Design and Evaluation of a Constraint-based Energy Saving and Scheduling Recommender System

Seán Óg Murphy[12], Òscar Manzano[12], and Kenneth N. Brown[12]

[1]International Energy Research Centre, Cork, Ireland
[2]Insight Centre for Data Analytics, Department of Computer Science, University College Cork, Ireland.
{seanog.murphy, oscar.manzano, ken.brown}@insight-centre.org
http://www.insight-centre.org

Abstract. Development of low-cost and inter-operable home sensing products in recent years has motivated the development of consumer-level energy and home monitoring software solutions to exploit these new streams of data available to end-users. In particular, this opens up the home energy space as an area of high potential for the use of consumer-level energy optimisation with home-owners actively engaged with data about their energy use behaviour. We describe the development of a tablet-based home energy cost saving and appliance scheduling system which calculates behaviour change suggestions that save occupants on their energy bills while avoiding disruption to their regular routines. This system uses a Constraint Satisfaction Problem Solver to compute savings based on real-world sensor data, and to generate revised schedules in a user-friendly format, operating within a limited computing environment and achieving fast computation times.

1 Introduction

In-home networked sensing has, in recent years, reached a level of cost, reliability and protocol maturity that enables consumer-level home sensing to a degree not possible before. The increasing variety and availability of Zigbee-compatible wireless sensing allows for rich sensor coverage of homes with minimal impact or installation, providing rich data streams from around the home, monitoring use of space, appliance use, temperature and more. Providing these data streams in an easy to understand format is of substantial benefit to home-owners to help understand inefficient heating of spaces, electricity efficiency and overall energy costs [24].

Furthermore, developments in energy infrastructure and generation have lead to a greater variety of energy production sources and tariff schemes, leading to a greater incentive (both in terms of cost and in environmental impact) for home-owners to take pay attention to their energy consumption over different time periods. Some regions use variable tariffs based on the time of day [25], and some schemes feature time-variable peak-demand limits with punitive charges incurred where total concurrent energy use exceeds a limit at a particular time of the week [1].

In this work we describe a Constraint-based system that recommends adjustments to energy use patterns to provide targetted cost savings in the home, based on historical and current observation of appliance use patterns in that home. The data is obtained from a deployed sensor network and energy monitoring system. The adjusted appliance use pattern accepted by the home-owner then forms the basis of a appliance use scheduling application, which recommends a set of activation times that correspond to observed historical behaviour while respecting limits of overall capacity and variable demand pricing schemes. This system operates on a low-power Android tablet device, as part of a Reporting Tool app which provides robust, easy-to-understand feedback to home-owners from the data streams provided by the home sensor deployment.

2 AUTHENTIC Smart home project

AUTHENTIC is a multi-institutional initiative to design and deliver a Home Area Network (HAN) infrastructure capable of supporting opportunistic decision making for effective energy management within the home.

The Authentic HAN consists of a collection of Zigbee wireless sensors communicating over a multi-hop wireless network with a low-power gateway Linux PC. This hub stores the sensor readings in a MySQL database. The readings are retrievable through a REST (Representational State Transfer) interface [19] in the form of JSON (Javascript Object Notation) Objects [6] which are transformed and can be displayed to the users via a reporting tool (discussed below). The first phase of deployments included smartplug sensors attached to power sockets which report appliance activations and energy consumption, passive infrared motion detectors reporting space occupancy, temperature, light and humidity sensors, and contact sensors reporting when doors and windows are opened or closed (Figure 1). Future phases include wireless smart meters for measuring total electricity, gas and oil consumption, sensors for monitoring the use and settings of heating, ventilation, air conditioning and water heating, and sensors and meters for monitoring the available of energy from renewable sources.

In order to assist the users to monitor and change their energy consumption, and to track their energy bills, a Reporting Tool app has been developed for Android tablets. The app provides direct and instant feedback, converting complex sensed data into customisable human-readable reports; sample screenshots are shown in Figure 2. Responding to user queries for specified time periods, the app can create visualisation of, for example, total energy consumption, consumption by room, consumption by appliance, room occupancy patterns, correlations between room occupancy and temperature or appliance use, and comparisons between different appliances and different time periods. The app includes a facility for real-time alerts when specified or anomalous patterns occur (e.g. windows opening in a heated room, or instantaneous consumption above a threshold). Fi-

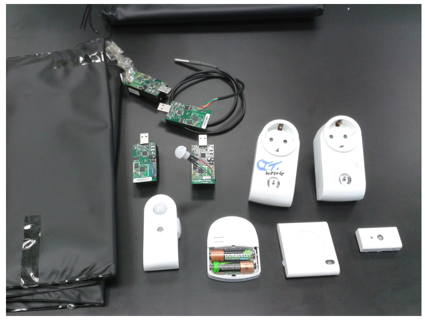

Fig. 1. AUTHENTIC sensors (movement, light, humidty, contact and smartplug sensors)

nally, the app provides facilities for guiding the user to change their behaviour. At a basic level, the app allows the user to set annual energy goals, and to track progress towards those goals. More sophisticated feedback, including advice on how to achieve specified reductions, is the main topic of this paper, and is discussed in the remaining sections.

3 Requirements

Users need guidance to help change their energy consumption patterns in accordance with their own goals and preferences. We consider two high level goals: reducing cost, and reducing total energy consumption[1]. The first steps are understanding how much energy is consumed by each activity and what cost is incurred by those activities at different times. Reducing consumption itself is easy – appliances can simply be switched off – but the difficulty is in balancing the reduction in consumption with the users' preferences for safe and comfortable living. Lowering a heating thermostat by 2 degrees may be unsafe in winter, while switching off an entertainment console may be impractical.

The aim is to propose changes to a user's behaviour which achieve their energy goals, while still being acceptable for their health, comfort and enjoyment. Further, those proposed changes should be personal, tailored to the users in a specific home. Thirdly, the interaction load imposed on the users must be manageable: systems which require extensive elicitation of preferences and utility functions under different scenarios before any proposal can be made are simply impractical in most domestic settings. Finally, the system should respect the privacy of the users, and should be capable of operating entirely within the home, on devices with limited computing power, without releasing data or relying on data obtained from other users outside the home.

[1] Additional goals of reducing carbon footprint and reducing energy consumption from non-renewable sources are being developed.

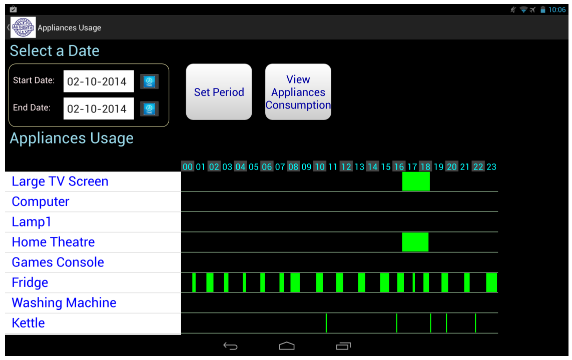

(a) Appliance Use History

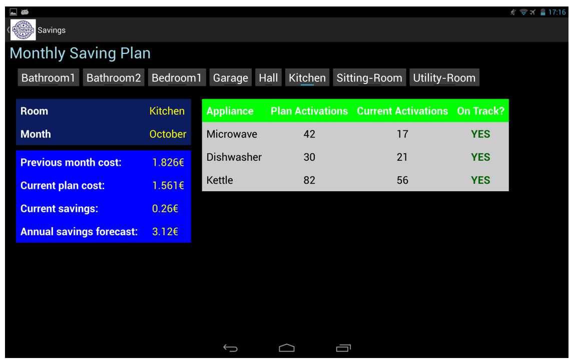

(b) Plan Recommendation (number of appliance activations)

Fig. 2. Android reporting tool.

Our intention is to guide the users based on the observed behaviour of those users in their home. The data streams from the sensors and smart plugs provide a history of appliance use, from which we can extract daily and weekly patterns. Given a goal (e.g. reduce the energy bill by 10%), we will then search for minimal changes to those patterns which satisfy the goal. Our proposals will be on two levels. The first level will recommend the total number (or duration) of activations for each appliance over a specified period, and if differential tariffs are in use, the total number of appliance activations in each price band.

The second level will propose a schedule of use for each day, ensuring the total simultaneous load is within any threshold, again representing a minimal change to observed behaviour patterns. Thus in each case an initial proposal is tailored to each home, without requiring any interactive preference elicitation apart from setting the initial high level goal. In each case, the users should then be able to interact with the system, imposing new constraints or objectives if the proposal is not satisfactory, and requesting a new recommendation. Further, the system will then monitor energy consumption within the home, alerting the user when targets are not going to be met, and offering recomputation of the schedules or activation levels.

Our problem thus involves multiple levels of decision support and optimisation, over constrained variables. The problems range from relatively simple linear assignment problems to constrained scheduling, and in each case the problems may be extended by additional user constraints. Thus we choose to use constraint programming as the single technology for all problem variants because of the dedicated support for modelling cumulative scheduling problems and because of its flexibility in adding what ultimately could be arbitrary side constraints and preferences. The Authentic system also imposes some technical limitations, which influence our choice of solver. The Reporting Tool app (Section 2) runs under Android, and is limited to JDK 1.6 for library compatibility issues. Thus we model and solve the decision and optimisation problems using the Choco 2.1.5 Java library, as it is also JDK1.6 compatible. This places a restriction on the variable types and constraints we use.

4 Constraint-based Recommender System

4.1 System Overview

The overall system is structured as follows. The Authentic HAN produces a database of sensor readings for the home. This data is transformed into appliance activation reports. Information is presented to the user, who specifies an energy or cost goal for the next period. From the appliance reports, the system constructs a constraint optimisation model, and proposes a high-level energy consumption plan to the user. It iterates with the user until an acceptable plan is agreed. From the high-level plan, the system then creates a scheduling problem and proposes a schedule to the user that respects the plan and requires a minimal change to previous behaviour. Again, the user iterates with the system until an acceptable schedule is agreed. The user is in control of all activations, and the system monitors compliance, alerts the user when targets are to be missed, and propose new plans or schedules in response. We describe the constraint-based modules in Sections 4 and 5, and evaluate performance in Section 6.

4.2 Data Preprocessing (Data Analysis Module)

To pre-process the sensor data, we developed a Data Analysis module which interfaces with the Authentic RESTful API [2] [18] to retrieve home sensor data which it discretises into forms more amenable to analysis and optimisation.
Sensor readings and events are retrieved in the form of JSON Objects which are then parsed by the module. By parsing the sensor information, data manipulation can be performed to convert this information into alternative forms more suitable for optimisation problems or efficiency analysis. In general, the facility is provided to convert Occupancy, Temperature, Light, Humidity and Appliance

Activation information into arrays of samples (at a user-configurable sampling rate). This allows for the like-for-like comparison of readings between arbitrary timepoints and aribtrary sensing types (e.g. plotting Occupancy vs Television, Light vs Lamp, or Humidity vs Dryer).
An additional function provided by this module is the generation of "Activation-Tuples", data objects representing the usage pattern for appliances in the home which we use as input for the Appliance Use Optimisation Module (Section 4.3).

ActivationTuples contain the following information:

- Appliance Name
- Average Consumption Per Activation (Watt-Hours)
- Number of Activations per Tariff (in the case of Day/Night or other variable tariffs)
- Power Draw Profile (time series, Watts)

Activations are considered from the time an appliance sensor reports instantaneous demand above a minimum threshold until the instantaneous demand returns to below that threshold. The average consumption is based on the watt-hour (WH) consumption per activation period. In the case of appliances with variable activation profiles (for example, different washing machine settings), we can subdivide the appliance into seperate sub-appliances, one for each setting or pattern. The Power Draw Profile represents the instantaneous power demand in watts over a time series with configurable sampling rate (e.g. 5 minutes). The Power Draw Profile is used with the Scheduling Module to ensure that the total concurrent power load in a home or on a circuit at a given time can be accounted for in the constraint satisfaction problem. As the reporting interval of the smart-plug sensors is generally 5 minutes, higher resolution sampling rates are not required.

4.3 Appliance Use Optimisation module (Solver)

The aim of the Appliance Use Optimisation module ("Solver") is to propose high-level patterns of appliance use which will achieve desired cost or energy savings while minimising disruption to the household. These patterns are simply counts of each appliance use in each price tariff. The module retrieves usage patterns from the analysis module in the form of activation tuples, and uses these to create a constraint satisfaction or optimisation problem. In this work we consider three tariffs (high, medium, low) available each day, although the solver is flexible and any number of price bands can be represented. The solver then searches for a modified pattern of use, by moving activations to lower price bands or reducing the number of activations, with a constraint on total cost. To avoid disruption, and to promote acceptance of the plan, we include an objective to minimise the deviation between the historical activation use pattern and the proposed pattern, represented as the sum of the squared differences of the appliance use counts). The constraint model is shown in Tables 1 and 2.

The only input required from the user for the first solution is to specify the required energy or cost reduction, and thus the interaction load on the user is minimal. However, it is likely that the proposed changes may not satisfy the user's requirements. Thus after the first schedule is proposed, the user is invited to specify additional constraints. The main type of user constraint is a domain restriction, specifying for a single appliance the acceptable range of activations in any price band or in total. Additional constraints, including for example, that the total number of activations for one appliance must be not less than the number for another can easily be handled by the solver, but are not yet implemented in the interface. The intention is to regard these user constraints as critiques of the proposed solutions, and thus the user is not required to specify any preference that is already satisfied. However, the added user constraints are stored, and will be applied to future solving instances (and the user is invited to remove any constraints that no longer apply).

5 Scheduling Module

The output from the optimisation module is a high level plan specifying the number of activations of each appliance in each price band over a specified period. The user is in control, and will choose when to activate an appliance. However, as discussed above, the system will monitor the activations and alert the user if the observed activations are not on track to meet the goals at the end of the period. In cases where there are multiple price bands during a day, or where there are peak power thresholds which invove a higher price, it may be difficult for the user to balance their appliance use appropriately. Thus, we also provide a task scheduler, which recommends start times for appliance use over a number of days in order to respect the cost constraints. As with the high level plan, there is an objective to find appliance start times which are as close as possible to the observed historical patterns for the user. The system allows users to modify aspects of the schedule, by adding or modifying tasks and constraints. The user is in control, but as before, the system will monitor appliance use compared to the schedule and alert the user when thresholds are likely to be breached, and offer the option of rescheduling.

The first aim is to generate a schedule which respects the cost constraints. The number of activations for each appliance is obtained from the high-level plan, and the required number of task instances are generated. In order to handle the time granularity for different constraints, each task instance is then broken up into multiple sub-tasks each of fixed duration, and sequencing constraints are applied between the subtasks. Disjunctive constraints[8] are imposed on multiple instance of the same appliance, to ensure that they cannot be scheduled to operate simultaneously. The high level plan also specifies the number of activations to take place under each tariff, and so for each task instance, we designate its highest permissible tariff and impose constraints preventing that task running

Variable Name	Description
	Constants
$InputCost$	Cost of the historical period
$Target$	Reduction in cost required (e.g. 0.9 of the previous cost)
H, L, M	Price per unit(wH) of electricity at High, Medium and Low tariffs
$A_iInputAct$	Original Number of activations of Appliance i in input
A_iCons	Average Consumption of Appliance i (wH) per activation
$A_iInputH$	Original number of activations at High tariff for Appliance i in input
$A_iInputM$	Original number of activations at Medium tariff for Appliance i in input
$A_iInputL$	Original number of activations at Low tariff for Appliance i in input
	Variables
A_iHAct	Number of activations of Appliance i at High tariff times
A_iMAct	Number of activations of Appliance i at Medium tariff times
A_iLAct	Number of activations of Appliance i at Low tariff times
	Auxilliary Variables
A_iHDiff	The difference between historical High use and the new plan for Appliance i
A_iMDiff	The difference between historical Medium use and the new plan for Appliance i
A_iLDiff	The difference between historical Low use and the new plan for Appliance i
A_iHCost	Total Cost for Appliance i at High tariff
A_iMCost	Total Cost for Appliance i at Medium tariff
A_iLCost	Total Cost for Appliance i at Low tariff
$A_iTotalCost$	Total cost for Appliance i
$A_iTotalDiff$	Sum of the squared differences for Appliance i
$TotalCost$	Total cost of the new plan
TotalDiff	Objective variable. Sum of squared differences
A_iAct	Total Number of activations of Appliance i

Table 1. Appliance Use Optimisation variables

in a higher tariff's time period. Note that this is a permissive approach, allowing the solver to find schedules with a lower cost than expected.

The next constraint to be considered is the peak power threshold. We handle this by imposing a cumulative constraint [3] over all tasks. We associate heights with each appliance task relative to their power demand (and in cases where the power demand varies over the activation cycle, the heights vary across the different sub-tasks). The cumulative constraint height parameter is set to to the maximum power threshold, and in the case of variable maximum demand schemes we create dummy tasks with appropriate heights to represent temporary lower thresholds over appropriate time periods.

Constraints	**Description**
$A_iHCost = A_iHAct * H$	Set the cost at H for Appliance i
$A_iMCost = A_iHAct * M$	Set the cost at M for Appliance i
$A_iLCost = A_iHAct * L$	Set the cost at L for Appliance i
$A_iHDiff = (A_iHAct - A_iInputH)^2$	Set the Squared Difference in H.
$A_iMDiff = (A_iMAct - A_iInputM)^2$	Set the cost at Squared Difference in M
$A_iLDiff = (A_iLAct - A_iInputL)^2$	Set the cost at Squared Difference in L
$A_iTotalDiff = A_iHDiff + A_iMDiff + A_iLDiff$	Set the sum Difference for Appliance i
$A_iTotalCost = A_iHCost + A_iMCost + A_iLCost$	Set the total cost for Appliance i
$A_iAct = A_iHAct + A_iMAct + A_iLAct$	Set the total number of activations for Appliance i
$TotalDiff = \sum_{i=1}^{n}(A_iTotalDiff)$	Set the total sum of Differences
$TotalCost = \sum_{i=1}^{n}(A_iTotalCost)$	Set the total Cost
$TotalCost \leq InputCost * Target$	Ensure the TotalCost is below the target price
$Minimise(TotalDiff)$	Objective is to minimise the sum of squared differences

Table 2. Appliance Use Optimisation constraints

To ensure each schedule mimics prior behaviour, we analyse historic usage patterns and infer a time series of the frequency with which an appliance was active at a given time. The aim is then to minimise the abnormality of appliance activation times; that is, to avoid recommending appliance use at times with very low historic use frequency. To achieve this, we divide each day of the week into timeslots based on the sub-task time granularity (in this work, 5 minutes). For each appliance activation in the historical data, we increment a count in each timeslot that matches the day and time (e.g. for a 15-minute activation we might increment Tuesday 4:00. Tuesday 4:05 and Tuesday 4:10). Having done this for several weeks worth of historical data, we now have a time series representing the frequency any appliance was in use at any given time during the 7-day week. We invert this time series (where F = the highest frequency in the series, all entries E are changed to F-E, Figure 5) which we use to produce a set of dummy constant Task Variables (each of height equal to the appropriate inverted frequency) for use with another cumulative constraint.

The peak value for the cumulative constraint is then a variable, with the aim being to minimise that peak, subject to the lower bound equal to the maximum height of the inverse frequencies. For each appliance activation, the height of all of its sub-tasks are scaled to be the gap between the highest and lowest

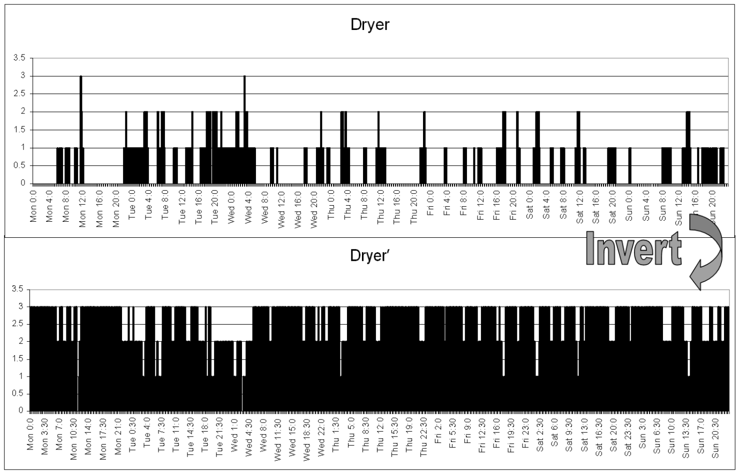

Fig. 3. Frequency Time Series and inversion to produce Masking Gradient

frequencies for that appliance, and the sub-tasks are added to the cumulative constraint (Figure 5). Thus, if we schedule a task so that the full duration of a task (i.e. each of its sub-task components) is at a time of highest frequency, the cumulative height variable is set to the lower bound; if the appliance task is scheduled so that it is active at the time of lowest historic frequency, then the cumulative height is at its maximum. All appliances are then normalised, and the optimisation objective is to minimise the sum of the appliance cumulative heights; that is, we minimise the sum of the 'abnormalities' of the activation times over each appliance use.

The initial solution to the scheduling can now be generated and displayed to the user. Although it may seem counter-intuitive to schedule uses of appliances like televisions or games consoles, it is important to account for their expected use, in order to manage the use of other appliances around that time. Thus we learn when those devices are most often used, and build that expected usage into the schedule. This initial solution does not require any interaction with the user to elicit preferences or objectives. We then allow the user to interact with the scheduler to improve the schedule, if necessary. The user can add new tasks or delete existing tasks for a particular day, can extend the duration of a task, can fix the start time of a task (by sliding the task position on the displayed schedule), and can add limited temporal constraints between tasks (for example, a specific dryer task must start within 30 minutes of the completion of specific washing task, or that the shower and washing machine cannot be active at the same time), and then can request a rescheduling. For new tasks, the same historical frequency pattern is applied; for user task time changes, the start time is fixed, and the frequency mask is adjusted to that no penalty is applied.

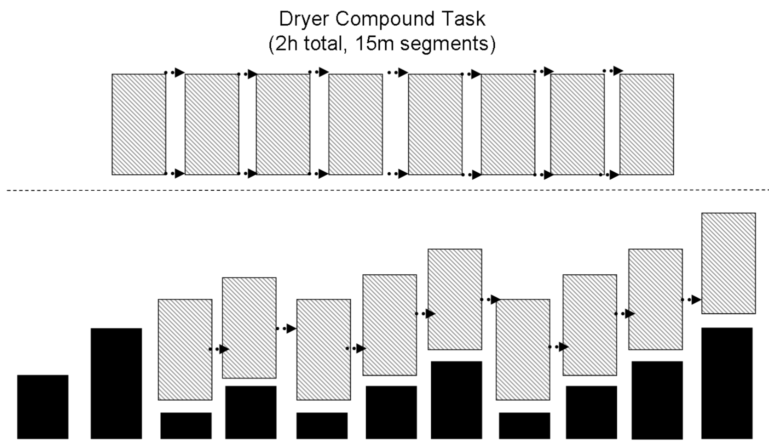

Fig. 4. Scheduling appliances using Masking Gradient

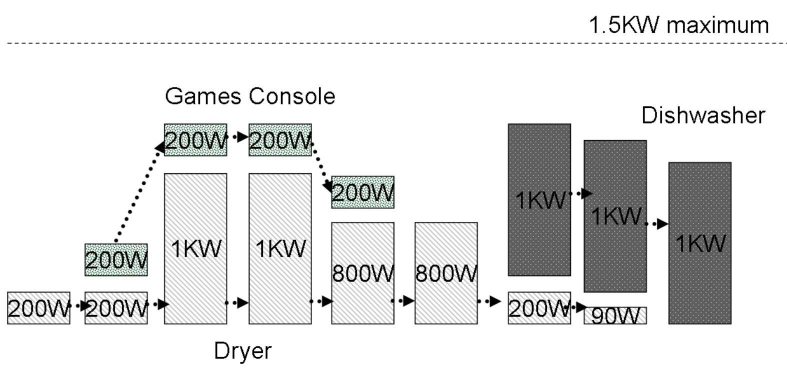

Fig. 5. Scheduling three appliances accounting for concurrent power demand

6 Performance Evaluation

A series of experiments were performed to determine the performance of the optimisation module on an Android tablet (Samsung Galaxy Tab 4 , Android version 4.4.2). Three months of data from a real-world home deployment formed the basis for the historical appliance use in these experiments, and the computation time taken for the Savings module was typically approximately 900ms and the Scheduler completed its task in approximately 17 seconds (using 30 minute time slots). These timings are well within the limits for acceptable response time, and demonstrate the success of the constraint programming formulation of the problems. As there is the facility for users to influence the savings problem through the addition of constraints on the number of appliance activations at different tariffs (i.e. the user can specify more limited ranges for the number of activations of appliances during particular tariffs, or to specify limits on the amount of adjustment made), we investigated the impact of the introduction of preferences on the running time and solvability of the savings problem. We also investigate the impact of larger and smaller time-slots on scheduler performance, with and without some user constraints applied.

6.1 Solver Performance

To investigate performance on the extreme end of user interaction with the appliance savings module, we performed experiments where we gradually restrict the scope of adjustment available to the solver. From savings targets of 10% through to 60%, we performed the savings routine while gradually restricting the range of acceptable values (or "freedom") for appliance activations at each tariff. For instance, at "40%" freedom, the solver is free to adjust the number of activations for appliances at any particular tariff band to within 40% of the historical values. As the freedom is reduced, solutions at high savings targets become unavailable, and eventually when the scope for adjustment is very low (low freedom value) no solutions can be found at any savings target, as shown in Figure 6. In these results, we observe that the imposition of user-specifiable restrictions on scope for adjustment to appliance use has little-to-no impact on the computation time (all solutions found took between 900 and 1150ms to discover), with solutions unavailable where the limits on value adjustment imposed by preferences prevent the discovery of solutions in the range of the target savings amount. Where no solution is available, the solver takes a consistent 780ms to determine this.

6.2 Scheduler Performance

To evaluate the performance of the scheduling module, we took the output of the activations saving module (without preferences) as the basis for scheduling the next week of appliance activations. In this evaluation, we schedule using a range of time-slot resolution values, from 5 minutes up to 60 minutes, and observe the computation time. As the users can introduce custom constraints to the scheduler, we re-perform the scheduling preocedure at each time-slot resolution level with the addition of an Inter-Appliance Disjunction constraint between two appliances (two particular appliances cannot operate simultaneously), and the further addition of a temporal link between two appliances (wherever one appliance ends, a particular other appliance should start shortly afterwards). The results of these experiments are shown in Figure 7. We observe that the performance with the addition of user-specified constraints remains approximately the same as without ("Default"), and that while computation time for small time-slot resolution is extreme (up to 400 seconds), with timeslots of 20 minutes or larger the computation time is much more reasonable (from around 30 seconds computation time at 20-minute slots down to 5 seconds at 60 minute slots).

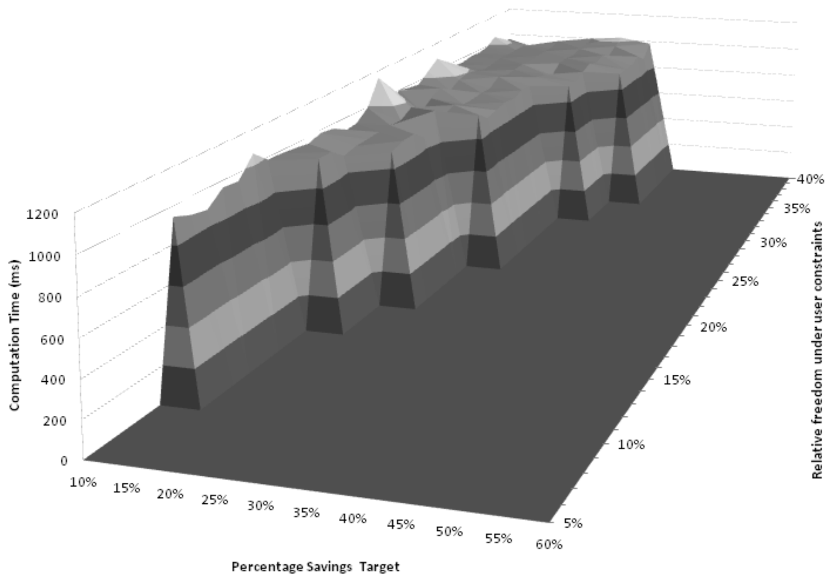

Fig. 6. Savings module performance under restricted scope

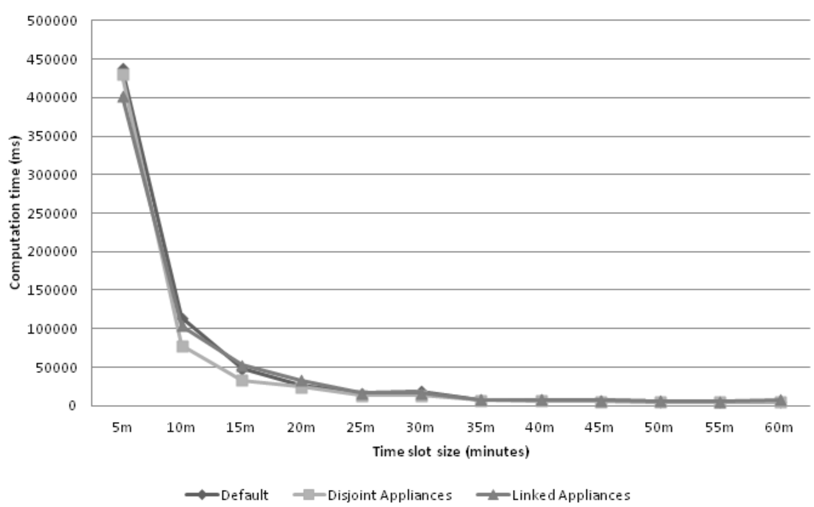

Fig. 7. Scheduler Results

7 Related Work

A number of examples of work in optimisation of energy use to balance demand and reduce costs exist intended to operate at the supply side. These approaches are typically intended to integrate with "smart" appliances in the home to allow for load balancing across the so-called "smart grid". As these approaches typically operate centrally at the utility, Mixed Integer Linear Programming (MILP) is commonly used to compute power use schedules based on energy availability [23,4,22,5,13,14]. MILP approaches take some time to compute (several minutes generally), and as they are solving over several homes and with respect to the utility company's energy supply requirements, there is limited scope for the home-owner to understand their personal energy use.

While not applicable to operating on tablets, Lim et al [17] describes a mixed-integer programming approach to air conditioning scheduling based on scheduled meetings in buildings, demonstrating the integration of energy-based requirements with real-world user requirements in scheduling. Scott et al [21] present a MILP-based framework featuring the concept of "comfort" in addition to the consumption and demand considerations that we also consider.

Felfernig & Burke[12] investigated the position of constraint-based models as the basis of Recommender Systems, noting that limited attention or understanding of complex recommender systems on the part of the user is a major consideration in the design of recommender systems. In our work, we use a simple-to-understand tablet-based interface to allow users to enter a dialogue with the system, introducing requirements which are interpreted as constraints, but without requiring particular expertise on the part of the user. Like our work, Ha et al [15] proposed a constraint-based approach for demand-side appliance scheduling, though this work was demonstrated using randomly generate synthetic data and compution time was in the order of minutes.

He et al [16] consider the case where electricity pricing is highly variable in time and the active use of energy by homes impacts the pricing, a potential future direction for energy production. They developed a market model combined with optimisation of community electricity use (using both constraint programming and mixed integer programming) for demand management and cost saving. Tushar et al [26] described a game-theory-based approach to demand-side energy management, modeling the interaction between home-owners and utilities through the trade and consumption of variable energy supplies and costs.

Darby et al [11] and Carroll [9] note that end-user perception has a significant role to play in motivating changes in energy use behaviours, with greater engagement and savings achieved simply through making more information available to the homeowner as found in smart-metering, from variable tariff trials in Ireland [10]. Bu et al [7] presented an MILP-based adaptive home energy scheduling approach with an initial discretisation step similar to the role of data pre-processing model in our work, though using synthetic test data rather than real-world data streams, and with relatively high computation time and resources compared to our approach. De Sá Ferreira et al [20] investigate highly variable tariff design at the utility level, generating optimised tariff pricing on a day-to-day basis. Highly variable tariffs generated through such a scheme could function well combined with the energy saving and scheduling described in our work, allowing for dynamic schedule generation on a daily basis while respecting the home-owner's typical behaviour patterns.

8 Conclusions and Future Work

We presented a constraint-based energy saving recommender system which uses real-world appliance energy use data to generate recommended behaviour changes to achieve energy saving goals in the home. The system is computationally efficient and operates on a low powered tablet. Integrating user preferences allows for a continuous, interactive optimisation process with users introducing additional requirements as constraints in the problem, and through periodic recalculation of solutions when the sensor-observed energy use behaviour differs from the suggested adjustments. This aspect allows for users to interact with the model and solutions without any special knowledge, and could lead to greater understanding on the part of the home-owner as to their energy use habits and their impact on costs and emissions.

In future we will expand the deployments and models, incorporating add room and water heating sensing, and the capacity to remotely actuate appliances attached to smart-plugs. Remote actuation would allow for the optimised schedules to be automatically implemented in the home without requiring the users to physically attend the devices, and would substantially ease the burden on the occupant to conform to the scheduled plans. As the AUTHENTIC project expands, further home deployments and long-term feedback from users will motivate expansing the range of user-specified constraints and further investigation into the performance of the scheduler in these scenarios. We will expand on the peak-load model for cases where temporally-restricted renewable energy or home-energy-storage is a feature. In the case that a home has a large energy storage solution (e.g. an electric car battery), a limited capacity of cheap energy could be stored overnight and used to power appliances the next day during higher tarriff times. Similarly, solar panels could augment the energy supply depending on the weather, and this could motivate opportunistic appliance scheduling in reaction to volatile cheap energy availability.

Acknowledgements This paper is funded by Enterprise Ireland through grant TC20121002D. We are grateful to the IERC and our academic and industry partners in Authentic for the design and development of the Authentic HAN, and to the anonymous occupants of our pilot homes for trialling the Authentic HAN, providing their data, and testing the Reporting Tool and Recommender System.

References

1. Application of delivery prices, orion group energy. http://www.oriongroup.co.nz/downloads/ApplicationOfDeliveryPrices.pdf, issued February 2 2015, accessed April 21, 2015
2. Authentic restful api. http://authentic.ucd.ie/documentation, accessed: 2014-12-01
3. Aggoun, A., Beldiceanu, N.: Extending chip in order to solve complex scheduling and placement problems. Mathematical and Computer Modelling 17(7), 57–73 (1993)
4. Agnetis, A., de Pascale, G., Detti, P., Vicino, A.: Load scheduling for household energy consumption optimization. Smart Grid, IEEE Transactions on 4(4), 2364–2373 (2013)
5. Bajada, J., Fox, M., Long, D.: Load modelling and simulation of household electricity consumption for the evaluation of demand-side management strategies. In: Innovative Smart Grid Technologies Europe (ISGT EUROPE), 2013 4th IEEE/PES. pp. 1–5. IEEE (2013)
6. Bray, T.: The javascript object notation (json) data interchange format (2014)
7. Bu, H., Nygard, K.: Adaptive scheduling of smart home appliances using fuzzy goal programming. In: ADAPTIVE 2014, The Sixth International Conference on Adaptive and Self-Adaptive Systems and Applications. pp. 129–135 (2014)
8. Carlier, J.: The one-machine sequencing problem. European Journal of Operational Research 11(1), 42–47 (1982)
9. Carroll, J., Lyons, S., Denny, E.: Reducing household electricity demand through smart metering: The role of improved information about energy saving. Energy Economics 45, 234–243 (2014)
10. CER11080a: Electricity smart metering customer behaviour trials (cbt findings report) (2011)
11. Darby, S.J., McKenna, E.: Social implications of residential demand response in cool temperate climates. Energy Policy 49, 759–769 (2012)
12. Felfernig, A., Burke, R.: Constraint-based recommender systems: technologies and research issues. In: Proceedings of the 10th international conference on Electronic commerce. p. 3. ACM (2008)
13. Georgievski, I.: Planning for coordination of devices in energy-smart environments. In: Proceedings of the Twenty-Third International Conference on Automated Planning and Scheduling (ICAPS 2013) (2013)
14. Ha, D.L., Joumaa, H., Ploix, S., Jacomino, M.: An optimal approach for electrical management problem in dwellings. Energy and Buildings 45, 1–14 (2012)
15. Ha, L.D., Ploix, S., Zamai, E., Jacomino, M.: Tabu search for the optimization of household energy consumption. In: Information Reuse and Integration, 2006 IEEE International Conference on. pp. 86–92. IEEE (2006)
16. He, S., Liebman, A., Rendl, A., Wallace, M., Wilson, C.: Modelling rtp-based residential load scheduling for demand response in smart grids. In: International Workshop on Constraint Modelling and Reformulation (Carlos Ansotegui 08 September 2014 to 08 September 2014). pp. 36–51 (2014)
17. Lim, B.P., van den Briel, M., Thiébaux, S., Bent, R., Backhaus, S.: Large neighborhood search for energy aware meeting scheduling in smart buildings. In: Integration of AI and OR Techniques in Constraint Programming, pp. 240–254. Springer (2015)
18. O'Sullivan, T., Muldoon, C., Xu, L., O'Grady, M., O'Hare, G.M.: Deployment of an autonomic home energy management system. In: Consumer Electronics (ISCE 2014), The 18th IEEE International Symposium on. pp. 1–2. IEEE (2014)

19. Richardson, L., Ruby, S.: RESTful web services. " O'Reilly Media, Inc." (2008)
20. de Sá Ferreira, R., Barroso, L.A., Rochinha Lino, P., Carvalho, M.M., Valenzuela, P.: Time-of-use tariff design under uncertainty in price-elasticities of electricity demand: A stochastic optimization approach. Smart Grid, IEEE Transactions on 4(4), 2285–2295 (2013)
21. Scott, P., Thiébaux, S., Van Den Briel, M., Van Hentenryck, P.: Residential demand response under uncertainty. In: Principles and Practice of Constraint Programming. pp. 645–660. Springer (2013)
22. Sou, K.C., Kordel, M., Wu, J., Sandberg, H., Johansson, K.H.: Energy and co 2 efficient scheduling of smart home appliances. In: Control Conference (ECC), 2013 European. pp. 4051–4058. IEEE (2013)
23. Sou, K.C., Weimer, J., Sandberg, H., Johansson, K.H.: Scheduling smart home appliances using mixed integer linear programming. In: Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on. pp. 5144–5149. IEEE (2011)
24. Sweeney, J.C., Kresling, J., Webb, D., Soutar, G.N., Mazzarol, T.: Energy saving behaviours: Development of a practice-based model. Energy Policy 61, 371–381 (2013)
25. Torriti, J.: Price-based demand side management: Assessing the impacts of time-of-use tariffs on residential electricity demand and peak shifting in northern italy. Energy 44(1), 576–583 (2012)
26. Tushar, W., Zhang, J.A., Smith, D.B., Thiebaux, S., Poor, H.V.: Prioritizing consumers in smart grid: Energy management using game theory. In: Communications (ICC), 2013 IEEE International Conference on. pp. 4239–4243. IEEE (2013)