2013 IREP Symposium-Bulk Power System Dynamics and Control -IX (IREP), August 25-30, 2013, Rethymnon, Greece

A Healthy Dose of Reality for Game-Theoretic Approaches to Residential Demand Response

Archie C. Chapman, Gregor Verbič and David J. Hill
School of Electrical and Information Engineering, University of Sydney, Australia
Email: {archie.chapman,gregor.verbic,david.hill}@sydney.edu.au

Abstract

This paper addresses the assumptions underpinning many control schemes for *residential demand response* (RDR), with particular focus on those that adopt the framework of non–cooperative games. We propose four principal assumptions that we believe are necessary to give a realistic grounding to research on RDR, so that they might be more readily applied to the real problems faced by aggregators and households in a future energy network. These are that: (i) The energy use levels of households do not take continuous values, they take discrete and hybrid values; (ii) In addition to the system state variables, each household has a private state, representing the states of the goals it addresses in consuming electricity; (iii) Households have private preferences that are state-based, and therefore non–convex and combinatorial, and moreover, the monetary costs imposed by the system operator represents only part of their preferences for electrical energy use; and (iv) Household behaviour is strategic, both at the level of equilibrium analysis and algorithmic design. For each assumption we argue why it is necessary that a RDR scheme satisfy it, and illustrate the effects of violating our proposed assumption, with reference the existing literature on RDR schemes. We also provide several examples of techniques that satisfy each assumption, and illustrate our assumptions by developing a model that satisfies all four.

Introduction

Demand response refers to methods for controlling the amount of energy used by end users of electrical power. Demand response schemes are employed to provide additional capacity to the electricity network without costly new infrastructure, and to facilitate greater penetration of renewable generation, as increasingly flexible energy use is able to better track the intermittent supply provided by wind and solar generation. In particular, this paper focuses on *residential demand response* (RDR). These are demand response schemes that are constructed specifically for large numbers of households spread across a distribution network. Typically, these schemes are built on a framework composed of an RDR *aggregator* that coordinates, schedules or otherwise controls part of participating households' loads.

Given this context, the aim of this paper is to characterise some necessary assumptions on the physical and economic interaction between households and aggregators participating in a RDR scheme. We adopt a non–cooperative game approach to this domain [1], and present a critical survey of existing proposed RDR schemes and frameworks, with an outline of a set of cohesive measures for their improvement. Specifically, we propose four principal assumptions that we believe are necessary to give a realistic grounding to research on RDR, so that it might be more readily applied to the real problems posed by a future energy network, to aggregators and households alike.

In more detail, there are four aspects of the RDR domains where the simplifying assumptions that are often applied in RDR schemes are not justified. The first regards the nature of the household's control variables. In particular, these are often assumed to be continuous [2], [3], [4], [5], but, in reality, the control variables within a household are typically discrete or hybrid. Continuity allows powerful analysis and continuous optimisation methods to be applied to the problems, which is a useful approach to the analysis of stylised models of aggregate household behaviour. However, these techniques will be rendered useless in actual future grid deployments if the variable domains to which they are applied do not conform with their requirements.

Second, household state variables are typically overlooked or ignored in RDR schemes. For example, it is often assumed that the only relevant state variables are system state variables (i.e. power generation cost parameters, wind power generation levels and forecasts, etc; e.q. [4], [6], [7]). This is tantamount to implicitly assuming that households are static, and do not adjust their levels of demand for electricity in response to previous and future scheduled allocations.

Third, typical assumptions about household preferences simplify the models to a point where they actually obscure the real problems in RDR. For example, preferences for power use are commonly left to be completely represented by the the monetary costs imposed by the power system operator (albeit, often with hard constraints around what times electricity can be used [2], [3], [8], [5]). Moreover, in the cases where household utility models are used, these are often taken from either industrial load preference models or from aggregate models (as in [9]), which in both cases are typically concave [7], [10]. However, in either case, the household models used miss the considerable diversity in usage patterns and the combinatorial structure to their demand, because they are either obscured by aggregation or simply not present in most industrial loads.

978-1-4799-0199-9/13/ $31.00 ©2013 IEEE

Fourth, the nature of households' behaviour is often assumed to be non–strategic, or cooperative. This is particularly the case with respect to the *algorithms* proposed to RDR schemes; although non-cooperative game methods are employed to analyse the equilibria of the proposed RDR models, game-theoretic reasoning is not extended to the design of algorithms [2], [3], [4], [5], [8]. Thus, issues of *incentive compatibility* and *truthful implementation* are ignored [11].

The types of assumptions above disconnect the proposed methods from reality and make it difficult to see how they can be directly applied to real-world RDR domains.

Contribution of the Paper

Against this background, we argue for a non–cooperative game–theoretic approach to RDR that is built on firm principles and valid assumptions about the RDR domain. To be perfectly clear, we are not arguing against a game–theoretic approach. Rather, our contention is that such an approach is necessary, but that the abstractions listed above do not advance our understanding of RDR scenarios, and need to be replaced with assumptions that better reflect reality. Thus, the main contribution of the paper is to construct assumptions for the four aspects of the RDR domain listed above that are realistic and appropriate. Our proposed list of necessary assumptions are that:

Household energy use values — The electrical energy use levels of households is an aggregation of continuous, discrete and hybrid variable types, so are not necessarily drawn from a continuous set of values;

Household state variables — A household has a private state that collectively represents the states of the goals it pursues or tasks it wishes to complete, thereby consuming electrical energy;

Household preferences — Households have private preferences that are state-based, and therefore non–convex and combinatorial; moreover, the monetary costs imposed by the system operator represents only part of their preferences for electricity use; and

Household behaviour — Household behaviour is strategic, both at the level of equilibrium analysis and algorithmic implementation.

In this paper, for each of these four assumptions, we provide a justification for the proposed assumption, show why the existing assumptions are not justified and need replacing, and discuss ways to analyse models and compute solutions in scenarios that use these new assumptions. We believe this will better enable the transfer of smart grid research to real smart grid deployments than those assumptions that are currently regularly invoked. Note that the assumptions here are defined loosely, and as such, numerous different representations or formulations could accommodate them; we are not prescribing or proscribing any particular approach at a detailed level.

A final point to note: our analysis is based on an assumption that the interaction between households and an aggregator is non–cooperative. An alternate point of view is to model the interaction with a *cooperative* or *coalitional game*, in which a household forms a binding agreement with the aggregator to undertake some RDR activities [12]. It is our contention that these types of games are best suited to direct load control of the households appliances (e.g. through broadcast signals), rather than the types of schemes listed above where households are always given the discretion to choose their level of participation in the demand response activity.

Outline of the Paper

The paper progresses with a Preliminaries section that introduces notation and some important concepts. Then, each of the four assumptions above are discussed in each successive section: Household Energy Use Values; Household State Variables; Household Preferences; and Household Behaviour. Each section includes some examples of models and methods that satisfy the assumption, an examination of the effects of violating the proposed assumption, and the consequences of the assumption for some recently proposed RDR schemes. In addition, as we progressive, we develop an example of a model that admits our assumptions, which we use to illustrate our arguments. We also use this model to demonstrate how ignoring our assumptions can be detrimental to alternative RDR schemes; by giving examples of proposed RDR schemes that ignore one or more of the four features listed above. We conclude the paper with a brief discussion of directions for future work and some broader issues for game-theoretic methods in RDR. Since the model we develop is spread across several sections, a table of notation is included at the end of the paper for reference.

Preliminaries

Throughout, the set of (positive) real numbers is denoted ($\mathbb{R}_+$) $\mathbb{R}$, and the probability of an event is denoted $\mathbb{P}(\cdot)$. We adopt a discrete-time model, where operations are divided into $\mathcal{H} = \{0,\dots,h,\dots,H-1\}$ time slots over the decision horizon. Consequently, all electrical quantities are stated as blocks of energy; for example a 100W appliance running for 30 minutes is described by a 0.050kWh demand block.

We consider the interaction of one residential load aggregator (e.g. a retailer) with its associated residential loads. The interaction of the aggregator and households with the broader energy system and market are illustrated in Fig. 1.

Generation and Distribution

Regarding the transmission and distribution of electrical energy, there are considerable network effects to consider in

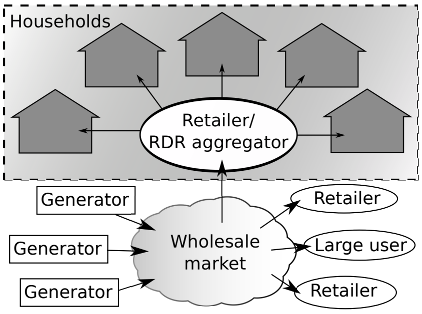

Fig. 1. Model of interaction between an RDR aggregator and households (the focus of this paper, bound by dashed lines), and the broader energy market.

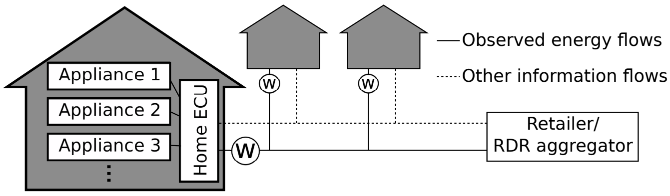

Fig. 2. Detailed model of interaction between an RDR aggregator and households: control of appliances remains completely under the control of the household, via the Home ECU; energy flows are observable to the aggregator, via the metering infrastructure (solid lines); other information flow are also facilitated, via a communications infrastructure (dashed lines).

future demand-response scenarios. However, as indicated by Fig. 1, we focus solely on demand-side modelling, so adopt a copper-plate network model.

We assume that the aggregator purchases energy in a wholesale electricity market comprising energy suppliers and other bulk purchasers (including some large end users), and also assume that the wholesale market that the aggregator buys from operates efficiently. Alternatively, we could make the same assumptions in the context of a microgrid with few generators, while still retaining our focus on the interaction of the residential aggregator and its residential loads.

Let $\mathcal{J}$ denote the set of electricity *suppliers*. Each energy provider $j \in \mathcal{J}$ supplies s_h^j units of energy in a slot h. Associated with each provider is a generation *cost function* $c^j : \mathbb{R}_+ \to \mathbb{R}_+$. We assume that, for all $j \in \mathcal{J}$, the cost functions $c^j(\cdot)$ are twice differentiable in $(0,\infty)$, strictly increasing and convex. For example, a typical conventional generator can be approximated by a *quadratic* cost function. Under the efficient market and cost function assumptions above, the price schedule faced by an aggregator buying energy in the wholesale market is also strictly increasing and convex.

Aggregator and Household Models

Our model of the interaction between an RDR aggregator and households, as described in Fig. 2. In this model, control over a household's load, via its appliances, remains completely under the control of the household, possibly intermediated by an *Home Electrical Control Unit* (ECU). In this model, energy flows can be seen by the aggregator using "smart meter" infrastructure. A two-way communications infrastructure allows for other information flows between the households and the aggregator, which will be used to facilitate and coordinate scheduling and pricing schemes in advance of the actual time that electrical energy is used. However, it should be noted that the information passed on these channels regarding predicted future energy flows is only verifiable at the time that the electricity is actually used.

Let I denote the set of *households* supplied by the aggregator. A household $i \in I$ uses $d_h^i \in \mathbb{R}_+$ units of electrical energy in slot h. Then, the total electrical energy supplied by the aggregator in slot h is given by: $x_h = \sum_{i\in I} d_h^i$. Household electricity use is associated with utility benefit or *reward*, $r_h^i \in \mathbb{R}_+$. This value can be thought of as indicating the household's "willingness to pay," or the user's fictitious monetary value derived from, consuming a certain amount of electrical energy. It is typically given a functional form, which is possibly also dependent on factors such as the particular time slot in which the electricity is used, or the timing of other energy use. However, since the composition of this function is a topic of contention in this paper, we leave a precise specification of a *reward function* $r_h^i(\cdot)$ to later sections.

Given a total amount of electrical energy supplied to all households by an aggregator, x_h, denote the corresponding price schedule that the aggregator faces by $C_h(x_h)$ (i.e. ignoring other market participants). The aim of an RDR scheme is to derive the method by which an aggregator structures its interaction with households. That is, an RDR scheme defines how the aggregator divides the costs it faces among its end users, and induces them to use electricity at the most appropriate times. Let $\mathbf{d}^i = [d_0^i, \ldots, d_{H-1}^i]$ be the vector of household i's electrical energy use over the horizon, and $\mathbf{d} = \{\mathbf{d}^i\}_{i\in I}$ be the collection of vectors of all households electrical energy use over the horizon. In general, we can define the cost division used by the aggregator as a vector function, $\phi(C_h(x_h), \mathbf{d})$, which returns a vector of costs, one for each end-user household (the aggregator calculates x_h from $\mathbf{d}$).

Each household has instantaneous *utility function* $u_h^i(\mathbf{d}, \cdot)$, which takes a quasi–linear form that combines their rewards and costs for consuming electricity. Moreover, this can be summed over the the decision horizon, giving a function for total utility:

$$U^i(\mathbf{d}, \cdot) = \sum_{h=0}^{H-1} u_h^i(\mathbf{d}, \cdot) = \sum_{h=0}^{H-1} r_h^i(\cdot) - \phi^i(C_h(x_h), \mathbf{d}) \qquad (1)$$

where $\phi^i(C_h(x_h), \mathbf{d})$ is the i^{th} component of the aggregator's cost division function. As in the value function, $(\cdot)$ represent a yet-to-be defined additional factors, the discussion of which make up part of the contribution of the paper.

The aggregator and households' actions are coupled through the dependence of their utilities on the vector of total loads $\mathbf{x}$. Thus, their interaction results in a game. The solution concept typically applied to non–cooperative games is the

Nash equilibrium solution. Let us write $\mathbf{d} = \{\mathbf{d}^i, \mathbf{d}^{-i}\}$, where $\mathbf{d}^{-i}$ denotes the loads of all households apart from i.

Definition 1: A joint electricity use profile, $\bar{\mathbf{d}}$, is a *pure strategy Nash equilibrium* if it satisfies:

$$U^i(\bar{\mathbf{d}}^i, \bar{\mathbf{d}}^{-i}, \cdot) - U^i(\mathbf{d}^i, \bar{\mathbf{d}}^{-i}, \cdot) \geq 0 \quad \forall \; i \in N.$$

That is, in a (pure strategy) Nash equilibrium, every households electricity use profile is a *best response* to the other households' usage profiles. In the remainder of the paper, when we use the term Nash equilibrium we are referring to a pure strategy Nash equilibrium. In order to retain our focus on demand–side modelling, we do not consider other solution concepts, such as mixed strategy Nash equilibrium or correlated equilibrium.

Each of the next four sections covers a specific proposed assumption. Each section has the following structure. First, the new assumption is described and justified. Then the implication of the proposed assumption on existing models and methods is explained, that is, we show how each assumption invalidates some proposed methods of analysis and computational procedures used to solve existing models. We then discuss how the difficulties of implementing the new household model assumptions can be overcome, and indicate some specific methods and representations that admit the new assumption. Finally, as we move through the four assumptions, we progressively develop an example of a model that satisfies all of them, illustrating their feasibility.

Household Energy Use Values

The first assumption we propose relates to the nature of control variables available to households. We start by noting the reality that households use power by operating appliances. Moreover, that household appliances have discrete, continuous and hybrid levels of power usage. For example, appliances with discrete power levels include an electric stove with discrete operating points, while others have fixed profile or patterns of power use, such as a refrigerator's compressor cycle, or a washing machine with a certain number of programs. Moreover, even in our discrete-time model, this means that in a fixed block of time only discrete amounts of electrical energy can be used.

Assumption 1: The electrical energy use levels of households is an aggregation of continuous, discrete and hybrid variable types, so are not necessarily drawn from a continuous set of values.

That is, we propose that it is necessary to assume that a household's choice of power level at different times, or energy requirements over a particular interval, is an aggregation of continuous, discrete and hybrid variable types, which may not allow for continuous choice of energy (or power) use. Our assumption stands in contrast to the common approach of treating the household decision on energy use as a continuous variable. When considering large industrial power users, assuming continuity is justifiable, due to the large nature of the loads in question. However, it is unreasonable in RDR, as most appliance loads cannot be arbitrarily split, due to the nature of the appliances' operation.

In order to make Assumption 1 concrete, we need to formalise some concepts. Let:

- $\mathcal{M}^i$ be the set of household i's appliances;
- $A^i = \underset{m \in \mathcal{M}^i}{\times} A_h^{i,m}$ be the set of actions that a household can undertake using its set of appliances, where $A^{i,m}$ is the set of action that can be undertaken by appliance $m \in \mathcal{M}^i$; and
- $d_h^{i,m}$ denote the load of appliance $m \in \mathcal{M}^i$ in slot h.

Now, having adopted a discrete-time modelling framework, we define loads as blocks of energy used in a time-slot. We make the not unrealistic assumption that appliances actions are completed within one time-slot (and this usually would be satisfied in the case that each slot is one hour long).[1] If appliance m can complete an action by operating at one of a set of discrete power levels, then the amount of energy it can use in a slot is one of a finite set of values; that is: $d^{i,m} \in \{a, b, \ldots\}$, where $a, b, \ldots \in \mathbb{R}_+$. If the appliance can complete a task by operating at a continuously variable level of power, then the amount of energy it uses over a slot also varies continuously; that is: $d^{i,m} \in [a, b]$, with $a, b \in \mathbb{R}_+$. Hybrid or mixed levels of power produce a level of electrical energy use over a slot that is between these two cases.

Household demand is necessarily an aggregation of these loads, and is given by:

$$d_h^i = \sum_{m \in \mathcal{M}^i} d_h^{i,m}.$$

Importantly, it does not follow that d_h^i can be represented by a continuously variable block of energy, because of the constraint that for some appliances, $d_h^{i,m}$ is drawn from a set of discrete values.[2]

We now discuss the effects of moving away from assuming continuous energy use variables on the analysis of RDR models and computation in RDR schemes.

Consequences for Analysis

When all decision variables are continuous, the analysis of a RDR scheme is often simple. For example, some early models of RDR adopted a convex game framework [2]. Since convex games have a unique Nash equilibrium [13], analysis of the solution is straightforward: it is equivalent to finding the

[1] For shorter slot lengths or more slots over the decision horizon, we would need to introduce additional variables to capture the couplings between energy used over the duration of an appliance's operation, which would distract from the main points of the paper.

[2] We note that households have flexibility in the timing of their energy use, and that they can use this flexibility to make the value of their energy use in a particular time slot continuous. However, taking these types of actions induces additional combinatorial structure to a household's energy use preferences.

minimum of a convex function on a convex feasible solution space.

More recently, several authors have adopted a potential game (or congestion game) framework for analysing energy use scheduling and load balancing schemes [4], [3]. In the case of *infinite potential games* with continuous actions, these models are easily analysed: By construction these games have a Nash equilibrium, since all potential games are guaranteed to have pure strategy Nash equilibria (see [14] for static games and [15], [16] for dynamic games). In contrast, in [5], I is taken to be an infinitely-large set of households. Thus, the effect of each household is infinitesimally small, and the aggregate joint action can be treated as continuous, to the same effect. Moreover, under certain convexity restrictions on the potential function, uniqueness of a Nash equilibrium can be guaranteed in games with continuous actions (e.g. [17]), however, these conditions are not always satisfied by RDR schemes.

In addition, simple tools from calculus can be used to bound the efficiency of these equilibria in potential games. This involves deriving a lower bound on the ratio of the worst–case Nash equilibrium to the "social optimum" of the game, which may or may not be an equilibrium itself.[3]

However, if the levels of demand are broken up in different ways, such that the games no longer have continuous or convex action spaces, then these methods of analysis cannot be applied. For example, even if every household's levels of electricity use can be expressed as a multiple of some arbitrarily small value $\eta > 0$ (so-called *integer-unit* congestion games), guarantees on the existence of pure strategy Nash equilibria can only be made with the addition of convexity conditions on the households utility functions (see [18], [19]).

Consequences for Computation

When all decision variables are continuous, not only is the analysis of a RDR scheme usually simplified, but so is the computation of a solution, because standard continuous (typically convex) optimisation methods can be used for solving the problems at hand. However, once variables that take discrete values are included, computation is made more difficult.

For example, convex games and infinite potential games can be solved using straightforward search techniques, such as gradient descent or binary search. In the case of convex games, these methods will find the unique equilibrium, and furthermore, they are usually easy to distribute across the households. For potential games, because of the possibility of multiple Nash equilibria, some researchers have proposed centralised solvers that select a specific, optimal Nash equilibrium [4]. Beyond this, other models that are neither convex or potential game–based also rely on continuous action spaces, for their proposed algorithms to operate [6]. However, in all of these cases, moving to discrete variables adds more complications depending on the precise nature of the variables' domains. For example, solving integer–unit congestion games involves a carefully designed search–based procedure that avoids falling into cycles [19]), which only exist as an effect of the discrete action space.

[3]For an overview of methods for analysing the efficiency of equilibria, see Section III of [11].

Household State Variables

Building on the previous assumption, the nature of household demand for appliance use is that it is *goal–* or *task–based*. That is, energy is an *intermediate good*: households do not gain utility for consuming electricity itself, rather by completing tasks using electrical energy. This lends itself to a representation of the state of a household in terms of the states of the goals associated with each appliance. These states could comprise tasks completed by particular appliances, such as clothes or dish washing, or goals for appliance states, such as heating the water in a hot water water systems to a sufficient level or ensuring that a refrigerator stays within a certain temperature range, and so on. We call the states of these goals, collectively, the *household state*.

Assumption 2: A household has a private state that collectively represents the states of the goals it pursues or tasks it wishes to complete, thereby consuming electrical energy.

We can formalise this assumption by requiring that RDR model includes a vector, s^i, of household state variables:

$$s^i \in S^i = \bigtimes_{m \in M^i} S^{i,m} \times S^{i,0},$$

where the space of values s^i can take comprises:

- the set of states of each appliance, $S^{i,m}$, and
- an appliance–independent household state, $S^{i,0}$.

For each appliance, the set $S^{i,m}$, is comprised of a set of states associated with the tasks or goals that the household wishes to complete or achieve using that appliance, as well as any intermediate or precursor states to these. In addition, the appliance–independent household state, $S^{i,0}$, could include the relevant household characteristics that are not affected by any appliance's actions, such as occupancy, outdoor temperature, and so on. This is very similar to the automatic systems used by thermostatically–controlled loads, such a refrigerators and stored electric hot water systems, which operate using dead–band controllers (e.g. "heating", "maximum temperature","cooling" and "minimum temperature").

An example of a *state-action* diagram for a dishwasher is given in Fig. 3. In this, the dishwasher has states and actions listed in Table I. Although this diagram, and subsequent ones, may contain more detail than is required for all RDR schemes, they do demonstrate one way that household states could be represented.

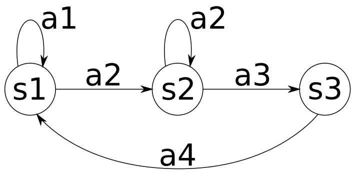

Fig. 3. A state-action diagram for an appliance.

State/Action	Description
s1	Not ready to start (items washed/empty)
s2	Ready to start (sufficiently full of dirty items)
s3	Dishwasher running
a1	No dirty items collected
a2	Collected dirty items
a3	Start dishwasher
a4	Wait to complete

TABLE I
EXAMPLE STATES AND ACTIONS FOR A DISHWASHER, CORRESPONDING THE STATE-ACTION DIAGRAM IN FIG. 3.

Single appliance models often adopt a state-based formulation implicitly. For example, models of households charging plug-in electric vehicles rely on the state of the vehicles battery and whether it is at a charging station (e.g. [20], [5]). For more complicated domains involving several tasks, task-based representations have not been widely employed; finding appropriate representations is a key area of focus for future work.

Sometimes the goal is to simply model household behaviour, in order to make better predictions of future loads. For this setting, some avenues investigated so far include factored models for appliance usage modelling using hierarchical time-series models [21], and identification and clustering of load types using Dirichlet distributions [22].

Appliance State Transition Functions

On the face of it, it appears that Assumption 2 has few direct consequences other than the requirement that household state is considered in RDR schemes. However, it is an important precursor to the third assumption on household preferences.

Specifically, admitting Assumption 2 leads us to define a *state transition function* for each appliance. The transition function for appliance m is given by:

$$T^{i,m} : \times_{l \in \mathbf{m}} S^{i,l} \times S^{i,0} \times A^{i,m} \to S^{i,m}. \quad (2)$$

where $\mathbf{m}$ is the set of appliances with states that are coupled with that of appliance m, e.g. through tasks that require the use of more than one appliance (for many tasks and appliances, $\mathbf{m}$ will contain only appliance m itself). This function defines the transition between appliance states as a result of taking an action $a^{i,m} \in A^{i,m}$ to complete a task while in joint state $s^{i,\mathbf{m}} \in \times_{l \in \mathbf{m}} S^{i,m} \times S^{i,0}$.

It is also worth noting that the set of permissible appliance actions can, and likely does, vary with the appliances state. This can be seen in Fig. 3, where only actions $a1$ and $a2$ are possible in state $s1$, and likewise in the other states.

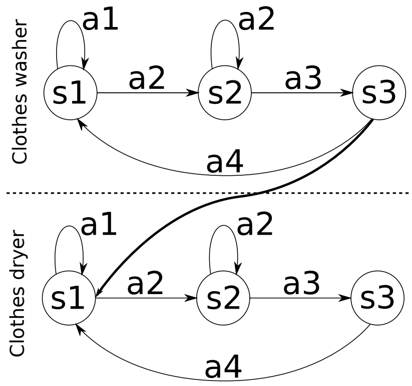

Fig. 4. A coupled state-action diagram for the joint task of "wash and dry clothes."

The transition function above gives a model of the flow of tasks and states in a household. The next assumption adds "rewards" or "preferences" to actions and state transitions. However, before this, we illustrate Assumptions 1 and 2 with a concrete example.

Running Example: Illustrating Assumptions 1 and 2

It is clear that including household state variables can significantly complicate an RDR model. However, as much as they complicate the model, the task-based couplings of household state variables often present an opportunity to prune the space of policies, as we demonstrate in the example developed in this section. Specifically, we consider an example of a power scheduling problem on a micro-grid, similar to the model developed in [2] and [7]. We will continue to use this as a running example through the paper as we add further assumptions to our RDR model.

Our example problem comprises a conventional gas generator and residential electricity users on a copper-plate microgrid, collectively selecting a discrete-time usage and generation policy. In our (minimal) example, the day is divided into three time-slots, low, shoulder and high, denoted: $h \in \{lo, sh, hi\}$.

We consider two households, A and B. Both have three tasks to complete each day: *dish washing* (DW), *clothes washing* (CW) and *clothes drying* (CD), which are undertaken by their respective appliance type. Thus, associated with each of the appliances has a set of states, such as three states for the dishwasher in Fig. 3. However, the clothes washer and dryer appliances are coupled, because for the joint task *wash and dry clothes*, CD necessarily occurs sometime after CW. This is illustrated in Fig. 4, where the thick line from CW $s3$ to the origin of CD $a2$ indicate the requirement that the clothes washer needs to have completed washing before the clothes dryer can begin.

We assume that both *DW* and *CW* use a relatively low (L) amount of energy, and that *CD* has a high (H) energy requirement, thereby satisfying Assumption 1. For demonstration purposes, we set these levels to $d^{i,CD} = H = 4$kWh and $d^{i,DW} = d^{i,CW} = L = 1$kWh. Given the loads for each appliance above, the full set of possible load schedules is given below, where each entry represents $(d^i_{lo}, d^i_{sh}, d^i_{hi})$:

$(H+2L,0,0)$	$(0,H+2L,0)$	$(0,0,H+2L)$
$(H+L,L,0)$	$(L,H+L,0)$	$(L,0,H+L)$
$(H+L,0,L)$	$(0,H+L,L)$	$(0,L,H+L)$
$(2L,H,0)$	$(H,2L,0)$	$(H,0,2L)$
$(2L,0,H)$	$(0,2L,H)$	$(0,H,2L)$
(H,L,L)	(L,H,L)	(L,L,H)

However, this does not take into account the task-based nature of the energy demand, as per Assumption 2. Specifically, *CD* is only permitted after *CW* is completed, implying that the H load must be preceded by a L load. Recognising this reduces the set of feasible schedules from 18 to the following eight:

$(2L,H,0)$	$(2L,0,H)$
$(L,H+L,0)$	$(L,0,H+L)$
(L,H,L)	(L,L,H)
$(0,L,H+L)$	$(0,2L,H)$

We have now constructed a basic model of the flow of tasks and states in a household and illustrated it with a concrete example. In the next section, we associate rewards to the actions and state transitions; that is, we introduce the notion of households' preferences for power use.

Household Preferences

In the Preliminaries section of this paper, we defined a household's utility benefit or reward as its "willingness to pay" for consuming a certain amount of electrical energy. In order to correctly model household preferences, we now extend this reward idea to our task- and state-based model of household activity.

Under assumptions 1 and 2, we argue that household control variables are of mixed type and residential electricity use is intrinsically task–related. As a consequence, a household's *preferences* for consuming electrical energy has multiple, task–related attributes that are coupled, giving it an inherently combinatorial and non–convex character.

Assumption 3: Households have private preferences that are state-based, and therefore non–convex and combinatorial; moreover, the monetary costs imposed by the system operator represents only part of their preferences for electricity use.

Specifically, this assumption should be interpreted as allowing for changes in appliances' states over time as a result of completing tasks, and allowing for couplings between the rewards for some appliances' uses. Combinatorial structure over time can be induced by substitution effects, such as tasks that are completed only once per day (e.g. running the dishwasher) and deadlines for task completion, among other things. Task-based couplings include complementarities such as preferred precedence orders on related tasks (e.g. a preference to run the clothes dryer immediately after running the clothes washer).

This assumption can be formalised by defining a reward function, given by:

$$r^i_h : S^i \times S^0 \times A^i \to \mathbb{R}_+ \tag{3}$$

such that $r^i_h(s^i_h, s^0_h, a^i_h)$ is the instantaneous reward for taking action a^i_h in joint local and global state $\{s^i_h, s^0_h\}$. Note that this formulation implies that the immediate rewards satisfy the Markov property, in that they are only dependent on the current state and the action taken. This reward function is then combined with the (aggregator supplied) cost division function, $\phi^i(C_h(x_h), \mathbf{d}^i)$ where each element of $\mathbf{d}^i$ is given by the load d^i_h in time-slot h. resulting in an instantaneous utility function of:

$$u^i_h(\mathbf{d}^i, x_h) = r^i_h(s^i_h, s^0_h, a^i_h) - \phi^i(C_h(x_h), \mathbf{d}^i). \tag{4}$$

There are a number of things to note about this formulation. First, (4) shows that coupling within a household, across the set of appliances, appears in both the reward and cost functions, while coupling across households arise through the cost function only. Second, it is expected that the cost schedule that the aggregator passes to the households is anonymous, in that only aggregate load values are reported, not other users' load information; thus, privacy is (at least partially) assured. Third, the cost division function here is left unspecified, as it comprises one of the main components of the RDR scheme design problem, but it is assumed to be dependent on the entire forecast of loads $\mathbf{d}^i$ over the decision horizon $\mathcal{H}$, not only the current usage level d^i_h.

The existence of combinatorial structures in the household's utility function implies that a household's preferences cannot be truly represented by the simplistic, static or single dimension preference models that are commonly employed in proposed RDR schemes. However, while there is no one particular combinatorial preference representation that is clearly a best fit for household preferences, some relevant examples of what may be suitable include *goal-base languages*, *CP-nets* and *dynamic preferences* including graphical models of dynamic preferences (see [23] for a survey of preference representations). One particularly appealing line of investigation is to extend the state-based load predictors in [21], [22], discussed earlier, so that they can be used as models of households preference.

Consequences for Existing Preference Models

Many household preference models currently used miss the richness of energy demand patterns, and there are several common simplifications.

Many models reduce a household's energy demand profile to a monotone concave function, often with of a form with single parameter, and/or assume that the utility for consuming power at a particular time is independent of energy use at other times [3], [7], [10]. There are two explanations for why researchers use this approach. First, standard microeconomics uses the "representative agent" method of analysis, which often invokes monotone concavity in utility. These types of utility representations function implies that households have an always increasing but marginally decreasing utility for consuming more energy, irrespective of the state of the household's task states. Microeconomic consumption also typically operates at a different time-scale to electricity use: electrical power is unique in that it cannot be (easily) stored (without significant losses), which is different from almost every other commodity, which can be stored and used over time. As such, state-based preferences are not often considered (other than in capital accumulation and investment problems). Second, early efforts to derive demand response schemes focused on large-scale and industrial energy users, as in [9]. For these types of users, monotone concave functions with intertemporal independence may be a good approximation of their utility, as industrial loads are significantly different in their characteristics to household loads. However as argued above, the same models cannot be extended to household loads, because when considerations of the combinatorial character of a household's preference are included, this style of preference modelling is inadequate.

Another common, but unjustified, simplification is that some models rely entirely on energy costs to model household preferences, and do not consider household benefit or willingness-to-pay at all (e.g. [2], [4], [24], [5]). This can be interpreted alternatively as assuming that households have *hard* requirements on their energy use, such that they have no discretion in the total energy used, only over the timing of the use. In other words, to take this approach is to implicitly apply the same reward to completing every task; that is, the models treat the completion of certain tasks as hard constraints on an energy allocation.[4] Under either interpretation, the only optimisation that is present is to minimise collective energy costs. This approach may well have its origins in the communications literature (it is also seen in the operations literature on shared processor load balancing), where delay or latency is typically the measure that both the systems operator and the users are trying to minimise, and a user's preference for its packets to reach their destination are assumed to be given by a hard quality of service requirement. However, scheduling several tasks, using only opening and deadline times, is not truly representative of the nature of demand for energy for completing coupled combinations of tasks, such as washing and drying clothes. Similarly, nor can tasks that have a discretionary component be captured by this type of preference model, such as being flexible to adjusting the temperature in a house in response to variations in energy costs.

[4]It should be noted that the *open–loop coordinated plug-in electric vehicle charging* scheme of [5] is open to users interrupting their charging when the cost is too high, i.e. above their reward, However, their model analysis does not take this type of disruption into account, nor does their price setting mechanism.

Label	$\mathbf{d}^i$	r^A	r^B
$d1$	$(2L,H,0)$	90	90
$d2$	$(2L,0,H)$	85	85
$d3$	$(L,H+L,0)$	100	100
$d4$	$(L,0,H+L)$	100	100
$d5$	(L,H,L)	95	100
$d6$	(L,L,H)	100	95
$d7$	$(0,L,H+L)$	105	110
$d8$	$(0,2L,H)$	110	100

TABLE II
TOTAL REWARDS OVER THE THREE TIME-SLOT HORIZON FOR EIGHT DIFFERENT LOAD SCHEDULES.

In several of the cases above, the major consequence of making these unwarranted assumptions is to open the predictions and proposed actions of an RDR scheme amalgamation errors, such as Simpson's paradox (the Yule–Simpson effect). These errors can adversely affect the incentives of a cost-division scheme, and thereby reduce the overall effectiveness of the RDR scheme.

The answer to these shortcomings is to adopt an appropriate combinatorial preference representation. However, one drawback of using combinatorial representations is that they suffer from exponential growth in the size as additional attributes or tasks are included. However, it is possible to choose an adequately rich but analytically and computationally feasible representation that does not lead to an exponentially large representation or intractable optimisation problem (see, e.g., [23]).

Running Example: Including State-based Preferences

We now develop our running example to include the two household agents' utilities for completing tasks. This begins with each agent's reward for completing tasks, and continues with an example of the costs of generation and an associated cost-division scheme employed by an aggregator, and is completed by integrating the two components in a quasi-linear utility model.

As argued above, electrical energy is used to complete household tasks, and the reward of consuming electrical power is the value placed on completing these tasks. Table II gives A and B's total reward over the three time-slot horizon, for completing tasks in the different slots. The tasks are completed as indicated by their discretionary power usage $\mathbf{d}^i = [d^i_{lo}, d^i_{sh}, d^i_{hi}]$. We use the labels as indicated above to denote the difference load schedules in the remainder of this example. These rewards express the preferences households have for completing tasks during different times during the day.

Combining the cost and value functions for the agents in a micro-grid produces a electricity usage scheduling game.

$\mathbf{d}^A$	$\mathbf{d}^B$ $d6$	$d7$	$d8$
$d6$	27, 22	43, 53	25, 35
$d7$	48, 38	12.5, 17.5	20.5, 25.5
$d8$	35, 20	25.5,25.5	31.5,31.5

TABLE III

UTILITIES FOR THE REDUCED SCHEDULING GAME.

Assume that there is an amount of non–discretionary aggregate energy use (from lights, refrigerator, and other uninterruptable or shiftable loads) in each slot of $\tilde{\mathbf{d}} = (1,2,3)$, which is in addition to the amounts given by the table above. Now, let the aggregate *generator cost* per time slot, $c_h(x_h)$, be a typical quadratic function of total demand in each slot, given by: $c_h(x_h) = ax_h + bx_h^2 + c$. Let where $\mathbf{x}$ is the vector of total power levels for one slot, which includes the non-discretionary loads, i.e.: $x_h = \sum_{i \in I} d_h^i + \tilde{d}_h$. Then, total generation costs per day are:

$$C(\mathbf{x}) = c_{lo}(x_{lo}) + c_{sh}(x_{sh}) + c_{hi}(x_{hi}). \tag{5}$$

For this example, set the cost parameters to $a = 0$, $b = 1$ and $c = 0$, so that $c(x) = x^2$. The total generation costs for the possible combinations of demand levels for our two-player example are given in Appendix A.

Using the quasi-linear utility function to combine these costs with the rewards from earlier gives the utilities for the two households over the three-step decision horizon, which can be found in Appendix A. Also there, we show that since the households have schedules that are *dominated*, we can apply the standard method of iterated elimination of dominated strategies to reduce the number of joint schedules requiring analysis. This eliminates all actions other than $d6$, $d7$ and $d8$ for both A and B, and the resulting rational strategic opportunities for the households is succinctly represented by the bi-matrix in Table III.

If this payoff matrix was known to the aggregator and/or to the other households participating in the RDR scheme — that is, if it was a game of *complete information* — then it would be a potential game. It would have three pure-strategy Nash equilibria, at $(\mathbf{x}^A, \mathbf{x}^B) = (d7, d6)$, $(d6, d7)$ and $(d8, d8)$. Note that the equilibrium at $(d6, d7)$ optimises the system's *social welfare*; that is, $\mathrm{argmax}[\sum_{i \in N} U^i(\mathbf{d}^i, \mathbf{x})] = (d6, d7)$.

However, since the households' rewards for electricity usage are not known to either the aggregator or the other households, that is, they are *private information*, the game is one of *incomplete information*. Therefore the RDR scenario is not adequately represented by the type of game above, as the Nash equilibrium analysis is not able to give a complete picture of the households' behaviour. For example, one approach to remedy the incompleteness of game information is for the aggregator to ask the households to reveal their rewards, and then instruct the households to play the socially optimal Nash equilibrium. We build on this idea in the next section, where we argue that households are *strategic*, and by employing such a method in the presence of incomplete information, the aggregator provides an incentive for households to benefit by mis-reporting their rewards.

Household Behaviour

In order to achieve their aims and reach their goal, households will be strategic; it is in their self-interest to be so. Assuming that they will not try to game the system in place is not a reasonable approximation of reality and clashes with motivation for using a non–cooperative game framework.

Assumption 4: Household behaviour is strategic, both at the level of equilibrium analysis and algorithmic implementation.

This is by far the most intricate of the four assumptions: By assuming that households will act rationally and strategically to pursue their own self-interest, the problem of RDR aggregation falls squarely in the realm of algorithmic game theory, and in particular, *algorithmic mechanism design*. The general aim of mechanism design is to derive or construct an allocation rule and a cost division (or payment) scheme that, together, allocate a good to those participants that value it the most — the *efficient* allocations. A mechanism is typically proven to be efficient by showing that, under the allocation rule and payment scheme, making a *truthful* preference report is the best action for every agent.

However, the task of deriving allocation and pricing rules that implement an electrical energy allocation over time that is (approximately) truthful and efficient is difficult. Indeed, the "approximately" here is an acknowledgement of the fact that the application of mechanism design principles is complicated by several additional constraints that are inherent to RDR aggregation, which are partially covered by our preceding assumptions. These are:

1) Electrical power usage is often at discrete levels (Assumption 1), which has consequences for choice of allocation rules that can be applied, as different rules will have different equilibria and aggregator revenue, and only certain methods of computation can be applied when usage levels take discrete or hybrid values (cf. the discussion of integer-splittable congestion games earlier).

2) Preferences over power usage profiles are combinatorial, and vary over time as a function of the tasks being undertaken in a household (Assumptions 2 and 3). Furthermore, the state of a household may not be directly observable to a mechanism, which may prevent the application of some approaches. For example, online and dynamic VCG mechanisms require accurate models of participants behaviour, but these may not be available if household states cannot be accurately observed by the aggregator [25], [26].

3) It is not clear how to express values over alternate power usage profiles (Assumption 3). A fully expressive representation of a households preferences over all possible power

allocations would be impractically large (implying infeasibly large communication times and computation requirements); while lesser representations may miss important aspects of households behaviour if they are not carefully constructed.

4) Computation and communication available to the aggregator and the households is constrained.

These four constraints proscribe the direct application of standard mechanisms. In the remainder of this section, we examine how strategic household behaviour affects the operation of existing optimisation routines for RDR, and discuss some way that it can be accommodated in RDR schemes.

Consequences for Existing Optimisation Routines

Many RDR schemes fail to ensure that incentive compatibility is preserved, or even approximated, during the actual running of their proposed algorithms and optimisation routines. Consider [4], [6], [2], [3]: All assume that, during negotiation or distribute computation, the agents act truthfully in response to their preferences. However, with private preferences for different slots, mis–reporting during the procedure can be beneficial for some households, and this may lead to an inefficient allocation of the energy resources (a similar effect is noted in [27]).

In other distributed optimisation domains where non–cooperative game models and techniques have been deployed, the system is composed of *cooperative* agents. Like several of the papers above, these models and techniques have typically focused on convex, potential and congestion game representations. Examples of this approach include consensus problems [28], resource allocation games [29], [16] and distributed constraint optimisation problems [30].

Given their success, the non–cooperative methods have been seen as useful in the context of RDR optimisation and control. However, there is a key difference between the settings described above the the RDR scenario. That is in RDR scenarios, households have their own private preferences, which are unknown to the system; whereas the techniques themselves were developed for cooperative problems composed of artificial or computational agents *without private preferences*. As such, their utility functions could be defined absolutely by the system designer.

This clearly cannot be done in the domain of RDR: As we argued in the previous section, households' have their own preferences for power use. These preferences cannot be ignored by the system designer and must be addressed in some manner. We now consider an example that clarifies this point.

In framework proposed in [2],[5] households are assumed to minimise their costs for using electrical energy. Put another way, the model derived assumes that the only preference that households hold is to minimise their energy costs, and that they have no other preferences for the timing of their energy use. Given this, the households' utilities can be derived so that they admit a suitable convexity condition (derived by [13]), which ensures that conventional convex optimisation techniques can be used to solve the game for its Nash equilibrium. Although several techniques could have been employed, the paper proposes that the game be solved iteratively, by each household broadcasting its best-reply energy use schedule (computed by the household using an interior-point method) in response to the most recently calculated aggregate energy use schedule.

[5]In [4], a similar model is derived that also incorporates stochastic wind power generation.

Now, under the assumption that the agents are truthful, the procedure in [2] does converge to a Nash equilibrium. Moreover, under the (spurious) assumption that the households have no preference for use of energy other than cost minimisation, the authors also claim that their method is strategy–proof: that is, there is no benefit to any household or a group of households to announce an incorrect usage schedule.

However, when households do have private preferences in addition to cost minimisation, the household (specifically, a subset of them) can game the system to their advantage. This could happen in the following way:

1) A subset of the users $Z \subset I$ collude and all report that they will use more energy than they actually intend to in a slot, k, that is particularly desirable to them.
2) The energy aggregator increases the price in slot k.
3) Other 'naïve' users in $Y = I \setminus Z$ respond by shifting some of their load out of k.
4) At the actual time of energy use, the members of S don't meet their power usage forecasts, therefore energy costs are lower than forecast.
5) Thus, the costs to members of Z are lower than if those members of Y that shifted their loads from k had not shifted. Moreover, the costs to the naïve users is greater, and the overall system cost greater, than if the manipulation had not taken place.

In addition, the complications of adding in private preferences — specifically the non–convexities argued earlier — would also likely cause the game model to fall outside the class of convex games (and possibly potential games), so the optimisation routine may fail to converge for this reason.

In the next section we give a concrete demonstration of this type of manipulation in our running example model. Beyond this, the type of manipulation above could be applied in the context of many proposed RDR methods, including those that do not adopt a game–theoretic framework, such as message passing schemes based on variational inequalities [6], distributed Lagrangian methods and proximal message passing [31].

Running Example: Failing to Consider Strategic Behaviour

Let us assume that an aggregator is used to coordinate the allocation to energy in our microgrid running example. Moreover, assume that the aggregator uses the method outlined at the end of the Household Preferences section (which is also the style of algorithm employed in [32], [5], [8], and others):

1) The households communicate their rewards $r^i(\mathbf{d}^i)$ (not utilities) for different energy use schedules to the aggregator (i.e. their values from Table II;
2) The aggregator computes utilities agents (i.e. the tables given in Appendix A), and deduces the social welfare optimising Nash equilibrium that maximises the sum of household utilities;
3) The aggregator instructs each household to play its component of this solution; and
4) The households play what they are told, because it is a Nash equilibrium.

As noted in the previous Running Example section, the joint schedule $(d6, d7)$ is social-welfare optimal, so if the households were truthful, this is the solution that the aggregator would prescribe (see Table III).

Now assume that household A is strategic and does not necessarily truthfully report its rewards to the aggregator, while B is truthful. Household A would prefer the equilibrium at $(d7, d6)$ to be chosen over that of $(d6, d7)$, as this maximises its own utility. In order to have the aggregator choose this equilibrium, A can manipulate the coordination scheme above by misreporting its reward values. There are many ways it could undertake the manipulation, including the following two:

- Household A could move the aggregators choice of equilibrium by misreporting a higher reward for schedule $d7$, for example, by reporting $r^i(d7) = 120$.
- Less obviously, A could lower its reported reward for $d6$ and raise its value for $d8$, for example by reporting $r^i(d6) = 90$ and $r^i(d8) = 120$. This would move the equilibrium at $(d6, d7)$ to $(d8, d7)$, which, even with they higher reported reward for $d8$, has a lower social welfare than $(d6, d7)$.

Either way, this results in an economically inefficient allocation of the power to the householders. If both households were strategic, the ability of the scheme above to allocate energy use schedules would be seriously curtailed. Above and beyond this, it is very possible that in a large power distribution system with many households, a significant group of households could coordinate to undertake similar manipulations on a large scale.

This type of behaviour is possible in other coordination schemes, including those where the households provide forecasts of their future energy needs (e.g. [2], [6]). It is not immediately clear how to prevent manipulations of these types, however in the next section, we discuss some promising lines of research in mechanism design that could be put towards developing RDR schemes that are robust to strategic households.

Accommodating Strategic Household Behaviour

Earlier in this section we discussed some complicating features of the RDR domain, which make it impractical to apply off-the-shelf or standard mechanisms to RDR scenarios. We now give three brief examples of problems that arise when applying standard mechanisms to RDR problems, to illustrate the need for new approaches in order to develop useful RDR schemes.

The first is the VCG (Vickrey–Clarke–Groves) mechanism [11]. VCG is the gold standard in mechanism design and can be thought of as a generalisation of the well known second-price sealed-bid auction. Now, we have argued that a household's utility for electrical energy use is combinatorial, in that the value of using energy now depends on the household's state, which itself depends on how much energy is allocated in the future or was in the past (e.g. dishes or clothes only need to be washed once). With reference to our running example, assume that the aggregator operates on a more-realistic hourly division of a day into of 24 time-slots, and computes an allocation one day at a time. Continue to assume that each household can ask for only high or low power in each slot. Under VCG, each household would be asked to submit a valuation on all of its possible power usage profiles over 24 hours; unfortunately, this tallies to $2^{24} - 1 > 16$ million values for each household. Thus, given the scenarios time constraints, directly applying VCG is infeasible, since it will not scale to the numbers of households likely to be involved in RDR schemes. Moreover, dynamic VCG mechanisms for various specific settings have derived by [25], [26], and, recently, [27] proposed a dynamic VCG–like mechanism that induces generators to implement the optimal control as computed by the systems operator: It is noted by the authors of all three of these works that the mechanisms they derive are computationally intractable for large problems.

A second example of an infeasible method for aggregating RDR are the uniform-price supply function market mechanisms used in wholesale electricity markets for dispatchable generation (i.e., in Australia, the market to determine the economic dispatch order run 40 hours in advance of dispatch time [33]). This is a reverse auction, where bids are generation levels for each price level; that is, generators submit complete supply schedules for all price levels. If generally applied, these auctions would consider all possible supply function forms, which is clearly not possible. Thus, in practice they restrict the bids to certain forms of supply functions that are well-behaved (e.g. affine or piece-wise linear functions). However, these restrictions on the preferences are too severe and lack the combinatorial character that household preference representations require. However, if this constraint on supply function form is relaxed, the communication required to transmit bids to the aggregator may be very large, and, moreover, the computation required would be great or infeasible, thus, again, failing to scale.

Third, as an example of effective practical combinatorial auction design, and also a partial contradistinctive example to the formats above, is the development of combinatorial clock auctions (CCAs) in the sale of radio spectrum in the USA ([34], chapter 5). CCAs were used because the bidders had complicated combinatorial preferences over bundles of spectrum. Without going into detail, CCAs are not true mechanisms in the theoretical sense, so may be only approximately efficient or truthful. Rather, CCAs are sensibly engineered auction formats that have proved to work well in specific domains. In particular, they have the benefit of allowing for *price discovery* — that is, of focusing the auction participants' attention on the pertinent areas of the allocation and pricing space — without bidders having to submit their full demand functions to the aggregator, thereby reducing the aggregator's computational load and the system's communication requirements.

It should be noted that [6] do acknowledge the possibility of the type of manipulation discussed in the example above, and discuss ways to moderate this type of behaviour using ad-hoc penalties for misleading forecasts. A more principled approach may be use *scoring rules*, which have been proposed for the problem of eliciting costly probabilistic forecasts [35], [36].

Conclusion

In this paper we have emphasised the ground-up approach that is needed to correctly characterise the RDR domain. In doing so, we have proposed four assumptions that we argue are necessary for an RDR to satisfy for it to be usefully applied to real RDR problems. We have illustrated the effects of violating our proposed assumptions, with reference to other proposed RDR schemes. We also have given several examples of techniques that satisfy each assumption; a thorough investigation of these will provide the basis of our future work in this area.

Appendix A

This appendix is included for completeness, and contains the full matrices for the running example used through the paper, and specifically those from the Household Preference section.

According to the parameters given earlier, the total generation costs for the possible combinations of demand levels for the two-player example are given in the following (symmetric) matrix:

$C(\mathbf{x})$	$\mathbf{d}^B$							
$\mathbf{d}^A$	$d1$	$d2$	$d3$	$d4$	$d5$	$d6$	$d7$	$d8$
$d1$	125							
$d2$	110	146						
$d3$	137	114	153					
$d4$	116	160	122	178				
$d5$	132	105	146	126	134			
$d6$	114	146	122	162	122	146		
$d7$	122	162	132	182	134	164	185	
$d8$	122	146	134	164	132	150	169	157

Following [2], costs are divided in proportion to the households total energy use per day. This is the same for both households, so costs are divided equally between A and B. Combining these costs with the rewards in Table II according to equations (4) and (1) gives the following full utility matrix for household A:

U^A	$\mathbf{d}^B$							
$\mathbf{d}^A$	$d1$	$d2$	$d3$	$d4$	$d5$	$d6$	$d7$	$d8$
$d1$	27.5	35	21.5	32	24	33	29	29
$d2$	30	12	28	5	32.5	12	4	12
$d3$	31.5	43	23.5	39	27	39	34	33
$d4$	42	20	39	11	37	19	9	18
$d5$	29	42.5	22	32	28	34	28	29
$d6$	43	27	39	19	39	27	43	25
$d7$	44	24	39	14	38	48	12.5	20.5
$d8$	49	37	43	28	44	35	25.5	31.5

and for household B:

U^B	$\mathbf{d}^B$							
$\mathbf{d}^A$	$d1$	$d2$	$d3$	$d4$	$d5$	$d6$	$d7$	$d8$
$d1$	27.5	30	31.5	42	34	38	49	49
$d2$	35	12	43	20	47.5	22	29	37
$d3$	21.5	28	23.5	39	27	34	44	43
$d4$	32	5	39	11	37	14	19	28
$d5$	24	32.5	27	37	33	34	43	44
$d6$	33	12	39	19	39	22	53	35
$d7$	29	4	34	9	33	38	17.5	25.5
$d8$	29	12	33	18	34	20	25.5	31.5

Dominated strategies can be iteratively eliminated as follows: (i) Starting with A, its schedules $d1$, $d2$ and $d4$ are eliminated, all other schedules are best-responses to at least one schedule of B; (ii) For B, in addition to $d1$, $d2$ and $d4$, schedules $d3$ and $d5$ are also eliminated, as they are only best-responses to A's schedules $d4$ and $d2$, respectively, which were eliminated at the first iteration; (iii) By the same reasoning, A also eliminates $d3$ and $d5$. No other dominated strategies remain, and the resulting 3×3 bi-matrix is given in Table III.

Notation

Variable	Description
$h \in \mathcal{H}$	Decision slot and set of H slots to the horizon.
$i \in I$	Household i and set of I households.
$m \in \mathcal{M}^i$	Household i's set of appliances, indexed m.
$a^i \in A^i$	Household i's action and action space.
r^i_h	Household i's reward function.
$d^{i,m}_h$	Energy used by appliance m^i in slot h.
d^i_h	Total energy used by i in slot h.
$\mathbf{d}^i$	Vector of total energy use by i over $\mathcal{H}$.
x_h	Total energy used by all households in slot h.
$\mathbf{x}$	Vector of total energy use for each h over $\mathcal{H}$.
C_h	Aggregator's total cost function over $\mathcal{H}$.
ϕ	Aggregator's cost division vector function.
ϕ^i	Household i's component of the cost division.
u^i_h	Utility function of i for slot h.
U^i	Total utility of i over $\mathcal{H}$.
$s^i \in S^i$	Household i's state vector and state space.
$s^{i,m} \in S^{i,m}$	State and state space of appliance m^i.
$T^{i,m}$	State transition function for appliance $m \in \mathcal{M}^i$.
$\tilde{\mathbf{d}}$	Vector of total non-discriminatory energy use.

References

[1] W. Saad, Z. Han, H. V. Poor, and T. Basar, "Game theoretic methods for the smart grid," *IEEE Signal Processing Magazine, Special issue on Signal Processing for the Smart Grid*, vol. 29, no. 5, pp. 86–105, 2012.

[2] H. Mohsenian-Rad, V. W. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia, "Autonomous demand side management based on game-theoretic energy consumption scheduling for the future smart grid," *IEEE Trans. Smart Grid*, vol. 1, no. 3, pp. 320–331, Dec. 2010.

[3] C. Ibars, M. Navarro, and L. Giupponi, "Distributed demand management in smart grid with a congestion game," in *Proc. of 1st IEEE Int. Conf. Smart Grid Communications*, Oct. 2010, pp. 495 –500.

[4] C. Wu, H. Mohsenian-Rad, J. Huang, and A. Y. Wang, "Demand side management for wind power integration in microgrid using dynamic potential game theory," in *IEEE Int. Workshop on Smart Grid Communications and Networks*, 2011, pp. 1199–1204.

[5] Z. Ma, D. S. Callaway, and I. A. Hiskens, "Decentralized charging control of large populations of plug-in electric vehicles," *IEEE Trans. Control Systems Technology*, vol. 21, no. 1, pp. 67–78, 2013.

[6] I. Atzeni, L. G. Ordonez, G. Scutari, D. P. Palomar, and J. R. Fonollosa, "Day–ahead bidding strategies for demand–side expected cost minimization," in *Proc. of 3rd IEEE Int. Conf. Smart Grid Communications*, Tainan City, Taiwan, 5–8 Nov. 2012.

[7] P. Samadi, H. Mohsenian-Rad, R. Schober, and V. Wong, "Advanced demand side management for the future smart grid using mechanism design," *IEEE Trans.'Smart Grid*, vol. 3, no. 3, pp. 1170–1180, 2012.

[8] C. Wu, H. Mohsenian-Rad, and J. Huang, "Vehicle-to-aggregator interaction game," *Smart Grid, IEEE Transactions on*, vol. 3, no. 1, pp. 434–442, 2012.

[9] M. Fahrioğlu and F. L. Alvarado, "Using utility information to calibrate customer demand management behavior models," *IEEE Transactions on Power Systems*, vol. 16, no. 2, pp. 317–322, 2001.

[10] M. Roozbehani, M. A. Dahleh, and S. K. Mitter, "Volatility of power grids under real-time pricing," *Power Systems, IEEE Transactions on*, vol. 27, no. 4, pp. 1926–1940, 2012.

[11] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. Cambridge, UK: Cambridge Univeristy Press, 2007.

[12] W. Saad, Z. Han, and H. Poor, "Coalitional game theory for cooperative micro-grid distribution networks," in *Proc. IEEE International Conference on Communications (ICC'11), Workshop on Smartgrid communications*. IEEE, 2011, pp. 1–5.

[13] J. B. Rosen, "Existence and uniqueness of of equilibrium points for concave *n*–person games," *Econometrica*, vol. 33, pp. 347–351, 1965.

[14] D. Monderer and L. S. Shapley, "Potential games," *Games and Economic Behavior*, vol. 14, pp. 124–143, 1996.

[15] J. A. M. Potters, T. E. S. Raghavan, and S. H. Tijs, "Pure equilibrium strategies for stochastic games via potential functions," in *Advances in Dynamic Games and their Applications*, 2009, pp. 433–444.

[16] A. Chapman and P. Varakantham, "Marginal contribution stochastic games for dynamic resource allocation: Bounds and convergence," in *6th International Workshop on Optimisation in Multi-Agent Systems (OPTMAS '13)*, at AAMAS 2013, St Paul, MN, USA, May 2013.

[17] A. Neyman, "Correlated equilibrium and potential games," *Int. J. Game Theory*, vol. 26, no. 2, pp. 223–227, 1997.

[18] C. A. Meyers, "Network flow problems and congestion games: complexity and approximation results." Ph.D. dissertation, Massachusetts Institute of Technology, 2006.

[19] L. Tran-Thanh, M. Polukarov, A. C. Chapman, A. Rogers, and N. J. Jennings, "On the existence of pure strategy Nash equilibria in integer–splittable weighted congestion games," in *The 4th Symposium on Algorithmic Game Theory*, 2011.

[20] S. Stein, E. Gerding, V. Robu, and N. R. Jennings, "A model-based online mechanism with pre-commitment and its application to electric vehicle charging," in *Proc. of AAMAS '12*, 2012, pp. 669–676.

[21] P. P. Reddy and M. M. Veloso, "Factored models for multiscale decision-making in smart grid customers," in *Twenty-Sixth Conference on Artificial Intelligence (AAAI-12)*, July 2012.

[22] N. C. Truong, J. McInerney, L. Tran-Thanh, E. Costanza, and S. D. Ramchurn, "Forecasting multi-appliance usage for smart home energy management," in *23rd International Joint Conference on Artificial Intelligence (IJCAI 2013)*, April 2013.

[23] J. Uckelman, "More than the sum of its parts: Compact preference representation over combinatorial domains," Ph.D. dissertation, University of Amsterdam, 2009, ILLC Publication DS-2009-12.

[24] R. Hermans, M. Almassalkhi, and I. Hiskens, "Incentive-based coordinated charging control of plug-in electric vehicles at the distribution-transformer level," in *Proc. American Control Conference (ACC '12)*, 2012, pp. 264–269.

[25] D. Bergemann and J. Välimäki, "The dynamic pivot mechanism," *Econometrica*, vol. 78, no. 2, pp. 771–789, 2010.

[26] D. C. Parkes, R. Cavallo, F. Constantin, and S. Singh, "Dynamic incentive mechanisms," in *AI Magazine*. AAAI Press, 2010, vol. 31, no. 4, pp. 79–94.

[27] T. Tanaka, A. Cheng, and C. Langbort, "A dynamic pivot mechanism with application to real time pricing in power systems," in *Proc. American Control Conference (ACC '12)*, 2012, pp. 3705–3711.

[28] J. R. Marden, G. Arslan, and J. S. Shamma, "Connections between cooperative control and potential games illustrated on the consensus problem," in *Proceedings of the 2007 European Control Conference (ECC '07)*, July 2007.

[29] J. R. Marden and A. Wierman, "Distributed welfare games," in *Proceedings of the 47th IEEE Conference on Decision and Control (CDC–08)*, December 2008.

[30] A. C. Chapman, A. Rogers, N. R. Jennings, and D. S. Leslie, "A unifying framework for iterative approximate best response algorithms for distributed constraint optimisation problems," *The Knowledge Engineering Review*, vol. 26, no. 4, pp. 411–4, 2011.

[31] M. Kraning, E. Chu, J. Lavaei, and S. Boyd, "Dynamic network energy management via proximal message passing," *Foundations and Trends in Optimization*, vol. 1, no. 2, 2013.

[32] P. Samadi, A.-H. Mohsenian-Rad, R. S. V.W.S., Wong, and J. Jatskevich, "Optimal real-time pricing algorithm based on utility maximization for smart grid," in *Proc. 1st IEEE Int. Conf. Smart Grid Communications (SmartGridComm'10)*, 2010, pp. 415–420.

[33] Australian Energy Market Operator, "Pre-dispatch process description," 2010. [Online]. Available: http://www.aemo.com.au/Electricity/Market-Operations/Dispatch/Predispatch-Process-Description

[34] P. Cramton, Y. Shoham, and R. Steinberg, *Combinatorial Auctions*. Cambridge, MA, USA: MIT Press, 2005.

[35] A. Papakonstantinou, A. Rogers, E. H. Gerding, and N. R. Jennings, "Mechanism design for the truthful elicitation of costly probabilistic estimates in distributed information systems," *Artificial Intelligence*, vol. 175, no. 2, pp. 648–672, 2011.

[36] H. Rose, A. Rogers, and E. H. Gerding, "A scoring rule-based mechanism for aggregate demand prediction in the smart grid," in *Proc. of 11th Int. Conf. Autonomous Agents and Multiagent Systems (AAMAS-12)*.