

Residential Demand Response under Uncertainty

**Paul Scott and Sylvie Thiébaux and
Menkes van den Briel**
Australian National University
NICTA, Australia
{firstname.lastname}@nicta.com.au

Pascal Van Hentenryck
University of Melbourne
NICTA, Australia
pvh@nicta.com.au

Abstract

This paper considers a residential market with real-time electricity pricing and flexible electricity consumption profiles for customers. Such a market raises an optimisation problem for home automation systems where they need to schedule consumption activities to reduce costs, whilst maintaining a base level of comfort and convenience. This optimisation problem faces uncertainty in real-time prices, weather conditions, and occupant behaviour. The paper presents an online stochastic combinatorial optimisation algorithm that produces fast, high-quality solutions to this problem. This algorithm is compared with reactive control strategies and an approach using an expected scenario. Our results demonstrate the value of stochastic information and online stochastic optimisation in residential demand response.

1 Introduction

Electricity consumption in residential markets will undergo fundamental changes in the next decade due to the availability of solar panels and novel pricing mechanisms, progress in batteries and electric cars, and the emergence of smart appliances and home automation. These technologies provide residential customers with the ability to actively participate in smart grid activities such as demand response where loads are shifted to times favourable for the network as a whole.

Having an intelligent Home Automation System (HAS) within each home is a key component in this vision. The HAS receives information about device operating characteristics, usage requests and network signals, and can send control actions back to smart devices. Apart from providing useful feedback to occupants on their consumption habits, it can also autonomously make control decisions. Through this control the HAS can target one or more of the following objectives:

1. Improve occupant comfort,
2. Reduce overall electricity consumption,
3. Perform demand response for network.

These objectives are often conflicting, so in order to get the right balance, occupants need to indicate how they value comfort against cost savings. The task of the HAS is then to decide on a series of control actions to take over time,

which produces an optimal solution for the weighted combination of objectives. The HAS can implement simple policies to try and meet these conflicting objectives. Or, more interestingly, it can use sophisticated stochastic optimisation technology which exploits forecasts and observed patterns in prices, weather, residential activities and smart device usage.

This paper aims to determine the benefits of online stochastic optimisation for a HAS that is exposed to Real-Time Pricing (RTP) as a demand response mechanism. A number of research projects have started examining this very issue (see the related work section) but they often give an incomplete picture of the benefits of optimisation and the value of stochastic information. These projects often consider simpler uncertainty models, which give a partial understanding of the true benefits that optimisation can bring to this setting. In contrast, this paper makes two primary contributions: one conceptual and one algorithmic.

At the conceptual level, the paper presents a compositional architecture for HAS optimisation, where each device can be modelled independently in terms of a collection of functions that encapsulate its behaviour. These devices are then assembled into a model of a home, from which the HAS optimisation problems can derive.

At the algorithmic level, the paper presents a comprehensive study of the value of HAS optimisation in the presence of uncertainty about future prices, occupant behaviour, and environmental conditions. Our formulation uses models representative of physical devices and stochastic models trained on real weather and network demand data. These device and stochastic models are used in two online stochastic optimisation algorithms which are compared to simple control systems based on reactive policies.

The experimental results not only show the value of stochastic information, but also that stochastic optimisation provides solutions that are close to the clairvoyant solutions which have perfect knowledge of the future. The online stochastic algorithms using MILP technology are fast and they produce significantly better solutions than the reactive controllers. Also of interest is the comparison between the two online stochastic algorithms, and an experiment that shows the distance into the future that needs to be considered when making a decision.

The rest of the paper presents the deterministic HAS optimisation problem, its stochastic version, the stochastic mod-

els, and the experimental results.

2 Deterministic HAS Optimisation

A house contains a collection of controllable devices which influence the amount of power consumed in the house and the level of comfort that residents experience. We consider the operation of these devices over discrete time steps¹: $\forall i \in \mathbb{Z} : t_i \in \mathbb{R}$ where $t_i > t_{i-1}$ and $\forall i \in \mathbb{Z} : t_i^{stp} = t_i - t_{i-1}$.

Given a real time price for electricity and other input parameters (e.g., external temperatures and device requests), optimal operation of these devices is achieved by minimising the sum of monetary and comfort costs. The optimisation problem decision variables are the device actions at each time step, which are constrained by device characteristics and total power limits on the house.

2.1 Formal Definition

We start with a new formal definition of a device, which is a collection of functions that govern the device operation. These include functions for permissible device actions, state updates, the electrical power transferred with the house, and any non-power related operation costs. Operation costs are always positive and may include any occupant comfort costs, fuel consumption or wear and tear on the equipment. By convention power consumed by the device is negative, and power generated, e.g., by a rooftop photovoltaic system is positive.

Definition 1 (Device)

A device is a tuple $d = (A_d, S_d, R_d, q_d, g_d, f_d, l_d)$, where:

- $A_d \subseteq \mathbb{R}^{m_a} \times \mathbb{Z}^{m'_a}$ is the set of device actions
- $S_d \subseteq \mathbb{R}^{k_a} \times \mathbb{Z}^{k'_a}$ is the set of device states
- $R_d \subseteq \mathbb{R}^{w_a} \times \mathbb{Z}^{w'_a}$ is the set of device input parameters
- $q_d : S_d \times R_d \rightarrow \mathcal{P}(A_d)$ is the permissible action function
- $g_d : A_d \times S_d \times R_d \rightarrow S_d$ is the state update function
- $f_d : A_d \times S_d \rightarrow \mathbb{R}$ is the electrical power function
- $l_d : A_d \times S_d \times R_d \times \mathbb{R} \rightarrow \mathbb{R}$ is the operational cost function

A house is simply a set of devices, together with bounds on the instantaneous amount of power the house can transfer to or from the grid:

Definition 2 (House)

A house is a tuple $h = (D_h, \underline{p}_h, \bar{p}_h)$, where:

- D_h is the set of devices
- $\underline{p}_h, \bar{p}_h \in \mathbb{R}$ are the lower and upper power limits

We now turn to the deterministic formulation of the HAS optimisation problem which will be later used as a building block for our stochastic formulation. The deterministic formulation assumes that the input parameters are known over a horizon of n time steps. The task is to choose device actions at each step to reduce the total cost over the horizon.

¹Variable time step sizes will be used to focus computational time where most needed.

Inputs include the device initial states, the RTP, the house background power² and the device input parameters at each step. We account for the fact that the RTP is often different depending on whether power is bought from or sold to the grid. The optimisation variables at each time step include the device actions and states, and the device and house power consumptions and costs. These variables and inputs are linked together via the device functions in Definition 1 and house power limit constraints in Definition 2.

We use the following notation: $(a)_+ = |a|$ if $a > 0$ and 0 otherwise, and similarly $(a)_- = |a|$ if $a < 0$ and 0 otherwise, where $a \in \mathbb{R}$.

Definition 3 (Deterministic HAS Optimisation Problem)

Let $h = (D_h, \underline{p}_h, \bar{p}_h)$ be a house. The HAS optimisation problem over a horizon $n \in \mathbb{N}^*$ for h is the following:

Inputs:

for each device $d = (A_d, S_d, R_d, q_d, g_d, f_d, l_d) \in D_h$

- $s_{d,0} \in S_d$ is the device initial state
- for each device $d \in D_h$ and time step $i \in \{1 \dots n\}$
- $r_{d,i} \in R_d$ are the device input parameters
- for each time step $i \in \{1 \dots n\}$
- $p_{h,i}^b \in \mathbb{R}_-$ is the house background power
- $v_i \in \mathbb{R}^2$ is the real-time price (buying, selling)

Decision variables:

for each device $d \in D_h$ and time step $i \in \{1 \dots n\}$

- $a_{d,i} \in A_d$ are the device action variables

Other variables:

for each device $d \in D_h$ and time step $i \in \{1 \dots n\}$

- $s_{d,i} \in S_d$ are the device state variables
- $p_{d,i} \in \mathbb{R}$ is the device power
- $c_{d,i} \in \mathbb{R}_+$ is the device operation cost
- for each time step $i \in \{1 \dots n\}$
- $p_{h,i} \in [\underline{p}_h, \bar{p}_h]$ is the total power
- $c_{h,i} \in \mathbb{R}$ is the total cost

Constraints:

for each device $d \in D_h$ and time step $i \in \{1 \dots n\}$

- $a_{d,i} \in q_d(s_{d,i-1}, r_{d,i})$ is the action permissibility constraint
- $s_{d,i} = g_d(a_{d,i}, s_{d,i-1}, r_{d,i})$ is the state update constraint
- $p_{d,i} = f_d(a_{d,i}, s_{d,i})$ is the device power constraint
- $c_{d,i} = l_d(a_{d,i}, s_{d,i}, r_{d,i}, t_i^{stp})$ is the device cost constraint

for each time step $i \in \{1 \dots n\}$

- $p_{h,i} = \sum_{d \in D_h} p_{d,i} + p_{h,i}^b$ is the house power constraint
- $\underline{p}_h \leq p_{h,i} \leq \bar{p}_h$ is the house power limits constraint
- $c_{h,i} = \sum_{d \in D_h} c_{d,i} + t_i^{stp} v_{i,1}(p_{h,i})_- - t_i^{stp} v_{i,2}(p_{h,i})_+$ is the house cost constraint

Objective:

$$\min \sum_{i=1}^n c_{h,i}$$

²This aggregates uncontrollable electrical consumption, e.g., lighting, entertainment and cooking.

2.2 Modelled Devices

In our experiments we consider a modern house with electrical HVAC, hot water heating, solar panels, washing machine, clothes dryer and dish washer. We also include two devices that are expected to become standard within the next decades, an electric vehicle (EV) and a dedicated battery bank for storing electrical energy. Descriptions of these devices are given in this section. Some liberty has been used in these descriptions to aid understanding, but with slight reformulation they all fit into the rigorous device definition. Device electrical powers and operational costs are consistently represented by the variables p_i and c_i , where a negative power represents power consumed by a device.

The physical behaviour of devices has been approximated by linearising their physical equations and discretising time. Only significant steps of this process are mentioned in the device descriptions. For the experiments parameters were selected to be representative of typical devices. For example, the EV battery capacity is equal to that of a Nissan Leaf, and the house floor area for heating purposes is typical of an average-sized house. Due to the difficulty in obtaining some parameters, for example the charging efficiency of a Nissan Leaf, and since our models are an approximation of these real systems anyway, some estimates had to be made.

Battery. A battery has a stored energy state $E \in [0, \bar{E}]$ and a charge/discharge power $p \in [\underline{p}, \bar{p}]$ action variable. Energy is lost through a fixed efficiency η when power is charged into the battery. The stored energy state update function is given by:

$$E_i = E_{i-1} + t_i^{stp} (\eta(p_i)_- - (p_i)_+) \quad (1)$$

A battery lifetime cost c is associated with any power that is discharged from the battery through a lifetime price v :

$$c_i = v(p_i)_+ \quad (2)$$

Electric Vehicle. An electric vehicle (EV) is essentially the same as the battery just presented, but with a few additional constraints. Firstly the EV battery can only be charged/discharged when it is at home, which is indicated by the input parameter $x^h \in \{0, 1\}$:

$$x_i^h = 0 \implies p_i = 0 \quad (3)$$

The input parameter $p^d \in \mathbb{R}_+$ represents the power drawn from the battery whilst it is driving. This modifies the state update function as follows:

$$E_i = E_{i-1} + t_i^{stp} (\eta(p_i)_- - (p_i)_+ - p_i^d) \quad (4)$$

The final constraint is on the amount of energy stored in the battery. The house occupants provide an input parameter $E^m \in [0, \bar{E}]$ that represents the minimum energy that the EV battery should have in it at each point in time. This value represents how much energy the occupant expects to need if they drive away in the car at a given time. This is not a hard constraint as the draw from driving can bring the battery charge below this limit, but it ensures that if the battery power does fall below this limit, then it charges back up as fast as possible.

$$x_i^h = 1 \implies$$

$$E_i \geq \min [E_{i-1} + t_i^{stp} (-\eta \underline{p} - p_i^d), E_i^m] \quad (5)$$

Hot Water Heating. The hot water system is made up of a storage tank and an electric heating element. We ignore the details of the interaction between hot and cold water in the tank and consider the state of the tank as being the amount of energy $E \in [0, \bar{E}]$ it contains above the inlet cold water temperature. The tank is considered empty of hot water when this value is zero. The action variable is the power setting of the electric heater $p \in [\underline{p}, 0]$ at each time step. An amount of power, given by the input parameter $p^d \in \mathbb{R}_+$, is drawn from the tank at each time step in order to meet occupant demand. The energy state update function is given by:

$$E_i = E_{i-1} + t_i^{stp} (-p_i - p_i^d - p_i^l + p_i^u) \quad (6)$$

The variable $p^l \in \mathbb{R}_+$ represents thermal losses from the tank to the outdoor environment. The rate of loss depends on how full the tank is and the difference in temperature between the water set point $T^s \in \mathbb{R}$ and the outdoor temperature $T^o \in \mathbb{R}$ through a resistivity $R \in \mathbb{R}_+$:

$$p_i^l = \frac{1}{R} \frac{E_i}{\bar{E}} (T^s - T_i^o) \quad (7)$$

The variable $p^u \in \mathbb{R}_+$ is a recourse variable that is used to indicate the amount of hot water demand which goes unmet, i.e. water drawn from the tank when it is empty. This is heavily penalised as a cost c through an unmet demand price v :

$$c_i = v p_i^u \quad (8)$$

The hot water system has a minimum stored energy level $E^m \in [0, \bar{E}]$, much like the electric vehicle. If drawn water brings the energy level of the tank below this value then the heater must work as hard as possible to bring the energy back up. This value is fixed in time and is used to represent a safety margin that the occupants impose in order to reduce the likelihood of running out of cold water.

$$E_i \geq \min [E_{i-1} + t_i^{stp} (-\underline{p} - p_i^d - p_i^l + p_i^u), E^m] \quad (9)$$

Under Floor Heating/Cooling. The heating system of the house includes a heat pump which heats/cools water which is then pumped through piping embedded in the floor of the house. The temperatures of the floor and the air in the room $T^f, T^a \in \mathbb{R}$ are the device states. The action variable is the amount of thermal energy that is supplied to the floor of the house $p^t \in \mathbb{R}$. This is limited by the heat pump electrical power consumption $p \in [\underline{p}, 0]$ through heating and cooling Coefficients of Performance (COP) $\eta^h \in [\underline{\eta}^h, \bar{\eta}^h], \eta^c \in [\underline{\eta}^c, \bar{\eta}^c]$:

$$p_i = -\frac{1}{\eta_i^h} (p_i^t)_+ - \frac{1}{\eta_i^c} (p_i^t)_- \quad (10)$$

The COPs depend on the temperatures of the two thermal wells between which the heat pump is operating. We assume the internal thermal well is at a constant temperature, and the external well is at the outdoor temperature $T^o \in \mathbb{R}$. We have the COPs as linear functions of T^o for some constants $a^h, a^c \in \mathbb{R}_+$ and $b^h, b^c \in \mathbb{R}$, with hard upper and lower limits:

$$\eta_i^h = \min [\max [a^h T_i^o + b^h, \underline{\eta}^h], \bar{\eta}^h] \quad (11)$$

$$\eta_i^c = \min [\max [-a^c T_i^o + b^c, \underline{\eta}^c], \bar{\eta}^c] \quad (12)$$

Heat can transfer between the floor and the outdoor environment $p^{fo} \in \mathbb{R}$, the floor and the air in the room $p^{fa} \in \mathbb{R}$, and the air in the room and the outdoor environment $p^{ao} \in \mathbb{R}$. We use simple lumped thermal resistivities $R^{fo}, R^{fa}, R^{ao} \in \mathbb{R}_+$ to govern these heat flows:

$$p_i^{fo} = \frac{1}{R^{fo}} (T_i^f - T_i^o) \quad (13)$$

$$p_i^{fa} = \frac{1}{R^{fa}} (T_i^f - T_i^a) \quad (14)$$

$$p_i^{ao} = \frac{1}{R^{ao}} (T_i^a - T_i^o) \quad (15)$$

The temperature state update functions are given by:

$$T_i^f = T_{i-1}^f + \frac{t_i^{stp}}{m^f \kappa^f} (p_i^t - p_i^{fo} - p_i^{fa} + A^f I_i) \quad (16)$$

$$T_i^a = T_{i-1}^a + \frac{t_i^{stp}}{m^a \kappa^a} (p_i^{fa} - p_i^{ao} + p_i^g) \quad (17)$$

where $m^f, m^a, \kappa^f, \kappa^a \in \mathbb{R}_+$ are the floor and air, mass and specific heat capacity coefficients respectively. Sunlight enters through the windows at an irradiance of I_i , and lands on a floor area of A^f . The input $p_i^g \in \mathbb{R}_+$ is thermal power generated by occupant metabolisms and background electric appliances, that contributes to heating the air in the room.

The final relation we have is for the comfort cost c which depends on the difference between the air temperature and an occupant specified set point temperature $T^s \in \mathbb{R}$. Two occupant specified time varying comfort prices v^a, v^b are used, one of which is only included after a threshold temperature difference ΔT^b :

$$c_i = \begin{cases} v_i^a |T_i^a - T^s| & \text{if } |T_i^a - T^s| < \Delta T^b \\ (v_i^a + v_i^b) |T_i^a - T^s| & \text{otherwise} \end{cases} \quad (18)$$

Shiftable Loads. Shiftable loads are devices that need to run once within a time window. An occupant sets two input parameters: a start time i^s and a last allowed start time i^l , between which the controller must schedule the device to run. Examples of this kind of device include washing machines, clothes dryers and dish washers. We model non-preemptive shiftable loads which can have time varying power consumptions.

The start of run indicators $x \in \{0, 1\}$ act as both the device action and state variables. A shiftable load has a cumulative energy consumption function $\psi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ which takes a run duration and returns the cumulative amount of energy that the device has consumed for that duration. Constraints on the run indicator variables and the device power $p \in \mathbb{R}_-$ are given by:

$$\sum_{k=i^s}^{i^l} x_k = 1 \quad (19)$$

$$p_i = - \sum_{k=i^s}^i x_k \frac{\psi(t_i - t_{k-1}) - \psi(t_{i-1} - t_{k-1})}{t_i^{stp}} \quad (20)$$

Photovoltaics. The photovoltaic (PV) panels have no action variables, the amount of electricity they generate is purely determined by the solar irradiance input parameter. We model a PV system ignoring temperature and shading effects and by assuming the panels lay on a horizontal surface. The generated electric power $p \in \mathbb{R}_+$ is then a simple function of the panel area $A \in \mathbb{R}_+$, efficiency $\eta \in [0, 1]$ and global irradiance input parameter $I \in \mathbb{R}_+$:

$$p_i = \eta A I_i \quad (21)$$

3 Stochastic HAS Optimisation

So far we have considered the deterministic home automation formulation that requires perfect foresight about what will happen in the future. However in practice, almost all the input parameters are uncertain, and their uncertainty is only revealed in real time (e.g., outdoor temperature) or in some cases a few time steps in advance (e.g., RTP). This motivates the use of online stochastic optimisation (Van Hentenryck and Bent 2006), which exploits statistical models of the uncertain parameters in order to make the best decisions on average.

3.1 The Stochastic Model

In the stochastic HAS problem, the RTP v_i , background house power $p_{h,i}^b$ and device input parameters $r_{d,i}$ are random variables. We denote their real-world realisations (i.e. their values when the uncertainty is revealed) with the symbol $*$. For instance, T_i^{o*} denotes the real outdoor temperature at time step i . For notational convenience, all inputs are combined into one vector

$$z_i = (v_i, p_{h,i}^b, r_{d_1,i}, r_{d_2,i}, \dots)^T \quad (22)$$

where we index elements with a k (e.g., $z_{i,k}$). Random variables at time step i may be dependent on each other and on the variables at previous time steps. Therefore the joint distribution for random variables up to time step i is given by:

$$P(z_i, z_{i-1}, \dots) \quad (23)$$

Let t^* represent the current real world time. Each input $z_{i,k}$ is revealed a fixed amount of time $\Delta t_k^{rev} \in \mathbb{R}_+$ in advance (or in real time if $\Delta t_k^{rev} = 0$). This means, that for a given t^* an input $z_{i,k}$ is known to be $z_{i,k}^*$ if $t_i \leq t^* + \Delta t_k^{rev}$, otherwise it is a random variable. Given i and t^* we use $K_{i,t^*} = \{k | t_i \leq t^* + \Delta t_k^{rev}\}$ to denote the set of known input indices.

3.2 Online Stochastic Optimisation

In an online stochastic algorithm decisions are made one step at a time using stochastic information about future events. After each time step the uncertainty and the effect of all actions is revealed, updating the state of the system. Decisions for the next period are computed and the process is repeated. Online stochastic optimisation has been used successfully on a wide variety of problems (e.g., Powell, Simao, and Bouzaiene-Ayari 2012; Van Hentenryck and Bent 2006).

Our algorithms use a rolling finite horizon as illustrated in Figure 1, where the time steps $1, \dots, n$ are aligned to each

horizon with $t_0 = t^*$. Optimisation is performed within each horizon using stochastic information for any unrevealed inputs, and then the actions for the first time step are executed in the real world.

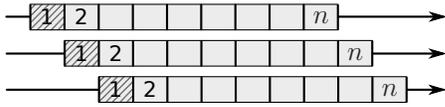


Figure 1: Rolling horizon for 3 consecutive iterations.

It might not be possible to execute actions produced by the optimisation if the real world input parameters z_1^* differ from what the optimisation anticipated. For example, if the optimisation decides to run the hot water heater at full power, and the tank unexpectedly reaches its capacity (due to less demand for hot water than expected), then the power of the heating action will need to be reduced so as to remain within the tank's capacity. Our HAS handles this automatically in the execution step, by using very simple executives for each device which select the closest feasible action.

In the following sections we introduce two approaches to solving the stochastic optimisation problem within each horizon: the expectation and the 2-stage algorithms.

3.3 Expectation Formulation

The expectation online stochastic algorithm takes the conditional expected value of any unrevealed inputs in the optimisation horizon, and solves the deterministic version of the problem given in Definition 3. We use the term expected value loosely because in truth we calculate the expected value only where it makes sense, which is typically for continuous inputs. For the rest of the inputs the most likely value is calculated instead. For example, expected value is used for outdoor temperatures and most likely value for the washing machine requests. Both of these calculations are performed using the joint distribution for inputs in the horizon, conditioned on any known inputs in and prior to the horizon:

$$P(z_n, z_{n-1}, \dots, z_1 | (z_{n,k}^*, \forall k \in K_{n,t_0}), \dots, (z_{1,k}^*, \forall k \in K_{1,t_0}), z_0^*, \dots) \quad (24)$$

3.4 2-Stage Formulation

In this algorithm 2-stage stochastic programming is used within each horizon. This provides an approximation to a full multi-stage stochastic program which are, in general, known to be extremely challenging computationally (Shapiro 2006). The first stage includes time step 1, and the second stage time steps $2, \dots, n$. Traditionally, in 2-stage stochastic programming there is no uncertainty in the first stage (Shapiro, Dentcheva, and Ruszczyński 2009). However in our problem, we are required to make decisions before some inputs in the first stage are revealed. To resolve this, first stage inputs are set to their real values if known, otherwise their conditional expected value is taken (as described in 3.3).

The second stage uses sampled scenarios to represent the uncertainty in the input parameters. We define a second

stage scenario s as being a sample from the joint distribution of random variables in the second stage, conditioned on any revealed inputs in the second stage, and inputs in and prior to the first stage:

$$s \sim P(z_n, z_{n-1}, \dots, z_2 | (z_{n,k}^*, \forall k \in K_{n,t_0}), \dots, (z_{2,k}^*, \forall k \in K_{2,t_0}), z_1, z_0^*, \dots) \quad (25)$$

We use the Sample Average Approximation (SAA) (Shapiro, Dentcheva, and Ruszczyński 2009) to limit the number of scenarios $S \in \mathbb{N}$ that we need to consider in the second stage. Each scenario in the second stage needs to have its own set of variables in the optimisation problem. For example we denote the power of device d at time step i in scenario s by $p_{d,i,s}$. The 2-stage objective function is given by:

$$\min \left[c_1 + \frac{1}{S} \sum_{s \in \{s_1, \dots, s_S\}} \sum_{i=2}^n c_{i,s} \right] \quad (26)$$

3.5 Stochastic Inputs

Stochastic inputs include the real-time pricing (RTP), outdoor temperature, solar irradiance, background power, internal heat generation, hot water demand, EV usage and shiftable load requests. Accurately modelling any of these random processes is a significant undertaking in itself. The models we developed, while not the most sophisticated, suit the purposes of our experiments by capturing the fundamental nature of these stochastic processes. We investigated a number of different model types before settling on Generalised Additive Models (GAM) (Hastie and Tibshirani 1990) for the continuous variables like temperature, and Markov Models for the more discrete occupant driven behaviours such as shiftable device requests.

Generalised Additive Models. In order to predict future values, the GAMs models take advantage of weather forecasts that can be readily obtained from national weather services. These forecast values include daily maximum and minimum temperatures, as well as morning and afternoon cloud cover and wind speed. They also take in the value from the previous time step and temporal information. The models were trained on data obtained from the Bureau of Meteorology³ and Australian Energy Market Operator⁴ relevant to the states of New South Wales (NSW) and the Australian Capital Territory (ACT) in Australia.

The best way of implementing RTP in retail markets is still an open question and so is worth particular mention. It is unlikely that it will be a simple replication of the wholesale spot market price due to its high volatility. More likely it will be controlled by the retailer, but have a shape representative of the wholesale market. We designed our RTP to be a quadratic function⁵ of the amount of power that fossil

³Bureau of Meteorology, www.bom.gov.au

⁴Australian Energy Market Operator, www.aemo.com.au

⁵The quadratic is representative of an increasing marginal supply price (Ramchurn et al. 2011).

fuel sources must supply to meet total network load. This is the total network demand minus the generation from renewable sources such as wind and solar. We used a GAM for the total network demand. The generation from renewables is a function of wind speed and solar irradiance. The RTP is only revealed to a house 30 minutes in advance.

Markov Models. Input parameters arising from the behaviour of house occupants were captured via semi-Markov models representing the activities and consumption patterns (e.g., hot water, shiftable load requests and EV usage) of the four occupants of a specific house in the ACT. Each model identifies the key activities of an occupant (e.g., sleeping, taking a shower and leaving for work), and specifies the probabilities of transiting from one activity to the next within certain time windows. Each activity is associated with a series of actions (e.g., watching TV, requesting the dish washer to operate) that trigger changes in input parameters. Conditional sampling through these models is used to generate scenarios.

Whilst this scheme was convenient for our experiments, other more data-driven options are possible: we could simply gather and use a database of raw scenarios, or learn model parameters from disaggregated demand data (Kolter, Batra, and Ng 2010; Parson et al. 2012).

4 Experiments

We implemented the 2-stage and expectation online algorithms using Gurobi as a backend to solve the MILP within each horizon. The devices in Section 2.2 were implemented and included in the experimental house, and conditional samplers were created for the uncertain input parameters in Section 3.5. We created a simple simulator that uses the same physical equations as the optimisation to simulate the execution of actions in the real world. We compare the performance of the 2-stage and expectation controllers with naive and smart reactive controllers, and a controller that has perfect information.

The *Naive* reactive controller represents a household that either has no ability or no desire to respond to a RTP. It starts shiftable devices as soon as a request is received, fills up the hot water tank in off-peak hours, charges the EV only if it is below the requested minimum level, maintains the room at the set point temperature and never uses the battery bank.

The *Smart* reactive controller uses simple device action policies to decide how to respond to changes in RTP. It delays running a shiftable device until it reaches a cheap price or the last available start time, uses thresholds about a moving average of the RTP to decide when to charge or discharge energy from the batteries, EV and hot water system, and maintains the room at the set point temperature like the naive controller.

The *Perfect* controller has perfect foresight about what will happen in the future. It optimises the deterministic problem in Definition 3 over the whole experiment duration with full knowledge of z_i^* . This controller (which is infeasible in practice) is used to give a lower bound on the objective that can be achieved by the other controllers.

4.1 Controller Comparison

Nine sets of input parameters typical for the month of February were generated. These were used in 9 separate experimental runs, each with a duration of 7 days. The online algorithms had 16 hour optimisation horizons, with 15 minutes for the first two time steps and 30 minute time steps for the remainder of the horizon⁶. The reactive and perfect controllers had 15 minute time steps. The 2-stage algorithm sampled 30 scenarios in its second stages.

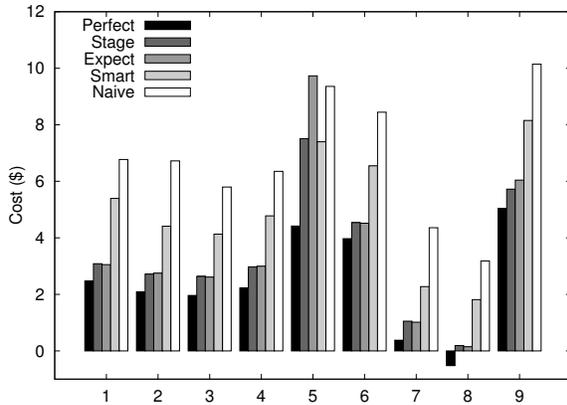
The controller costs are plotted in Figure 2a for each of the 9 experimental runs. These results are adjusted to account for any energy that remains in the battery, EV, or hot water system at the end of an experimental run. This is done by valuing the left-over energy at the average RTP for the last 24 hours. Without this adjustment it would not be fair to compare the controllers, since any controller that anticipates the need to store energy for a future purpose, would perform poorly if it does so just before the experiment ends. This is an artefact of the finite length of our experimental runs; with very long durations this problem goes away as the costs associated with left-over stored energy become insignificant.

We see that the 2-stage and expectation algorithms get quite close to the performance of the controller with perfect foresight. They produce significant cost reductions over the two reactive controllers. In run 5 the performance of the 2-stage and expectation controllers is drastically reduced. This is caused by large spikes in hot water consumption that occur during this run. The expectation algorithm fails to account for these spikes. The 2-stage anticipates them, but their effects are heavily discounted because of their low likelihood. When the spikes eventuate both algorithms are hit by a large cost for not being able to supply enough hot water. The 2-stage fares better because it has done some preparation, which is the amount that on average will give the optimal outcome. The two reactive controllers are not significantly affected by these unlikely spikes because they have to plan for the worst case and always keep the tank relatively full.

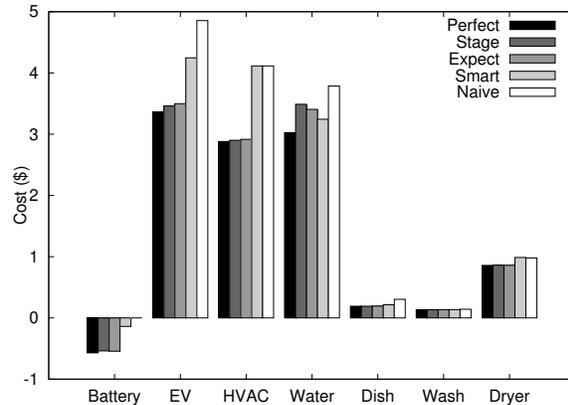
The expectation algorithm does just as well or better than the 2-stage controller for a majority of the experimental runs, but on average the 2-stage outperforms it by 7.4%. The expectation algorithm performs poorly when it fails to anticipate an important scenario, such as the possibility for large spikes in hot water consumption. The amount of time spent optimising in Gurobi per day is on average 64 seconds for the 2-stage algorithm and 0.9 seconds for the expectation algorithm (using a single core of an Intel i7-2600 3.4GHz CPU). Whilst the 2-stage is much slower than the expectation algorithm, its computational time remains insignificant when spread out over a day.

Figure 2b shows the costs for each controller in run 1 disaggregated into those associated with each device, ignoring the PV. We see cost savings for all devices when using the online stochastic controllers over the reactive ones. The greatest area where the 2-stage and expectation controllers can be improved is in hot water heating; in this experiment

⁶By using larger time steps for more uncertain values further into the future we reduce the computational burden with only a minor reduction to solution quality.



(a) Controller comparison.



(b) Disaggregated costs.

Figure 2: Costs for each experimental run and disaggregated costs for run 1.

the smart reactive controller is even outperforming them. Aside from this their performance is close to that of the perfect controller.

Figure 3a gives an example of the power exchanged between the house and the grid for one day, along with the RTP. As expected most consumption occurs when the prices are low, and when the price is high power is sold back to the grid from the battery, EV and PV. The 2-stage and expectation controllers follow the general trend of the perfect controller with some small divergences.

4.2 Parameter Tuning

Figure 3b shows the results of an experiment where we investigated how performance changes with the horizon duration. This plot shows the performance of the perfect controller running as an online algorithm where it is restricted to only having perfect foresight a certain distance into the future. The experiment is performed on run 1 for a number of different horizon durations, and the results are compared to the original perfect controller that can see the full 7 days. The results show that there is little to be gained by looking any further into the future than 20 hours.

We reran the 2-stage experiments 3 times using different starting sampler seeds. The standard deviation of the results was typically less than 2 cents for each run except run 5 where we had a standard deviation of 42 cents. This seems to suggest the 2-stage algorithm could have benefited from more samples in its second stage for run 5. Indeed initial experiments show some improvements when using more samples but more comprehensive testing is still required.

5 Related Work

Much of the existing literature on residential demand response focuses on deterministic formulations over fixed horizons where the scheduler has perfect foresight (Ramchurn et al. 2011; Gatsis and Giannakis 2012). Those that have considered uncertainty in the problem typically focus on just one aspect (e.g., real-time pricing) (Mohsenian-Rad

and Leon-Garcia 2010), or use very simple models for random variables (Tischer and Verbic 2011).

Model-predictive control has been used to account for the uncertainty of estimated device model parameters and measurement noise (Yu et al. 2012), but not the uncertainty of the type we model. In general, model-predictive control is best suited to unconstrained, purely continuous settings with limited uncertainty.

Dynamic programming (Tischer and Verbic 2011; Kim and Poor 2011) and Q-learning (Levorato, Goldsmith, and Mitra 2010) have been used in conjunction with Markov Decision Process (MDP) formulations of the residential load scheduling problem, to generate policies that allocate power to each device. MDP approaches suffer from severe scalability issues, especially since the state space needs to be discretised. Moreover, MDPs seem somewhat excessive for our problem, given that uncertainty does not depend on the decisions taken. Our stochastic programming approach which uses scenario sampling is more scalable and more natural in the presence of exogenous uncertainty.

One paper (Tischer and Verbic 2011) found that acting on the basis of the optimal dynamic programming solution did not provide any benefit over acting on expectations. For many of our experiments we found this to be true, but we did identify certain cases where the more cautious nature of the 2-stage algorithm is superior. The difference in our two results is thought to be due to our use of more sophisticated uncertainty models and differences in device models.

The paper closest to ours compares two-stage stochastic programming and robust optimisation techniques for scheduling residential loads (Chen, Wu, and Fu 2012). Uncertainty is restricted to the RTP which is known for the first stage but becomes uncertain thereafter. The objective includes minimising expected price and the probability mass of “risky” scenarios whose price exceeds a certain threshold. Comfort is handled by imposing hard constraints under which appliances must run, rather than by inclusion into the objective. In this setting, two-stage stochastic programming

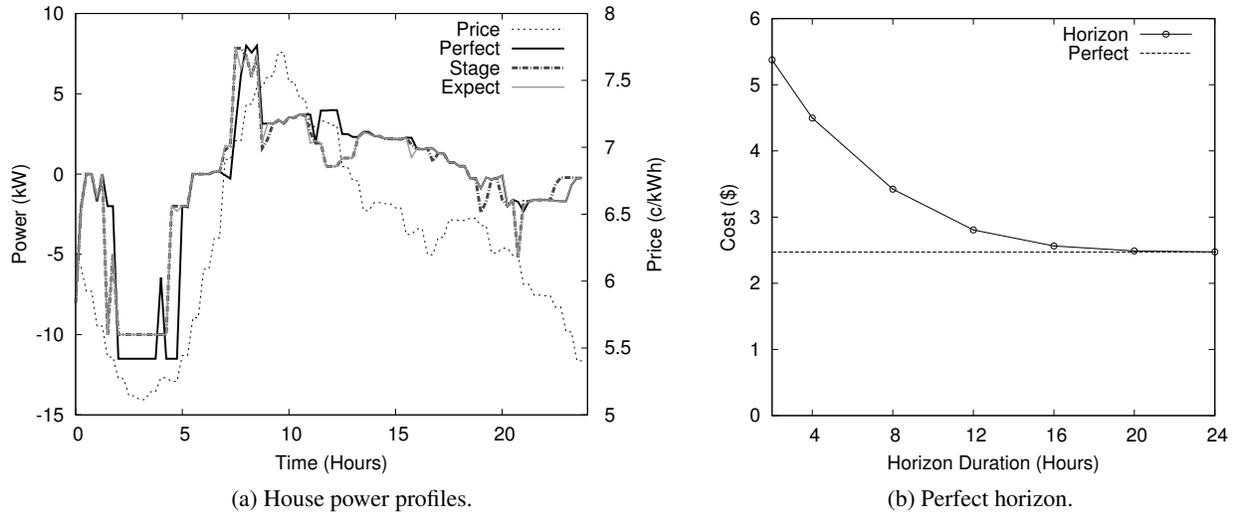


Figure 3: House power profiles over one day and performance for different horizon durations.

was observed to provide benefits over robust scheduling.

The scope of our analysis goes significantly beyond these results, by exploring uncertainty from a large range of sources and by identifying the value of stochastic information. We enable richer sources of uncertainty to be considered in our framework, by allowing inputs to be revealed at arbitrary points in time.

6 Conclusion and Future Work

This paper contributes to the growing body of work on residential control of loads and storage under real-time pricing, by developing a framework that accounts for uncertainty. To our knowledge, it is the first work that provides a scalable and accurate solution in the presence of uncertainty about future prices, occupant behaviour and environmental conditions. Using models representative of physical devices and random processes, we have shown the monetary and comfort cost savings that can be achieved by using online stochastic algorithms over reactive control, and the performance increase of a 2-stage approach over acting on expectations. Studies such as the one in this paper are import for rallying industry and customers towards more effective energy management schemes.

Further research is needed to investigate how closely reality can be modelled with random processes, and if in turn they are suitable for online learning. We also need to further investigate how time step sizes and the number of second stage scenarios influence performance, and to conduct more experiments for different months of the year. The experimental set up we have developed can be used to experiment with and compare different pricing schemes. For example, time of use pricing and schemes where the price offered for generation is lower than that for consumption. We also plan on investigating how multiple houses react to a RTP and what sort of emergent behaviour develops when they are all learning their statistical models online.

Commercially available residential DR solutions⁷ typically focus on direct load control or simple reactive policies. Such systems could experience more optimal DR performance and greater residential customer satisfaction by using our algorithms. However, there are two practical challenges that need to be addressed before our technology can have widespread adoption in homes. Firstly there needs to be a standard smart device interface for communication with HASs. The proposed Australian demand response appliance standards in AS4755 are a move in the right direction, but true two-way communication would enable better control outcomes. All new appliances should come with such an interface, while some existing appliances could be retrofit. The second requirement is for utilities to start offering RTP services to residential customers. The HAS could receive the RTP signal directly⁸ or use existing smart meters as an intermediate. Customers will be attracted to these schemes and to the purchase of a HAS with DR capabilities by the potential electricity bill savings.

Acknowledgements

This work is supported by NICTA's Optimisation Research Group as part of the Future Energy Systems project. We thank our project members and reviewers for useful discussions and helpful suggestions. NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

⁷e.g., comverge: www.comverge.com, nest: www.nest.com and Cooper Power Systems: www.cooperindustries.com

⁸e.g., using openADR Alliance protocols: www.openadr.org

References

- Chen, Z.; Wu, L.; and Fu, Y. 2012. Real-time price-based demand response management for residential appliances via stochastic optimization and robust optimization. *Smart Grid, IEEE Transactions on* 3(4):1822–1831.
- Gatsis, N., and Giannakis, G. B. 2012. Residential load control: Distributed scheduling and convergence with lost ami messages. *Smart Grid, IEEE Transactions on* PP(99):1–17.
- Hastie, T., and Tibshirani, R. 1990. *Generalized Additive Models*. Monographs on Statistics and Applied Probability Series. Chapman & Hall, CRC Press.
- Kim, T., and Poor, H. 2011. Scheduling power consumption with price uncertainty. *Smart Grid, IEEE Transactions on* 2(3):519–527.
- Kolter, J. Z.; Batra, S.; and Ng, A. Y. 2010. Energy disaggregation via discriminative sparse coding. In *24th Annual Conference on Neural Information Processing Systems (NIPS)*, 1153–1161.
- Levorato, M.; Goldsmith, A.; and Mitra, U. 2010. Residential demand response using reinforcement learning. In *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, 409–414.
- Mohsenian-Rad, A.-H., and Leon-Garcia, A. 2010. Optimal residential load control with price prediction in real-time electricity pricing environments. *Smart Grid, IEEE Transactions on* 1(2):120–133.
- Parson, O.; Ghosh, S.; Weal, M.; and Rogers, A. 2012. Non-intrusive load monitoring using prior models of general appliance types. In *Proceedings of Twenty-Sixth Conference on Artificial Intelligence (AAAI-12)*.
- Powell, W.; Simao, H.; and Bouzaiene-Ayari, B. 2012. Approximate dynamic programming in transportation and logistics: a unified framework. *EURO Journal on Transportation and Logistics* 1(3):237–284.
- Ramchurn, S. D.; Vytelingum, P.; Rogers, A.; and Jennings, N. R. 2011. Agent-based control for decentralised demand side management in the smart grid. In Tumer; Yolum; Sonnenberg; and Stone., eds., *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems - Innovative Applications Track (AAMAS 2011)*, 330–331.
- Shapiro, A.; Dentcheva, D.; and Ruszczyński, A. 2009. *Lectures on Stochastic Programming: Modeling and Theory*. MPS-SIAM series on optimization. Society for Industrial and Applied Mathematics (SIAM, 3600 Market Street, Floor 6, Philadelphia, PA 19104).
- Shapiro, A. 2006. On Complexity of Multistage Stochastic Programs. *Operations Research Letters* 34(1):1–8.
- Tischer, H., and Verbic, G. 2011. Towards a smart home energy management system - a dynamic programming approach. In *Innovative Smart Grid Technologies Asia (ISGT), 2011 IEEE PES*, 1–7.
- Van Hentenryck, P., and Bent, R. 2006. *Online Stochastic Combinatorial Optimization*. Cambridge, Mass.: The MIT Press.
- Yu, Z.; McLaughlin, L.; Jia, L.; Murphy-Hoye, M. C.; Pratt, A.; and Tong, L. 2012. Modeling and stochastic control for home energy management. In *2012 Power and Energy Society general meeting*.