, both techniques work only when a few (among all) issues are interdependent. This is not so in our case where energy storage makes all time periods interdependent. The second track (e.g., Hattori et al. 2007 and Ito et al. 2007) focuses on settings in which a mediator collects information about the agents' utility functions. This centre then finds the set of Pareto-optimal solutions, from which the agents choose one. However, these solutions require the presence of an independent mediator, capable of carrying out intensive computations. Such assumptions are hard to justify in our decentralized settings with no centre and where agents are required to negotiate directly with each other.

Against this background, we present a negotiation protocol to address the issue of negotiation over energy exchange. Our protocol imposes four key restrictions on the offers that agents can make and specifies the negotiation process in a way such that it leads to a subgame perfect Nash equilibrium (SPNE). Our work can be seen in line with Alam et al. (2011), however, our protocol is concurrent and scalable to a community. In addition, our protocol does not assume financial payments or a mediator which makes it applicable

in the decentralised remote communities. More specifically, we extend the state-of-the-art in the following ways:

1. We present a novel negotiation protocol for concurrent negotiation over energy exchange in a mutiagent system.

2. We show that this protocol leads to a subgame perfect Nash equilibrium where outcomes are Pareto-optimal.

3. We empirically evaluate our protocol against the Nash bargaining solution (NBS) and show that, in this instance, a community can use our protocol to reduce its overall battery charging by close to 37% (while via the NBS it is 49%) and improve its social welfare (sum of utilities) by 14% (via the NBS up to 17%).

The rest of paper is as follows. Section 2 presents a model of a single home while Section 3 shows a community model. Section 4 details our protocol and Section 5 discusses its properties. Section 6 establishes the benchmark and Section 7 shows an empirical evaluation. Section 8 concludes.

2   Model of an Individual Home

We model an individual home similar to that of Alam et. al (2011) and Vytelingum et al. (2011). Each home has a renewable generation unit, some loads and a battery to store electricity. Let agent a represent a home, with a *generation* $\boldsymbol{g} = (g_1, ..., g_t) \in \mathbb{R}^t_{\geq 0}$ denoting the energy it generates over $\boldsymbol{t} = (1, ..., t) \in \mathbb{N}^t$ time periods and a *load* $\boldsymbol{h} = (h_1, ..., h_t) \in \mathbb{R}^t_{\geq 0}$ denoting its load requirements. The battery is characterised by four parameters: (i) a maximum storage capacity, $q_{max} \in \mathbb{R}_{\geq 0}$, (ii) a maximum charging rate, $c_{max} \in \mathbb{R}_{\geq 0}$, (iii) a maximum discharging rate, $d_{max} \in \mathbb{R}_{\geq 0}$, and (iv) an efficiency $e \in \mathbb{R} | 0 \leq e \leq 1$. The efficiency describes the loss of energy when the battery is charged. The dynamic state of the battery is given by: the energy flow into the battery (charge) $\boldsymbol{c} = (c_1, ..., c_t) \in \mathbb{R}^t_{\geq 0} | \ \forall c_i \in \boldsymbol{c} \ 0 \leq c_i \leq c_{max}$, the flow going out (discharge) $\boldsymbol{d} = (d_1, ..., d_t) \in \mathbb{R}^t_{\geq 0} | \ \forall d_i \in \boldsymbol{d} \ 0 \leq d_i \leq d_{max}$ and the amount of charge stored in battery $\boldsymbol{q} = (q_1, ..., q_t) \in \mathbb{R}^t_{\geq 0} | \ \forall q_i \in \boldsymbol{q} \ 0 \leq q_i \leq q_{max}$. Finally, in some cases an agent may not be able to immediately use or store the available energy due to its limited battery flow or capacity. We call this the *wasted* energy and denote it by $\boldsymbol{w} = (w_1, ..., w_t) \in \mathbb{R}^t_{\geq 0}$.

Using the battery an agent can compute an *energy allocation*, $\boldsymbol{p} = (p_1, ..., p_t) \in \mathbb{R}^t_{\geq 0}$, allocating the generated energy g to loads h. The utility of agent a at time i is then load p_i that is powered at time i. The overall utility u^a is given by:

$$u^a = \sum_{i=1}^{t} p_i \qquad (1)$$

Thus, the goal of an agent is to power as much of its load as possible to maximise its utility. The battery is useful here as it gives the agent flexibility in deciding when to store and when to use energy and thus, it enables the agent to find an optimal energy allocation, $\boldsymbol{p}^*$, given by:

$$\boldsymbol{p}^* = \underset{\boldsymbol{p}}{\operatorname{argmax}} \sum_{i=1}^{t} p_i \quad \forall \ i \in \boldsymbol{t} \qquad (2)$$

This can be transformed to a linear programming (LP) model with the following constraints:

Constraint 1: At time i, the allocated power p_i depends on the generated power g_i, charging c_i and discharging d_i:

$$p_i = g_i - c_i + d_i - w_i \quad \forall \ i \in \boldsymbol{t} \qquad (o_1)$$

Constraint 2: The current battery state q_i depends on the last battery state $q_{(i-1)}$, charge $c_{(i-1)}$ and discharge $d_{(i-1)}$. The charge flow $c_i \in \boldsymbol{c}$ is subjected to the battery efficiency e. Also, the initial battery state q_1 is zero.

$$q_i = \begin{cases} q_{(i-1)} + e \times c_{(i-1)} - d_{(i-1)} & if \ i > 1 \\ 0 & if \ i = 1 \end{cases} \qquad (o_2)$$

Constraint 3: Allocated power p_i must not exceed load h_i:

$$p_i \leq h_i \quad \forall \ p_i \in \boldsymbol{p}, \ h_i \in \boldsymbol{h} \qquad (o_3)$$

Having outlined the model of a single agent, we now discuss how agents can be connected to form a community.

3   Connecting Agents to Build a Community

Connecting two agents requires a physical link between them to enable them to (i) communicate and (ii) exchange energy. However, the absence of a centralised infrastructure (e.g., the electricity grid or telephone networks) in remote communities makes it challenging to connect homes. We envision that this challenge can be addressed by establishing a light-weight *peer-to-peer* network of homes where each home owns an *exchange box* that connects it to other homes (see Figure 1). Two homes are connected if there exists a direct physical link between their exchange boxes. A single exchange box can be connected to multiple exchange boxe, forming a network of interconnected agents from the ground-up without a centralised infrastructure. When an agent is connected, the power available to it also includes the flow on the links (in short, the *flow*) between it and the agents to which it is connected to. Let M be a set of agents connected to a and let $\boldsymbol{l}^{a \leftrightarrow j} = (l_1^{a \leftrightarrow j}, ..., l_t^{a \leftrightarrow j}) \in \mathbb{R}^t$ denote the agreed flow between a and some agent $j \in M$. Then the total flow f_i available to agent a at time period i is:

$$f_i = z \times \sum_{j=1}^{|M|} l_i^{a \leftrightarrow j} \quad \forall j \in M, \forall i \in \boldsymbol{t}$$

Here $0 \leq z \leq 1$ is the efficiency of the physical link.[2] We can modify constraint o_1 to include this power in the total power that is available to a at time period i as follows:

$$p_i = g_i - c_i + d_i - w_i + f_i \qquad \forall i \in \boldsymbol{t} \qquad (o_4)$$

For a given flow $\hat{\boldsymbol{f}} = (\hat{f}_1, ..., \hat{f}_t) \in \mathbb{R}^t$, a can maximise its utility by using Equation 1 and constraint o_4 as follows:

$$u^a(\hat{\boldsymbol{f}}) = \max \sum_{i=1}^{t} (g_i - c_i + d_i - w_i + \hat{f}_i) \quad \forall i \in \boldsymbol{t} \quad (3)$$

Where $u^a(\hat{\boldsymbol{f}})$ denotes the maximum utility that a can get for $\hat{\boldsymbol{f}}$, subjected to constraints $\{o_2, ..., o_4\}$. Similarly, when a needs to compute $\boldsymbol{f}^*$ that maximises its utility it can use:

$$\boldsymbol{f}^* = \underset{f_i \in \boldsymbol{f}}{\operatorname{argmax}} \sum_{i=1}^{t} (g_i - c_i + d_i - w_i + f_i) \quad \forall i \in \boldsymbol{t} \quad (4)$$

Now that an agent can compute its optimal flow and evaluate its utility for any offered flow, it can negotiate with other agents to reach an agreed flow that increases its utility. Here, the *increase in utility* comes from the fact that via exchange an agent can avoid energy storage losses and utilise energy that will be unused otherwise. Note that, if an agent has 100% efficient battery and infinite storage, it cannot increase its utility via exchange. The negotiation is challenging for agents because it involves interdependent issues and multiple agents. To faciliate negotiation in this context, we next present a protocol that reduces this complexity and enables agents to reach agreements efficiently.

4 Energy Exchange Protocol (EEP)

The core idea behind our energy exchange protocol (EEP) is to divide agents into two power pools that need energy at alternate times, and impose restrictions on the negotiation to reduce complexity. These restrictions are engineered so that the negotiation ends in outcomes with certain properties.

Before defining the EEP, we define our terminology. We consider exchange over finite time (e.g., a day) which can be divided into *exchange periods*. An exchange period is an atomic unit of time (e.g., 12 consecutive hours) for energy exchange and consists of at least one time period. The EEP allows only *two* exchange periods and divides agents into two *exchange types* as per the exchange period in which they require energy. The negotiation starts with *round zero* where agents declare their exchange types followed by *offer rounds* at *specified times*. In each *offer round*, only one exchange type is allowed to make *simultaneous* offers to all connected agents in order to reach an *agreed flow*. If a makes an offer to b, we denote this offer as $l^{a \rightarrow b}$. The EEP imposes restrictions (r_1, r_2, r_3, r_4) (Figure 2) on the offers made (called *valid flow (VF)* offers). The receiver can accept this offer or any VF part of it (r_5), and the outcome (i.e., agreed flow) is denoted by $l^{a \leftrightarrow b}$. Given these terms, Figure 2 describes the EEP in detail. In order to negotiate, an agent needs to know its desired outcomes and its strategy which we discuss next.

Computing Valid Link Flows

As discussed in Section 3, the utility u^a of an agent a depends on its total flow $\boldsymbol{f}$. Let S be the set of all flows, then a can find $\boldsymbol{f}^* \in S$ that maximises u^a, via Equation 4 (see Section 3). However, under the EEP only *valid* flows can be agreed (since agents can make or accept only VF offers, the agreed flow if any, is also a VF). In this sense, the EEP reduces S to the set of all *valid* flows $S_{VF} \subset S$ that meets the restrictions (r_1, r_2). To find $\boldsymbol{f}^* \in S_{VF}$, a can use Equation 4 subjected to r_1 and r_2 (in addition of $\{o_2, ..., o_4\}$). Knowing $\boldsymbol{f}^*$, a can easily infer its *exchange type* (i.e., which exchange period it prefers to receive energy in).

Here, we note that r_1 and r_2 are designed such that S_{VF} is a convex set where all members lie on the same geometric line. More specifically, if $\boldsymbol{f} = (f_1, f_2, f_3, f_4) \in S_{VF}$, then r_1 requires the sum of energy in both exchange periods to be equal (e.g., $f_1 + f_2 = -(f_3 + f_4)$) while r_2 says that $|f_1| = |f_2| = |f_3| = |f_4|$. Now, any scalar multiple of $\boldsymbol{f}$, i.e., $c \times \boldsymbol{f}$:
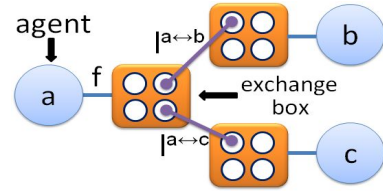
---

²For our experiments in Section 7, we have $z = 0.999$.

Figure 1: Agent a connected to b and c via exchange boxes.

$c \in \mathbb{R}$ also meets r_1 and r_2 and hence all scalar multiples of $\boldsymbol{f}$ are in S_{VF}. This also implies that if $\boldsymbol{f} \in S_{VF}$ then all $\boldsymbol{f}' \in S_{VF}$ can be described as $c \times \boldsymbol{f}$ (e.g., if $\boldsymbol{f} = (1, 1, -1, -1) \in S_{VF}$ then $\boldsymbol{f}' = (2, 2, -2, -2) \in S_{VF}$ can be expressed as $2 \times \boldsymbol{f}$. This geometric characteristic of S_{VF} ensures that if $\boldsymbol{f}^* \in S_{VF}$ maximises u^a, then $\boldsymbol{f}^*$ is unique and u^a is monotonically decreasing over $0 \leq \boldsymbol{f} \leq \boldsymbol{f}^*$ (see Lemma 1).

Making Valid Link Flow Offers

Having known $\boldsymbol{f}^* \in S_{VF}$ and its exchange type, a needs to know what VF offers to make. To reduce complexity at this stage, the EEP imposes r_3 which requires a to treat all agents (that it is making offers to) equally (see Figure 2). This reduces the strategy space for an agent, and together with other restrictions, entails an SPNE that we prove in Section 5.

Before we explain the properties of the EEP, we give an intuitive example to show how it will work in action:

Example 1. *Imagine that in a society of agents, the following are the already agreed on conventions:*

1. *Negotiation begins at 0200 hours every morning. Subsequent rounds take place every minute.*
2. *The total time of an exchange is 24 hours. The exchange starts at 0600 hours and ends at 0600 hours the next day.*
3. *This day is divided into two exchange periods, each consisting of 6 two-hours-long time periods.*
4. *Agents that need energy in the first exchange period are* exchange type 1 *and allowed to make offers in each round.*

Now, in a society of agents a, b and c, a finds that its optimal VF is $\boldsymbol{f}^a$=(4,4,4,4,4,4,-4,-4,-4,-4,-4,-4) and its exchange type is 1. Similarly, b and c find their exchange type to be 2 and their optimal VFs $\boldsymbol{f}^b$=(-1,-1,-1,-1,-1,-1,1,1,1,1,1,1) and $\boldsymbol{f}^c$=(-3,-3,-3,-3,-3,-3,3,3,3,3,3,3). At round zero, all agents declare their types simultaneously. At round 1, a (being exchange type 1) makes a VF offer $l^{a \rightarrow b}$=(2,2,2,2,2,2,-2,-2,-2,-2,-2,-2) to b and $l^{a \rightarrow c}$=(2,2,2,2,2,2,-2,-2,-2,-2,-2,-2) to c (Note: $l^{a \rightarrow b} + l^{a \rightarrow c} = \boldsymbol{f}^a$). Since $\boldsymbol{f}^b < l^{a \rightarrow b}$, b sends a PARTIAL ACCEPT message to a and the flow $l^{a \leftrightarrow b} = \boldsymbol{f}^b$ is agreed. While for c, $l^{a \rightarrow c} > \boldsymbol{f}^c$, it sends an ACCEPT message with FO=1(see Figure 2). In round 2, a makes a further offer of $l^{a \leftrightarrow c}$=(1,1,1,1,1,1,-1,-1,-1,-1,-1,-1) to c which it accepts and sends an ACCEPT message with FO=0. Thus, the overall exchange is agreed as per $l^{a \leftrightarrow (b,c)} = (\boldsymbol{f}^b, \boldsymbol{f}^c)$.

5 Properties of Our protocol

Here, we first show that the agents have a strategy profile in SPNE. We then discuss Pareto-optimality and scalability.

Subgame Perfect Nash Equilibrium: We model negotiation under the EEP as a *sequential* game where agents

make their moves in a well-defined sequence (e.g., declaring exchange types and then making offers) at specified times (i.e., rounds). A *subgame* is then a part or subset of this sequential game (e.g., some *offer rounds*). To show that a strategy profile is SPNE it is necessary to show that it represents an SPNE in *every* subgame of the original game. In the following, we first show that all agents have a best response (BR) for any given round (i.e., *round zero* or *offer round*) such that when all agents play their BR, it leads to a Nash equilibrium (NE) in that particular round. We then show that the strategy profile where all agents play their BR in all rounds is SPNE.

Theorem 1. *In* round zero*, all agents have a BR which is to declare their true exchange type.*

Proof. Only two exchange periods are allowed (r_1), thus an agent can be one of the two exchange types. This divides agents into two power pools that need energy in the alternate exchange periods. Agents in the same pool will not exchange energy between them as they prefer to receive energy in the same exchange period. If an agent misreports its exchange type, it will either make or receive offers from the agents in its own pool, hence no agreement will take place. Also, once declared, an agent cannot make offers that do not correspond to its prior declared exchange type (r_4). Thus, an agent's BR is to declare its true exchange type to negotiate with the opposite pool. Hence, the strategy profile where all agents declare their true exchange type is NE in *round zero*. □

Theorem 2. *In an* offer round*, an agent making offers to a set of agents has a BR which is to make an equal offer to each agent such that the sum of all its offers equals the VF that maximises its own utility.*

Proof. Let a be an agent with $\boldsymbol{f^a} \in S_{VF}$ that maximises its utility. Let X be the set of agents to which a would like to make offers and $\boldsymbol{l^{a \to j}}$ denote a VF offer to an agent $j \in X$. Under the EEP, j can either accept some or all of $\boldsymbol{l^{a \to j}}$ (r_5), depending on its optimal VF and the other offers it may have received. Let $\boldsymbol{l^{a \leftrightarrow j}}$ denote the agreed VF between a and j. Let $\boldsymbol{l^{a \to X}}$ be the sum of offers that a makes to each agent in X while $\boldsymbol{l^{a \leftrightarrow X}}$ be the sum of agreed VFs between a and all agents in X. Now, Table 1 lists all three possible strategies of a in offering $\boldsymbol{l^{a \to X}}$ and, assuming all agents in X play their BR (see Theorem 3), their outcomes. We note that offering $\boldsymbol{l^{a \to X}} = \boldsymbol{f^a}$ dominates $\boldsymbol{l^{a \leftrightarrow X}} < \boldsymbol{f^a}$ because it offers at least as much utility along with a potential outcome of $\boldsymbol{f^a}$. Note, when a offers $\boldsymbol{l^{a \to X}} = \boldsymbol{f^a}$ and the outcome is $\boldsymbol{l^{a \leftrightarrow X}} < \boldsymbol{f^a}$, then it can reach new agreements in further rounds by offering the remaining flow (i.e., $\boldsymbol{f^a} - \boldsymbol{l^{a \leftrightarrow X}}$). We also note that a can achieve $\boldsymbol{f^a}$ with $\boldsymbol{l^{a \to X}} > \boldsymbol{f^a}$. However, one potential outcome of this strategy is when $\boldsymbol{l^{a \leftrightarrow X}} > \boldsymbol{f^a}$ (i.e., a is now committed to a VF that exceeds its optimal VF) which not only means that a does not get its optimal VF $\boldsymbol{f^a}$ but also that the further negotiation rounds (thus agreeing to even greater flow) will not lead to $\boldsymbol{f^a}$. Therefore, $\boldsymbol{l^{a \to X}} = \boldsymbol{f^a}$ is BR for a such that it cannot do better by offering any other VF, provided all agents in X play their BR. Note, as r_3 mandates, a will make an equal offer to all (i.e., $\boldsymbol{l^{a \to j}} = \boldsymbol{l^{a \to (j+1)}} \ \forall j \in X$) so that $\boldsymbol{l^{a \to X}} = \boldsymbol{f^a}$. □

Energy Exchange Protocol (EEP)

1. Negotiation starts at a specified time with *round zero* where all agents declares their *exchange type* to connected agents.

2. Subsequent *offer rounds* take place at specified intervals. Only one exchange type is allowed to make offers in all *offer rounds*.

3. Agents make simultaneous *valid* flow offers to connected agents. An offer $\boldsymbol{l}$ is valid if it meets the following criteria.

   - The offer comprises of exactly two exchange periods. Each exchange period consists of an equal number of consecutive time periods. The amount of energy exchanged in each exchange period must be the same, with opposite flows. For example, if exchange starts at time period 1, then:

   $$\boldsymbol{l} = (l_1, ...l_n) \in S_{VF} : \sum_{i=1}^{n/2} l_i = - \sum_{i=n/2+1}^{n} l_i \qquad (r_1)$$

   - The amount of energy in each time period is equal.

   $$\boldsymbol{l} = (l_1, ...l_n) \in S_{VF} : \forall \, l_t \in \boldsymbol{l} : |l_t| = |l_{t+1}| \qquad (r_2)$$

   - If an agent makes offers to n agents in a round, then the amount of energy in each offer must be the same. $\qquad (r_3)$
   - The offer conforms to the agent's declared exchange type.(r_4)

4. On receiving an offer $\boldsymbol{l^{a \to b}}$ from a, an agent b has two options:.

   - Accept: b can accept $\boldsymbol{l^{a \to b}}$ and sends to a a two-part message. First part is the ACCEPT message to indicate that it accepts $\boldsymbol{l^{a \to b}}$. Second is a boolean signal FO (Further Offers) indicating if b is interested in receiving further offers from a. The agreed flow is $\boldsymbol{l^{a \leftrightarrow b}} = \boldsymbol{l^{a \to b}}$. The EEP round terminates.
   - Partial Accept: b can accept any partial VF $\boldsymbol{l} = y \times \boldsymbol{l^{a \to b}}$: $y \in \mathbb{R}|0 \leq y < 1$ (r_5). It can send a PARTIAL ACCEPT message with $\boldsymbol{l}$, indicating the partial acceptance of $\boldsymbol{l^{a \to b}}$. The agreed flow is $\boldsymbol{l^{a \leftrightarrow b}} = \boldsymbol{l}$. Note, $\boldsymbol{l} = 0 \times \boldsymbol{l^{a \to b}}$ is also a partial VF which b can choose to send to a to indicate that it does not wish to exchange. The EEP round terminates. (r_5)

Figure 2: The Energy Exchange Protocol

Theorem 3. *An agent considering the received offers from a set of agents has a BR which is to immediately accept offers in this round such that the sum of accepted offers is less than or equal to the VF that maximises its own utility.*

Proof. Let agent a with $\boldsymbol{f^a} \in S_{VF}$ that maximises u^a and $\boldsymbol{l^{X \to a}}$ be the sum of offers that it has received from agents in X. Now, if $\boldsymbol{f^a} \leq \boldsymbol{l^{X \to a}}$, then a can ACCEPT or PARTIAL ACCEPT offers such that $\boldsymbol{f^a} = \boldsymbol{l^{X \to a}}$, thus acquiring its optimal VF. However, if $\boldsymbol{f^a} > \boldsymbol{l^{X \to a}}$ then a can ACCEPT all offers in this round and then participate in further rounds to get its remaining VF. Note, an agent is never worse off by accepting offers immediately. On the contrary, if it delays the acceptance in a round then the other agents may reach agreements with each other in that round and thus in the further rounds the number of offers will reduce, reducing its chances of reaching agreements. Hence, the offer-accepting agent has a BR which is to accept offers immediately such that their sum is less than or equal to its optimal VF. □

Theorem 4. *The strategy profile where agents play their BR - as per Theorem 1,2 and 3 - is SPNE.*

Proof. We know that all agents have a BR in *round zero* which is to declare their true exchange type and the strategy profile where all agents declare their true exchange type is NE as no agent can do better by deviating individually (Theorem 1). We also know that in any *offer rounds*, each offer-making agent has a BR which is to make offers such that the sum of their offers equals its optimal VF (Theorem 2). Similarly, each offer-accepting agent has a BR which is to immediately accept the offers such that their sum is less than or equal to its optimal VF (Theorem 3). Note, no offer-making or offer-accepting agent can do better by individually deviating from its BR, and thus when all agents play their BR in an *offer round* it leads to a NE in that round. Now, consider the overall strategy profile (OSP) where all agents play their BR in *round zero* according to Theorem 1 as well as their BR in all *offer rounds* according to Theorem 2 (if they are making an offer) or Theorem 3 (if they are accepting an offer). Now, the OSP is SPNE for a sequential game because it defines a NE at every stage (i.e., round) of the game. Similarly, the OSP is SPNE for every subgame of the original game because any subgame consists an optional *round zero* and an optional number of *offer rounds*, and for any of these rounds the OSP defines a NE. $\qquad\square$

Pareto-optimal outcomes: Consider a, with the optimal VF f^a, connected to a set of agents X. Under EEP, its negotiation with other agents ends in two scenarios. First, where it agrees to VFs such that their sum equal to f^a in which any further change in the agreed VF will decrease its utility (f^a is unique, see Lemma 1) and thus the outcome is Pareto-optimal. Second, when the agreed VF is less than f^a but no other agent is willing to negotiate in further rounds (they already have reached their optimal VFs). In this case, though increasing the flow will improve u^a, the other agent(s) will no longer gain their maximum utilities; hence, Pareto-optimality ensues. Therefore, all outcomes under the EEP are Pareto-optimal in S_{VF}.

Tractability, Concurrency and Scalability: The EEP restrictions simplify negotiation such that it becomes tractable, concurrent and scalable. More specifically, r_1 and r_2 reduce S to S_{VF} where it becomes easier for an agent to compute the optimal VF (Equation 4) using an LP solver. This LP formation makes the computation tractable. If an agent knows its optimal VF and would like to make offers to some agents, r_3 dictates that it make equal offers to all. In this sense, r_3 makes it easier for an agent to negotiate with *multiple* agents *simultaneously* which ensures scalability and concurrency.

6 The Nash Bargaining Solution: Benchmark

The Nash bargaining solution (NBS) is a widely known axiomatic bargaining solution in cooperative bargaining that agents can use to find a common satisfying solution (Nash 1950). Its axioms define a unique solution that maximises the product of gains in utility of agents. In the context of energy exchange, computing the NBS involves finding the flow f for each agent so that the product of gains in their utility is maximised. Let $(d^1, ..., d^m) \in \mathbb{R}^m$ be the utilities that m agents obtain when they are disconnected (also called

Agent a offers	Potential $l^{a\leftrightarrow X}$	Utility $u^a(l^{a\leftrightarrow X})$
$l^{a\to X} < f^a$	$l^{a\leftrightarrow X} < f^a$	$u^a(l^{a\leftrightarrow X}) < u^a(f^a) \because l^{a\leftrightarrow X} < f^a$
$l^{a\to X} = f^a$	$l^{a\leftrightarrow X} < f^a$	$u^a(l^{a\leftrightarrow X}) < u^a(f^a) \because l^{a\leftrightarrow X} < f^a$
	$l^{a\leftrightarrow X} = f^a$	$u^a(l^{a\leftrightarrow X}) = u^a(f^a)$
$l^{a\to X} > f^a$	$l^{a\leftrightarrow X} < f^a$	$u^a(l^{a\leftrightarrow X}) < u^a(f^a) \because l^{a\leftrightarrow X} < f^a$
	$l^{a\leftrightarrow X} = f^a$	$u^a(l^{a\leftrightarrow X}) = u^a(f^a)$
	$l^{a\leftrightarrow X} > f^a$	$u^a(l^{a\leftrightarrow X}) < u^a(f^a)$ $\because \forall f \in S_{VF}: f \neq f^a \ u^a(f) < u^a(f^a)$

Table 1: Offering f^a is the best response for agent a.

disagreement utilities). These disagreement utilities are the maximum utilities that agents can get with no energy exchange. Let $F = (f^1, ..., f^m)$ denote the flows for all agents. Then, the NBS is the solution that maximises the following:

$$F_{NBS} = \operatorname*{argmax}_{(f^1, ..., f^m)} \prod_{j=1}^{m} \left[u^j(f^j) - d^j \right] \qquad (5)$$

When F is compact and convex then the solution F_{NBS} is unique (Nash 1953) and computing the optimal solution is straight forward using convex optimisation. However, interdependency between issues gives rise to a non-convex solution set with multiple NBS (Fujita et al. 2010) whereby stochastic optimisation techniques are needed. Note that, the NBS only defines *which* solutions (in the set of all solutions) meet the defined axioms and not *how* agents can reach such agreements. However, we can use it as the *theoretical* upper bound to evaluate the EEP as we discuss in the next section.

7 Empirical Evaluation

Having outlined the benchmark, we now set-up a realistic example to demonstrate the practical applicability of the EEP. To this end, we consider an example of energy exchange in a community of 20 agents where each agent has either a 1.5kW wind turbine or a 1.75kW solar panel with equal probability. The energy generation data for the wind turbine comes from a wind farm near Lugo, Northwest Spain (www.sotaventogalicia.com), while the output of the solar panel is estimated to be directly proportional to the daily radiance for the same region (www.re.jrc.ec.europa.eu/apps/radday.php). We use data for July 2011, estimate the average generation for a day and scale it to match the output of a 1.5kW wind turbine and a 1.75kW solar panel. The load requirements of homes in remote areas are not available so we use load data, recorded and provided by a UK electric company in low-income homes equipped with smart meters. Figure 3 shows this consumption along with the generation (solar and wind). The actual generation and consumption for each agent comes from a distribution over these profiles. More specifically, we model generation/consumption in each time unit as an independent Gaussian distribution (with scaled value as the mean and the variance within 10% of it). We assume that agents have identical batteries [$s = 20$kWh, $c = 4$kWh, $d = -4$kWh, $e = 90\%$]. Given these profiles, agents can compute their utilities without exchange using an LP solver.
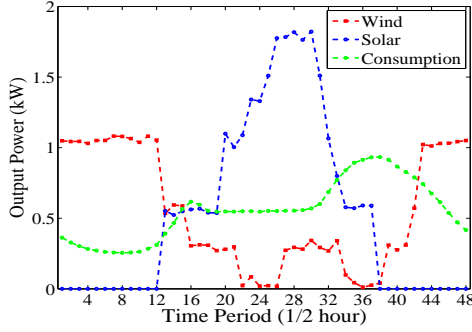
Figure 3: Avg. consumption and generation.

Figure 4: Social Welfare - EEP Vs NBS

Given this setup, our main objectives are to show that via energy exchange using the EEP, agents can (i) reduce their battery charging needs (ii) increase their utility even as the battery efficiency decreases. We simulate a community of 20 agents and calculate the total battery charging with (i) no exchange and when they exchange energy via (ii) the EEP and (iii) the NBS. We repeat this simulation 20 times and find that, on average, when agents do not exchange energy their overall battery usage is 222.7kWh while with energy exchange via the NBS their overall battery usage is 113.6kWh (i.e., 49% reduction compared to no exchange) and with the EEP their usage is 141kWh (37% reduction). This is important because electric batteries are expensive (costing as much as 500 USD/kWh) and have a limited number of charging cycles (3000 to 5000). Reducing the battery charging prolongs the battery life and reduces the need for frequent replacements and thus savings in maintenance cost.

The above reduction in battery charging via exchange becomes more useful as the battery efficiency deteriorates with time and usage. To show this, we simulate a community of 20 agents and calculate its social welfare function (sum of all agents' utilities) as the battery efficiency of *all* agents is reduced (other parameters remain unchanged). With a less efficient battery, the storage losses increases and therefore agents' utilities (without exchange) decrease. However, by exchanging energy such losses can be avoided. Indeed, Figure 4 shows that the exchange becomes increasingly useful as the battery efficiency reduces. In particular, we note that, on average, the NBS improves the social welfare by 16.5% while the EEP does so by 13.7%.

8 Conclusion and Future Work

The problem of negotiation over energy exchange is a complex interdependent multi-issue negotiation problem. Here, we presented a negotiation protocol, the EEP, which tackles this complexity by imposing certain restrictions over offers, as so that agents have a strategy profile in SPNE and negotiation is concurrent, scalable and entails Pareto-optimal outcomes (within allowed agreements). Using real-world data, we empirically evaluate the EEP and benchmark it against the NBS. Our results show that energy exchange via the EEP is useful in communities to improve the efficient use of energy and storage. Future work will investigate how relaxing the EEP restrictions affects the negotiated outcomes when the energy generation is uncertain and loads are deferrable.
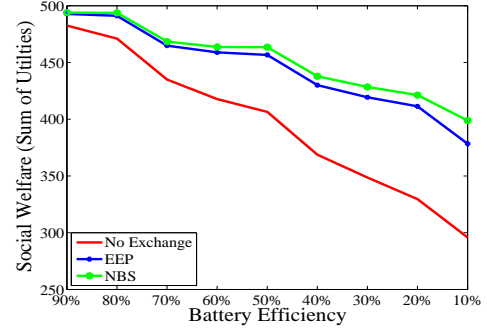
A Appendix

Lemma 1. *If $\boldsymbol{f} = (f_1, ..., f_n) \in S_{VF}$ maximises u^a then $\boldsymbol{f}$ is unique and u^a is a strictly monotonically decreasing function over $\forall \boldsymbol{f'} \in S_{VF} : (0, ..., 0) < \boldsymbol{f'} < \boldsymbol{f}$ range.*

Proof. Let a be an agent with the optimal VF $\boldsymbol{f} = (f_1, f_2, f_3, f_4) \in S_{VF}$ and with exchange periods ex_1 comprises of time period 1 and 2, and ex_2 comprises of time period 3 and 4. Let's assume that a prefers to receive energy in ex_1 and transfer in ex_2. Note that r_1 requires agents to receive and transfer the same amount in ex_1 and ex_2, i.e., $f_1 + f_2 = -(f_3 + f_4)$. The *gain* in the utility for a comes merely because a can utilise *more* of the same amount of energy in ex_1 than ex_2 (due to lack of demand in ex_2 or storage losses). This implies that for any VF $\boldsymbol{f'} \in S_{VF}$ to improve u^a, it must be that $\alpha_1|f_1'| + \alpha_2|f_2'| > \alpha_3|f_3'| + \alpha_4|f_4'|$ or $\alpha_1|f_1'| + \alpha_2|f_2'| - \alpha_3|f_3'| - \alpha_4|f_4'| > 0$ where $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \alpha_3, \alpha_4) : \forall \alpha_i \in \boldsymbol{\alpha}\ 0 < \alpha_i \leq 1$. Now, the greater the inequality, the greater the amount of energy *saved* that would otherwise be unused. Indeed, Equation 4 attempts to maximise this inequality and find the optimal VF for maximum increase in the utility (see Section 4).

Now, if $\boldsymbol{f} = (f_1, f_2, f_3, f_4) \in S_{VF}$ is the optimal VF then the maximum inequality is $\alpha_1|f_1| + \alpha_2|f_2| > \alpha_3|f_3| + \alpha_3|f_4|$. Let $0 < \beta_2 < \beta_1 \leq 1$. Since $\beta_2 < \beta_1$, the following will hold:

$$\Rightarrow \quad \beta_2(\alpha_1|f_1| + \alpha_2|f_2| - \alpha_3|f_3| - \alpha_4|f_4|) <$$
$$\beta_1(\alpha_1|f_1| + \alpha_2|f_2| - \alpha_3|f_3| - \alpha_4|f_4|)$$
$$\Rightarrow \quad \alpha_1|\beta_2 f_1| + \alpha_2|\beta_2 f_2| - \alpha_3|\beta_2 f_3| - \alpha_4|\beta_2 f_4|) <$$
$$\alpha_1|\beta_1 f_1| + \alpha_2|\beta_1 f_2| - \alpha_3|\beta_1 f_3| - \alpha_4|\beta_1 f_4|$$

Here, $\beta_1 f_i$ and $\beta_2 f_i$ are just the scalar multiple of $\boldsymbol{f}$ (see Section 4 on the geometric properties of S_{VF}). Let $\boldsymbol{f'} = \beta_1 \boldsymbol{f}$ and $\boldsymbol{f''} = \beta_2 \boldsymbol{f}$. Note, $0 < \beta_2 < \beta_1 \leq 1$ implies $0 < \boldsymbol{f''} < \boldsymbol{f'} < \boldsymbol{f}$. Therefore:

$$\Rightarrow \quad \alpha_1|f_1''| + \alpha_2|f_2''| - \alpha_3|f_3''| - \alpha_4|f_4''| <$$
$$\alpha_1|f_1'| + \alpha_2|f_2'| - \alpha_3|f_3'| - \alpha_4|f_4'| \qquad (i)$$

This shows that the amount of energy that can be saved (that would otherwise be lost) due to flow $\boldsymbol{f''}$ (i.e., $\alpha_1|f_1''| + \alpha_2|f_2''| - \alpha_3|f_3''| - \alpha_4|f_4''|$) is *less* than that of due to $\boldsymbol{f'}$ for all $0 < \boldsymbol{f''} < \boldsymbol{f'} < \boldsymbol{f}$. We know that the utility of an agent is the total amount of energy used (Equation 1). Therefore, if Inequality (i) holds, then $u(\boldsymbol{f''}) < u(\boldsymbol{f'})$ also hold for all $0 < \boldsymbol{f''} < \boldsymbol{f'} < \boldsymbol{f}$.

We know that if $\boldsymbol{f} \in S_{VF}$, then $\forall \boldsymbol{f'} \in S_{VF} : \boldsymbol{f} \neq \boldsymbol{f'}\ \boldsymbol{f} = c \times \boldsymbol{f'}$ where $c \in \mathbb{R} | 0 < c < 1$ (see Section 4), therefore, the amount of energy that is transferred by each VF (i.e., $f_1 + f_2$ or $f_3 + f_4$) is unique. Thus, when an agent needs to exchange a certain amount of energy (for optimal utility), there is only one corresponding VF in S_{VF}. Hence, optimal VF $\boldsymbol{f}$ is unique for an agent. $\qquad \square$

References

Alam, M.; Ramchurn, S. D.; and Rogers, A. 2013. Cooperative energy exchange for the efficient use of energy and resources in remote communities. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Alam, M.; Rogers, A.; and Ramchurn, S. D. 2011. A negotiation protocol for multiple interdependent issues negotiation over energy exchange. In *Proceedings of the AI for an Intelligent Planet*, AIIP Workshop, 1:1–1:6. New York, USA: ACM.

Fujita, K.; Ito, T.; and Klein, M. 2010. Representative-based protocol for multiple interdependent issue negotiation problems. *InTech* 347–363.

Hattori, H.; Klein, M.; and Ito, T. 2007. A multi-phase protocol for negotiation with interdependent issues. In *Proceedings of IEEE/WIC/ACM International Conference on Intelligent Agent Technology.*, 153 –159.

Hindriks, K.; Jonker, C.; and Tykhonov, D. 2006. Reducing complexity of an agents utility space for negotiating interdependent issues. *InterJournal for Complex Systems*.

IEA. 2010. World energy outlook. Technical report, International Energy Agency.

Ito, T.; Hattori, H.; and Klein, M. 2007. Multi-issue negotiation protocol for agents: exploring nonlinear utility spaces. In *Proceedings of the 20th International Joint Conference on Artifical Intelligence*, 1347–1352.

Nash, J. 1950. The bargaining problem. *Econometrica* 18(2):155–162.

Nash, J. A. 1953. Two-person cooperative games. *Econometrica* 21(1):128–140.

Rosenschein, J. S., and Zlotkin, G. 1994. Rules of encounter: designing conventions for automated negotiation among computers. In *Rules of encounter*. Cambridge, MA, USA: MIT Press.

United Nations. 2012. What will it take to acheive the millennium development goals. Technical report, United Nations Development Programme.

Vytelingum, P.; Voice, T. D.; Ramchurn, S. D.; Rogers, A.; and Jennings, N. R. 2011. Theoretical and practical foundations of large-scale agent-based micro-storage in the smart grid. *J. Artif. Int. Res.* 42(1):765–813.