Energy & Buildings 198 (2019) 275–290

Contents lists available at ScienceDirect

Energy & Buildings

journal homepage: www.elsevier.com/locate/enbuild

A behavior-centered framework for real-time control and load-shedding using aggregated residential energy resources in distribution microgrids

Andrea Mammoli[a,*], Matthew Robinson[a], Victor Ayon[a], Manel Martínez-Ramón[b], Chien-fei Chen[c], Joana M. Abreu[d]

[a] Department of Mechanical Engineering, University of New Mexico, Albuquerque, NM 87131 United States
[b] Department of Electrical and Computer Engineering, University of New Mexico, Albuquerque, NM 87131 United States
[c] NSF Center for Ultra-wide-area Resilient Electrical Energy Transmission Networks, University of Tennessee, Knoxville, TN 37996 United States
[d] Fraunhofer USA, Inc. Center for Sustainable Energy Systems, Boston, MA 02210 United States

A R T I C L E I N F O

Article history:
Received 19 September 2018
Revised 5 April 2019
Accepted 7 June 2019
Available online 7 June 2019

Keywords:
Residential load synthesis
Aggregated control
Demand response
Load shedding

A B S T R A C T

A bottom-up method to generate synthetic residential loads realistically, but with minimal computational resources, is presented. Six energy services, associated with high electricity use, are considered. Each energy service is characterized by number of events on a given day, event start time, and event duration. Distributions for number of events, start time and duration are proposed for four demographic categories: singles, couples, families and retired people. The distributions are augmented by elasticity parameters that allow load control and shaping. The distributions are based on information from focus groups and online surveys. In principle, the method can produce data at arbitrary temporal and topological resolution, and is thus suitable for a range of applications from machine learning of energy consumption patterns to detailed transient power flow analysis. Data can be aggregated as needed, for example by meter, by distribution transformer, or by substation transformer. In the present framework, loads for individual appliances, associated with individual electric meters, are generated at 1 Hz resolution, to explore two important applications that are relevant to the development of control paradigms for distribution microgrids. In such microgrids, a distribution feeder may be islanded from the bulk grid. The applications considered are aggregated real-time power dispatch and load shedding, both of which are needed for effective management of distributed energy resources in a microgrid setting. It is shown that aggregated loads can be shaped to follow a desired signal, for example to balance intermittent solar generation. Significant load reduction achieved by residents' behavioral response is also demonstrated. Such load reductions could be invoked in the case of low-probability, high-consequence events, and could contribute to increased energy resilience at the community level.

1. Introduction

The residential sector was responsible for 37% of the total energy delivered by the U.S. electric grid in 2015 [1]. This share is set to grow with the advent of electric vehicles [2], and with the electrification of many end uses needed to reduce carbon emissions [3]. While the contribution of the residential sector to the electric load is the largest, more than either commercial or industrial loads, its potential for the purposes of providing grid services such as storage and demand response currently is largely unexploited. Unlike commercial and industrial loads, residential loads are usually small and very heterogeneous. Their sheer number, small size and diverse use have hitherto made it uneconomical to install any kind of control capability, beyond that of the human user and of simple control devices such as thermostats or programmable timers. However, pervasive sensors and inexpensive distributed intelligence (the so-called Internet of Things) are rapidly changing this situation. Smart thermostats for home heating and cooling are now commonplace. Heat-pump water heaters with smart controllers are currently entering the market. Smart white goods including refrigerators and laundry appliances are sold by most major manufacturers. While this trend is setting the stage for using residential loads as controllable resources, a general framework

* Corresponding author.
E-mail addresses: mammoli@unm.edu (A. Mammoli), mr369@unm.edu (M. Robinson), vayon01@unm.edu (V. Ayon), manel@unm.edu (M. Martínez-Ramón), cchen26@utk.edu (C.-f. Chen), jabreu@cse.fraunhofer.org (J.M. Abreu).

https://doi.org/10.1016/j.enbuild.2019.06.021
0378-7788/

for this control is still in its infancy. Two necessary components of the framework are the ability to aggregate and control resources, and the ability to assess behavior-related responses.

In this work, the intent is to develop the modeling framework to simulate the operation of individual appliances on a hypothetical residential feeder of the near-future, with microgrid-like characteristics. Full electrification of energy services such as cooking, water heating, space heating and cooling is assumed. The framework is used to demonstrate both control algorithms and load shedding capabilities in such a setting.

Historically, the majority of residential load simulations follow a top-down approach, where measured feeder loads are distributed more or less equally among individual meters. An example is the work of Chan et al. [4], who use a physically-based methodology for synthesizing the hourly residential heating, ventilation, and air-conditioning (HVAC) load based on data from a utility. Significantly, the simulations are then used to evaluate the effect of load-management technologies. Calloway and Brice [5] present a first-order model of an HVAC load, and use it to evaluate the impact of load curtailment measures on the electric power system. The model can be calibrated as a function of the simulated buildings insulation value, infiltration, thermostat setting and air conditioner (AC) efficiency. The authors also suggest that this type of model can be used as a basis for bottom-up approaches. Modeling options for participation of thermostatically controlled loads (TCLs) in ancillary services were evaluated by Kamgarpour et al. [6]. In this work, the authors use methods from computer science and optimal control to produce various measures of control quality, including trajectory tracking and model error bounds. The balance of quality of control vs. cost of implementation is also addressed.

Bottom-up approaches present several advantages, including a better representation of the behavioral and socio-economic drivers behind electricity use, and a more realistic representation of power flow in a distribution feeder simulation. Many bottom-up models of residential energy consumption have been presented in the literature. Individual works focus on various aspects of bottom-up synthesis formulation, calibration, control or forecasting, often including models from the behavioral sciences. One group of bottom-up models broadly links technology and behavior to loads, particularly in aggregation. An early example of such models is the work by Capasso [7], who uses several probability functions to cover the close relationship existing between the demand of residential customers and the psychological and behavioral factors typical of the household. It includes psychological factors which affect individual use of the various electrical appliances. The work of Walker and Pokosky [8] follows similar principles: specifically, an availability function which statistically estimates the number of people in a household available to use an appliance, and a proclivity function which gives the probability that an individual will use that appliance at any given time of day. These functions are then used to drive models of the various appliances. In the work by Muratori et al. [9], the physical behavior of a dwelling is used to calculate energy consumption for heating and cooling, while energy-related activity patterns of individuals within the dwelling are used to construct energy consumption from services such as entertainment and appliances. The model predictions are verified against metered data. Data for energy consumption can be obtained at high temporal resolution using this model. An interesting perspective on constructing energy use patterns for individual households is provided by Widèn et al. [10], who focus on energy use resulting from activities of each individual in a household. To characterize activity patterns, hourly diaries were obtained from 464 individuals in 179 households. Output consists of hourly energy consumption for individual households. Many of the best qualities of residential load simulators presented in the literature, including high time resolution, socio-economic factors, seasonal variations and correlation between start-time and duration of an activity were incorporated recently into a single tool by Fischer et al. [11]. The tool, SynPro, was calibrated extensively against available data, and provides an excellent representation of real loads. Hourly load profiles are considered in a study by Ge et al. [12], who use multiple Gaussian distributions to reduce and simplify the data requirements for modelling. The advantage of the method presented is the requirement for only a very limited number of parameters to generate a household's hourly electricity load profile. The model essentially fits Gaussian curves to the aggregated loads, so it is not able to provide information on individual appliances, and while it can generate useful load data at a moderate frequency, it is not well-suited for implementing or planning demand-response (DR) measures. A bottom-up methodology to construct models of individual residential loads based on the use of inexpensive event loggers, is proposed by du Preez and Vermeulen [13]. Gotseff and Lundstrom [14] use the GridLAB-D agent-based distribution simulation environment to implement models of household electricity use that can provide high time-resolution data for individual appliances. The house model calibration was performed using transformer data for cases with similar weather. The modeled loads matched measured power consumption with respect to metrics such as daily average energy, average and standard deviation of power, power spectral density and load shape. Segmentation of households is identified as an essential component of bottom-up models to simulate electricity load profiles by Hayn et al. [15]. Categories based on household size, income and employment status are listed as most significant. A model providing hourly loads is used to evaluate the effect of technology, combined with tariff design, on aggregated electricity demand profiles. Bottom-up models intended to reconstruct loads are summarized in Table 1.

A second group of models is more focused on using bottom-up load simulations to investigate methods to control aggregated loads. Richardson [16] presents a high-resolution model of domestic electricity use that is based upon a combination of patterns of active occupancy and daily activity profiles that characterise how people spend their time performing certain activities. One-minute resolution synthetic electricity demand data are then created through the simulation of appliance use. The data are validated using actual electricity use recorded over the period of a year within 22 dwellings. An approach similar to the present work is presented by Cole et al. [17], who use a simulation of hundreds of houses to show the effect of four control schemes aimed at managing distribution feeder loads. Energy management is based on control of HVAC systems achieved by altering the temperature setpoint. Household energy use data from the Pecan Street Project were used to extract energy use patterns. The work is effective at showing how control scenarios influence energy use outcomes (including peak shaving and overall energy consumption), however real-time (i.e. on the order of 1 Hz) is not considered. An interesting outcome of these simulations is that peak reduction is accompanied by overall increase in energy use, due principally to higher energy consumption that results from pre-cooling, as well as, to some extent, to post-event rebound. Small-scale consumer models are proposed by Chrysopoulos et al. [18] for the purpose of bottom-up aggregation of appliance use, simulation of energy efficiency scenarios, and assessment of change in consumer habits resulting from tariff design. As in the work proposed here, appliance use is modeled using probability density functions (PDFs) for number of events, event start time, and event duration. The models can be calibrated with measured data if available. Demand-shaping mechanisms are also included in the models. Various time-of-use pricing schemes were tested, demonstrating that DR can be achieved effectively. As was also the case in the work of Cole et al. [17], the simulations shown that often peak reduction produced by DR schemes where energy cost reductions

Table 1
Survey of load modeling literature.

Reference	Load model	Behavior model	Calibration	Validation	Application
Capasso [7]	Appliance load profile	Occupancy, resource availability, income	Mail survey	Aggregated load comparison	Effect of technology
Walker and Pokosky [8]	Physical appliance model	Occupancy, proclivity to use appliance	Travel data, heuristic	Utility data	Load estimation for resource planning
Muratori et al. [9]	Physics-based models	Activity Markov chain	Time of use data	Comparison with individual and collective datasets	Policy decision support, effect of technology
Widen et al. [10]	Baseload and activity-based loads	Activity to end-use mapping	Time of use data with multiple data sets	Comparison to aggregate measurements	Alternative to load measurement for modeling
Fischer et al. [11]	Load profile, continuous	Demographics, frequencey-start-duration	Time of use data	Individual and collective load profiles	Effect of technology and extreme event
Ge et al. [12]	Aggregated load shapes	Activity times	Placement of Gaussian functions to fit load profile	N/A	Energy efficiency studies
du Preez and Vermeulen [13]	Appliance nameplate rating, ON-time	N/A	Appliance event logging data	Comparison with average aggregated household load	Residential load modeling
Gotseff and Lundstrom [14]	GridLAB-D house objects (load profile and physics), load data	Occupancy and cooling schedules, data-driven occupant behavior	Peak loading conditions on summer day	1-second resolution at secondary transformer	Modeling voltage regulation performance of batteries in high-PV feeders
Hayn et al. [15]	House with local generation and storage, with optimization	Fixed electricity and heat load profiles, electric vehicle	Heuristic	N/A	Effect of technology, socio-demographics

are offered off-peak often result in an overall increase in energy use, due to increased energy use in the longer off-peak period. The aspect of control design is addressed by Bahash and Fathi [19], who design a Lyapunov-stable sliding mode controller for the system governed by a set of bilinear equations, using a Monte Carlo model of TCLs. Demand Response-enabled load models were developed by Shao et al. [20], who considered individual appliances including space cooling and heating, water heating, clothes dryers and electric vehicles. The loads are physics based, and include the number of people in the dwelling as a model parameter. The models and methodology were validated against individual appliance loads, as well as in an aggregated sense by comparing with the RELOAD database [21]. A load model based on Markov chain Monte Carlo (MCMC) that combines accurate load electric characteristics and user behavior is presented by Collin et al. [22]. Significantly, the model is made freely available to the research community. A model of a system that can aggregate customer loads based on availability of controllable loads and customer preference is presented by Iria et al. [23]. The model uses a 5-minute time step, with a total of 30 households, and the appliances controlled are inverter-based AC units, hot water heaters and refrigerators. Jambagi et al. [24] present a bottom-up, high time and spatial resolution model based on time use survey data. Their model separates individual appliance use, and is therefore suitable for simulating the implementation of DR schemes. The model also accounts for demographic variables such as number of people in a household. The model is validated using minute-to-minute volatility, and coincidence factor, and shown good agreement between model and real data. Bottom-up models used for the purposes of developing control tools are summarized in Table 2.

In the framework presented here, the existing body of work is augmented by the collection of data that quantifies the willingness of individuals in households to alter their behavior so that certain common benefits can be achieved, including higher system resilience and improved ability to integrate intermittent renewable resources. Moreover, the data output provided by the synthesis tool can be used in conjunction with grid simulation tools to provide insight into infrastructure-related constraints that may affect results from resource optimization, as demonstrated in the recent work by Ayon et al. [25]. To the best of the authors' knowledge, no other work combines extremely efficient load synthesis, load forecasting and controllability in a single tool. Load control, in a real setting, would be based on a real-time stream of data, either from individual loads or aggregated in some way, and would work within the setting of microgrid or distribution system management. The load synthesis tool presented here effectively replaces the live data stream, and provides researchers with a means of building, calibrating and testing control and load-shedding schemes in a realistic setting, before field deployment. Sufficient detail is provided to allow the implementation of the tool, also providing access to calibration parameters. Such a framework could be especially valuable in the context of distribution-level microgrids.

2. Grassroots load simulation

The scenario considered in this work is one in which the electrification of energy services has already occurred, following decarbonization efforts. So, for example, water heaters consist of storage tanks with small heat pumps, cooking ranges are electric, and space heating and cooling is provided by reverse-cycle heat pumps. In addition, LED lighting has replaced both CFL and incandescent lighting. It is also assumed that major appliances are controllable remotely, rather than solely by the home occupants in manual mode. A bottom-up load simulation framework was chosen as the basis of the framework presented here for several reasons:

1. It is compatible with agent-based models such as GridLAB-D;
2. It is compatible with the need to model occupant behavior in response to technology and incentives;
3. It is compatible with the implementation of control strategies within the same framework;
4. It correctly captures power flow to and from customers, at the appropriate time scale and at the appropriate level of geographic detail;

Although there are tens of appliances that make up the total load measured by a meter, there are a relatively small set that dominate, either because of their collective energy consumption, or because they provide a useful means to control either energy or power demand. Here, the appliances and usage patterns are characteristic of conditions prevalent in the Unites States. Specifically, AC units were chosen as a means to provide fast demand response, due to their high power draw, while water heaters were chosen for the purposes of peak shifting, due to their high energy storage potential. In both cases, the energy service to the distribution is provided without affecting convenience and comfort. For the case of the AC units, the temperature remains

Table 2
Summary of bottom-up simulations for control algorithm development.

Reference	Load model	Behavior model	Calibration	Validation	Application
Richardson [16]	Appliance demand cycle, power factor	Occupancy activity profile	Ownership statistics, heuristics	Data from dwelling	Resource planning, demand-side management
Cole et al. [17]	Reduced-order physics-based model with weather, AC only	Occupancy price-sensitivity	Pecan Street dataset	Annual billing data	Load shaping, demand response for large synthetic community
Chrysopoulos et al. [18]	Pre-defined applince consumption model	Data-driven activity models	Data from instrumented houses	Comparison with measured data for individual houses	Effect of technology, customer habits, demand response
Bahash and Fathi [19]	Physics-based TCL model, AC only	Montecarlo	Heuristic	N/A	Control of aggregated TCLs
Shao et al. [20]	Physics-based models, controllable and critical loads	N/A	House structure-based	Comparison with real house measured data	Demand response strategies
Collin et al. [22]	Primarily appliance load database	User activity state	UK time use survey	Aggregated UK data	Distribution circuit analysis, effect of technology evolution, demand-side management
Iria et al. [23]	Physics-based models, EV state of charge	Customer activity patterns, presence, comfort settings	Known load parameters	Measured data for 30 houses	Demand response driven by energy aggregator optimization algorithm
Jambagi et al. [24]	Typical load profiles	Time of use, standard load profile, demographics	Based on smart meter data and heuristics	Comparison to standard load profiles	Demand-side management of large numbers
The present work	Markov-chain and physics	Activity frequency, start-time and duration, customer willingness to curtail appliance automatically or behaviorally	Focus group and online survey	Comparison with typical individual meter profiles	Distribution feeder power flow, critical events

within the deadband. For the case of the water heaters, sufficient hot water to satisfy demand is always available, while load shifting is achieved by better management of tank charging schedules. To test behavior-driven load shedding, all appliances are used, including refrigerators, clothes dryers, electric ranges and lighting. In this case, occupant comfort is affected by the load shedding measure, but here the assumption is that this is a voluntary demand reduction enacted by the energy user in a rare emergency situation, with the intent of enabling the community to survive the emergency while maintaining service to critical users.

2.1. Load models

Each load is modeled according to its physical nature, and to its interaction with the end user. Air conditioners and heat pumps for space thermal conditioning are modeled according to:

$$M_S \frac{dT}{dt} = \dot{Q}_L - \dot{Q}_R, \tag{1}$$

$$\dot{Q}_L = K_1[T_a - T(t)], \tag{2}$$

$$\dot{Q}_R = \Lambda \times \mathrm{COP} \times P_{AC}, \tag{3}$$

where M_S is the effective heat capacity of the conditioned space, $T(t)$ is the temperature of the space, T_a is the ambient temperature, t is time, Q_L and Q_R are the thermal losses from the structure and thermal gain from the air conditioner or heat pump, K_1 is a parameter that represents the building's thermal insulation (whose distribution is selected heuristically so that individual AC units cycle 2–3 times per hour), COP is the coefficient of performance, P_{AC} is the power of the air conditioner's compressor, and Λ is a state function that indicates whether the air conditioner or heat pump is ON or OFF. 100% penetration of AC units is assumed, with a mean electric demand of 3 kW for the compressor, and a standard deviation of 1 kW [26]. In addition, the penetration of 'smart' thermostats is assumed to be high, meaning that schedules are highly representative of life style and generally implemented correctly. It should also be noted that is that a first-order model of an AC unit with fixed COP and no internal or solar gains, and in principle this could significantly under-predict AC capacity during hot hours. More detailed models could easily include variable COP and internal gains. Solar gains could be more cumbersome to implement, but should be minimal in well-designed houses. Also, it is assumed here that the response of non-air-mass components in the thermal model is substantially slower than the timescale of the real-time control described here, and can be neglected to a first approximation.

In its basic form, the heat pump in cooling mode is controlled by the switching logic

$$\text{if } T(t) < T_L \text{ then } \Lambda = 0, \tag{4}$$

$$\text{if } T(t) > T_U \text{ then } \Lambda = 1, \tag{5}$$

where T_L and T_U are the lower and upper deadbands for the temperature control. When the temperature is within the comfort deadband, the state function Λ at a particular timestep is the same as its previous value, i.e. switching only occurs when the temperature goes outside the deadband. The human interaction with this device consists of the control of the temperature setpoint, while the deadband is automatically set to a narrower range when human presence is detected, and wider otherwise. If the house is occupied, temperature is controlled to remain within a comfortable deadband around a temperature setpoint, while if the house in not occupied, the deadband is reset to a much wider range. While in this work the occupied and unoccupied setpoints are fixed, a more realistic model could reflect the fact that occupants may adjust setpoints more dynamically. Moreover, the advent of better interfaces may facilitate this process, and perhaps even automate it.

Refrigerators share a similar model and control logic. The temperature $\Theta(t)$ of a refrigerator can be modeled by:

$$M_R \frac{d\Theta}{dt} = \dot{Q}_L + \dot{Q}_D - \dot{Q}_R, \tag{6}$$

$$\dot{Q}_L = K_2[T(t) - \Theta(t)], \tag{7}$$

$$\dot{Q}_D = K_3\,\Xi(t)[T(t) - \Theta(t)], \tag{8}$$

$$\dot{Q}_R = \Gamma \times \mathrm{COP}_R \times P_C, \tag{9}$$

where M_R is the effective heat capacity of the refrigerated volume (assumed constant here for simplicity, but this could change as a consequence of grocery shopping habits), Q_L is heat gain through the refrigerator walls, Q_D is heat gain through the open door, Q_R is the heat removed by the mechanical refrigeration, and P_C is the electric power of the refrigerator's compressor. The constants K_2 and K_3 model the refrigerator walls and the convection mechanism through the open door respectively. Γ indicates the binary state of the mechanical refrigeration, and its logic is the same as for the air conditioner / heat pump. $\Xi(t) = 1$ when the door is open, zero otherwise.

For the domestic hot water (DHW) load, the model represents a heat pump water heater with storage. This type of heater has the potential to be carbon-free (depending on the source of electricity) and also allows energy storage and load management opportunities. Assuming good stratification, the state of the water heater can be represented by its state of charge (SOC), which is related linearly to the location of the thermocline in the water tank. The system can be represented by:

$$\frac{d(\mathrm{SOC})}{dt} = -K_4 \times \mathrm{SOC} \\ -K_5\,\Pi(t) + \Delta\mathrm{COP}_H \times P_H, \tag{10}$$

where K_4 is a constant that quantifies the quality of the tank insulation, K_5 is related to the flow rate during water draw events, COP_H is the coefficient of performance of the heat pump, and P_H is its electrical power. $\Pi(t) = 1$ during water draw events, zero otherwise. The control is also similar to the case for thermostatic devices, however it is based on SOC rather than temperature, as described by:

$$\text{if } \mathrm{SOC} \geq \mathrm{SOC}_H \text{ then } \Delta = 0, \tag{11}$$

$$\text{if } \mathrm{SOC} < \mathrm{SOC}_L \text{ then } \Delta = 1, \tag{12}$$

As with the AC control, the state of the compressor Δ does not switch when the SOC is within the deadband. The mean heating capacity assumed for the water heaters was 4 kW (a typical value), with a COP of 2.2 at typical operating conditions [27], and a volume of 250 l. As was the case with K_1, distributions for the constants $K_2 - K_5$ were chosen to produce realistic behavior of the system, consistent with typical appliances.

The model for the clothes dryer is extremely simple: it is either on or off, and when it is on the power consumption is constant. Only the resistor (3 kW mean) is modeled, while the motor and fan are neglected for simplicity and because of their small relative magnitude. Event number, start-time and duration are modeled by statistical functions as described later, with the only constraint being that events cannot overlap.

For the lighting, the model also assumes either on or off state at constant power, for each light bulb in the house. The probability mass function of each light bulb is a Bernoulli distribution independent of the others. Power for the individual light bulbs (a mix of 8 W, 16 W and 24 W) is low because LED technology is assumed.

Heating elements in the electric range are modeled using a Markov-chain approach. During the 'active' part of the schedule (i.e. during cooking events), the transition probability matrix can be defined through probabilities P_{01}, which is the probability of transition from OFF to ON, and P_{10}, which is the probability from ON to OFF, plus their complementaries. This action simulates typical thermostatic activation of electric elements in a burner. Each electric element is associated with an individual power level, representing a small (1 kW), medium (2 kW) and large (3 kW) element. Up to three elements can be active at any one time.

2.2. Modelling household energy routines and their elasticity

The human interactions associated with the energy use models above (e.g. occupancy schedules for AC, door opening for refrigerators, etc.), are associated with events, characterized by number of events per day, event start time and event duration. In turn, these events are represented by probability distributions. To add realism to the simulation, and to afford the framework better portability to different locations and demographics, customer types were assigned to 'clusters,' each representing an appliance usage pattern that could be considered typical. Similarly to the approach taken by Fischer et al. [11], four demographic classes were used to cluster customers: single home occupant, working couples, families with children and retired people. To gather insights about household routines and the willingness to modify the use of the appliances that enable those routines, two semi-structured focus groups, each with eight participants, were conducted in Michigan (Plymouth), and New Mexico (Albuquerque) [28]. These cities were selected to ensure climatic variability. The focus group sessions were designed to build on each other and continuously dive deeper into the beliefs and concerns regarding shifting household activities, and their potential flexibility for the benefit of the individual or the community in various emergency scenarios. As shown in Fig. 1, turning off lights as needed, dishwashing, and heating and cooling would be easier activities to adjust while showering and cooking would be more difficult. Concerns raised regarding shifting appliance use during seasonal emergencies included technical issues that could prevent the utility from turning appliances back on, loss of autonomy and privacy. During the discussion new themes emerged: saving energy was regarded positively. Overall, participants were willing to take actions to change their household practices above and beyond what they would normally do for the welfare of their communities, but required advanced warnings. These insights were used to support the design of two statewide surveys (N=1500, from New York State and Texas), balanced for gender and income.

The survey contained 70 questions: socio-demographic, and aspects of energy use: heating and cooling, washing, showering, entertainment, willingness to participate in energy savings and energy curtailment events. A full analysis of the survey data is beyond the scope of this work, and is discussed in other venues (see e.g. [29]). However, some of the survey data were used to inform the model and ensure an adequate level of realism.

Four demographic clusters were identified, each with its specific energy patterns: families are households where with three or more occupants; retired people are identified by age greater than 64; working singles are non-retired people in households of one person; working couples are non-retired people in households of two people. While there may still be overlaps between categories and some mis-categorization (for example, there may be some households of three or more people some of whom are retired), the data show fairly distinct and self-consistent behaviors. The proportion of each demographic cluster for the entire sample is shown in Fig. 2. As expected, families are the largest individual demographic cluster, forming almost half of the sample population. That, compared with the fact that these households are generally the largest energy users, means that the family demographic cluster is expected to dominate the total electric load.

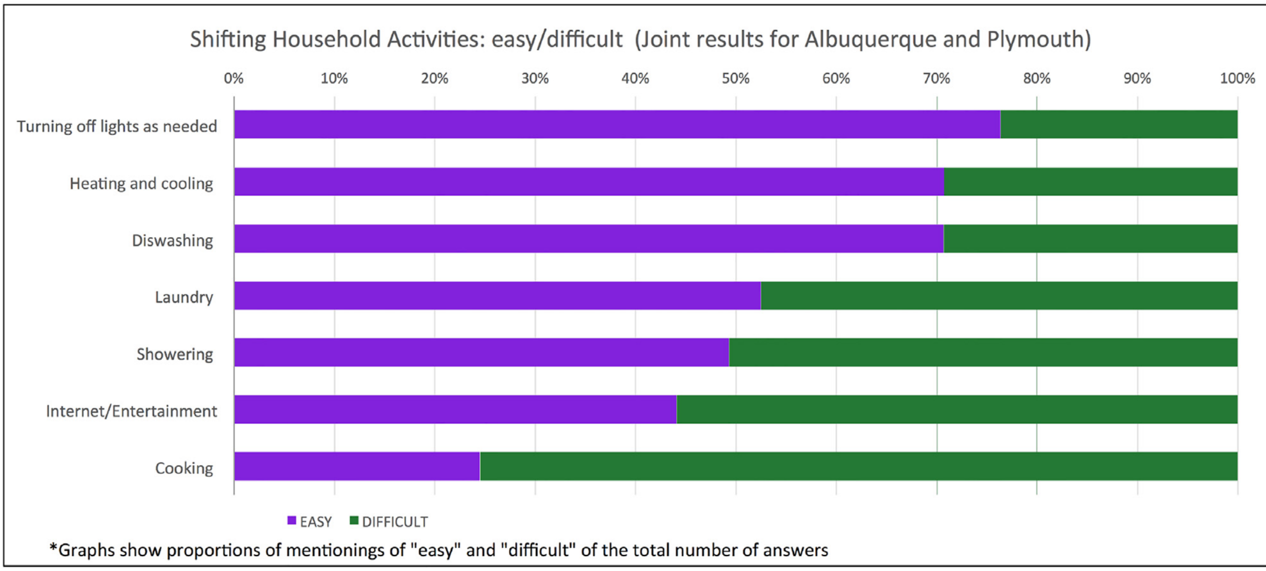

Fig. 1. Focus group results indicating ease / difficulty of shifting household activities in time.

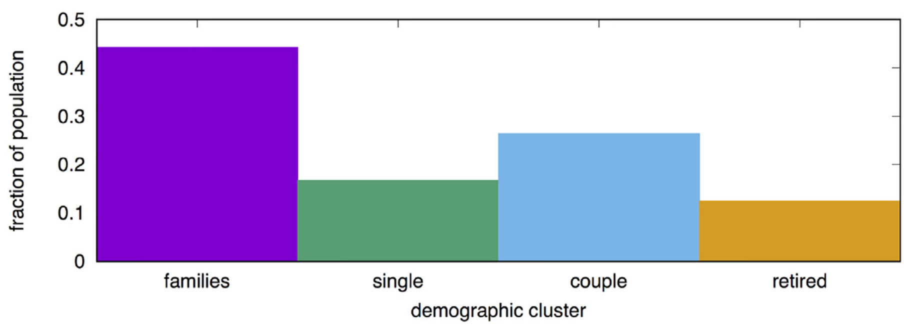

Fig. 2. Proportion of each of the families, retired people, working singles and working couples demographic clusters, among the 1500 respondents to the survey.

Air conditioning is one of the most energy intensive activity in the household, in warm regions. Energy consumption is highly dependent on the temperature setpoint. The ability to control or curtail air conditioning load is the extent to which temperature variations are tolerated by house occupants. The distribution of space temperature setpoints for the four demographic clusters is shown in Fig. 3, for the case where the house is occupied, and for the case that it is not. While all demographics show a similar distribution of setpoints, there are some clear distinctions. First consider the case of occupied houses. For families, the setpoint distribution function is skewed to the left (i.e. to lower temperatures), and is broader. Conversely, the distribution for retired people is skewed to the right and is more narrow. For singles, the distribution is similar to the one for retired people, but more skewed to the left. For working couples, the distribution is symmetric about the mean and centered in the range 70–72°F. A small percentage of the population, particularly singles, do not use air conditioning devices. For the unoccupied case, the trends are similar, but the underlying distributions are shifted to the right (i.e. warmer temperature). A much larger fraction of the population, on the order of 20%, do not use cooling at all when away.

Domestic hot water use is another major source of energy consumption, as well as a major opportunity for load shifting. Showers are usually the largest hot water draw. The temporal distribution of shower times obtained through the survey are shown in Fig. 4. All distributions are bimodal, with morning and evening peaks. For families, both morning and evening peaks occur earlier, and while the morning peak is higher, the magnitude of the two peaks is more similar than with other demographic clusters. For the retired demographic cluster, the broad morning peak occurs later in the morning, and the evening peak also occurs approximately two hours later than with other clusters. Both working singles and couples demographics are associated with a sharp morning peak between 6 and 7, while working couples are also associated with a larger than average evening peak between 20 and 21 h.

Occupancy time is an important parameter because it may be possible to shift the time that an appliance provides an energy service, or gets ready to do so by storing energy. The survey shows that houses with families are quite likely to be occupied at all peak times, with a likelihood of occupancy ranging from approximately 50% in the early afternoon, increasing to approximately 75% in the late afternoon and evening. For retired people, the probability of occupancy is always high, ranging from approximately 80% in the early afternoon to over 90% in the evening. For working singles and couples, the probability of occupancy rises from the mid 30% in the early afternoon to approximately 70% in the evening, with couples associated with slightly higher probability of occupancy (Fig. 5).

With an expanded goal-expectation theory framework, the survey provided data to examine how residents' willingness to reduce energy consumption and choices of giving up certain major household appliances relate to personal goals of maintaining power for family and neighborhood, expected level of cooperation from neighbors, and expected efficacy of energy-saving behaviors. Additionally, the survey provides data to examined residents' willingness to allow utilities to automatically control household

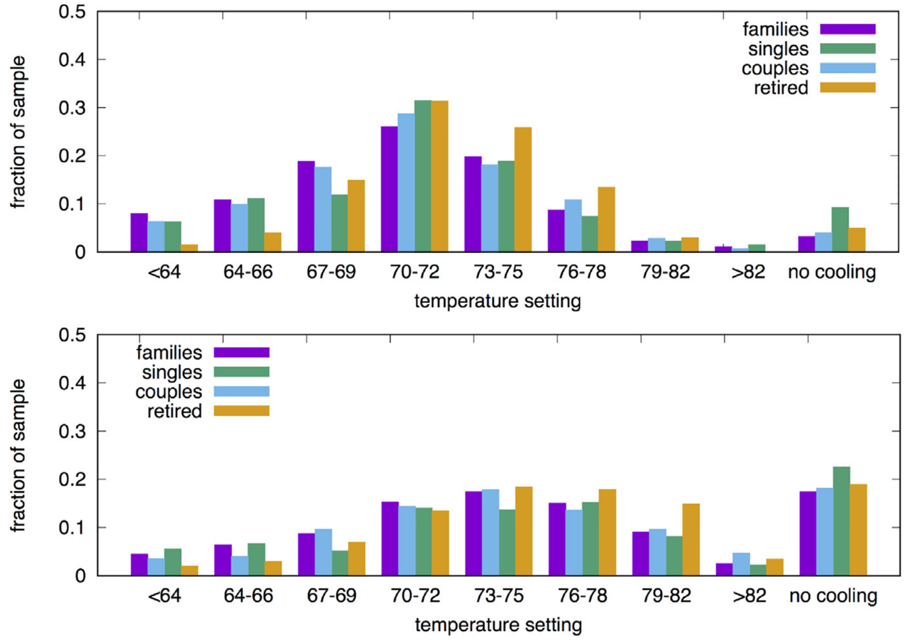

Fig. 3. Space temperature setpoints for occupied mode (top panel) and unoccupied mode (bottom panel), for the four demographic clusters, during the summer season. Seasonal differences were not captured in this survey.

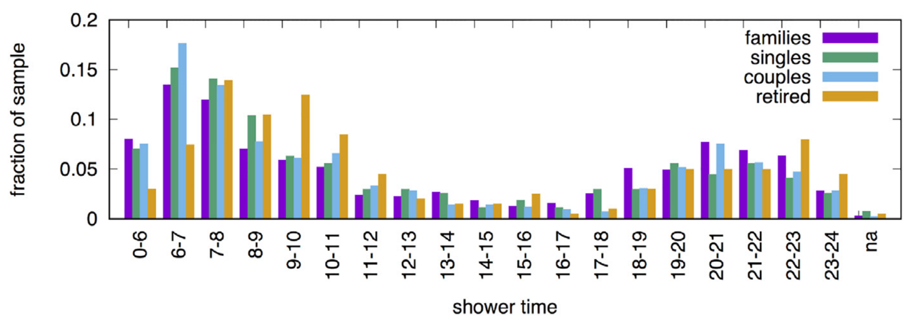

Fig. 4. Distribution of showering time for the four demographic clusters.

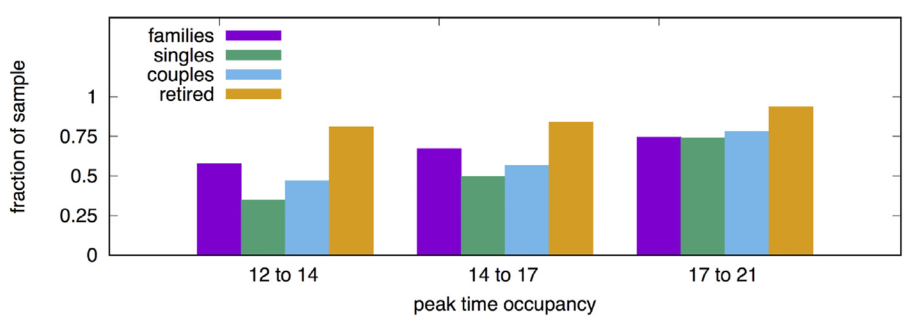

Fig. 5. Probability of occupancy during afternoon and evening peaks for the four demographic clusters.

appliances during normal and extreme conditions. While a detailed study of the survey results is beyond the scope of this paper, it is interesting to note that the willingness of the population to give up certain energy services to maintain basic service is generally high, as shown in Fig. 6. Most people are either likely or certain that they would give up some services. Approximately 20% of the population is as likely as not to give up energy services, while only a small percentage is either unlikely to forego energy services, or certain they will not (unless, of course, the power fails).

2.3. Transforming self-reported data in probability distribution functions

A probability density function for the start time and one for the duration were assigned to each combination of appliance and customer cluster, with shapes based on the survey response data. In addition, each cluster / appliance combination also was assigned with discrete probability values associated with the number of events per day. Start time distributions are a mixture of Gaussians, truncated at the edges, smoothed so that initial and final values

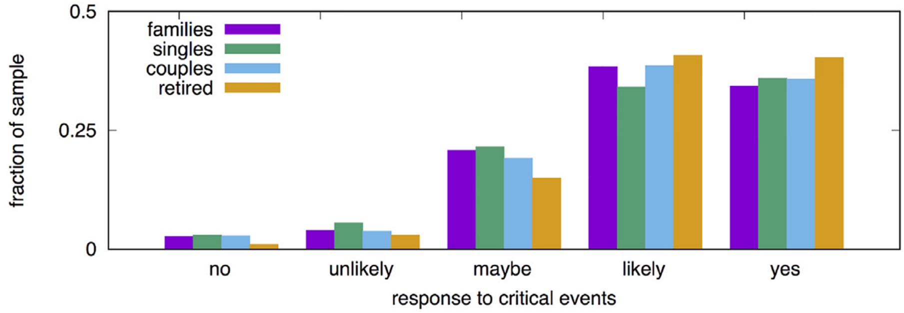

Fig. 6. Response to a question on whether the respondent would forego certain energy services to reduce the likelihood of total power failure.

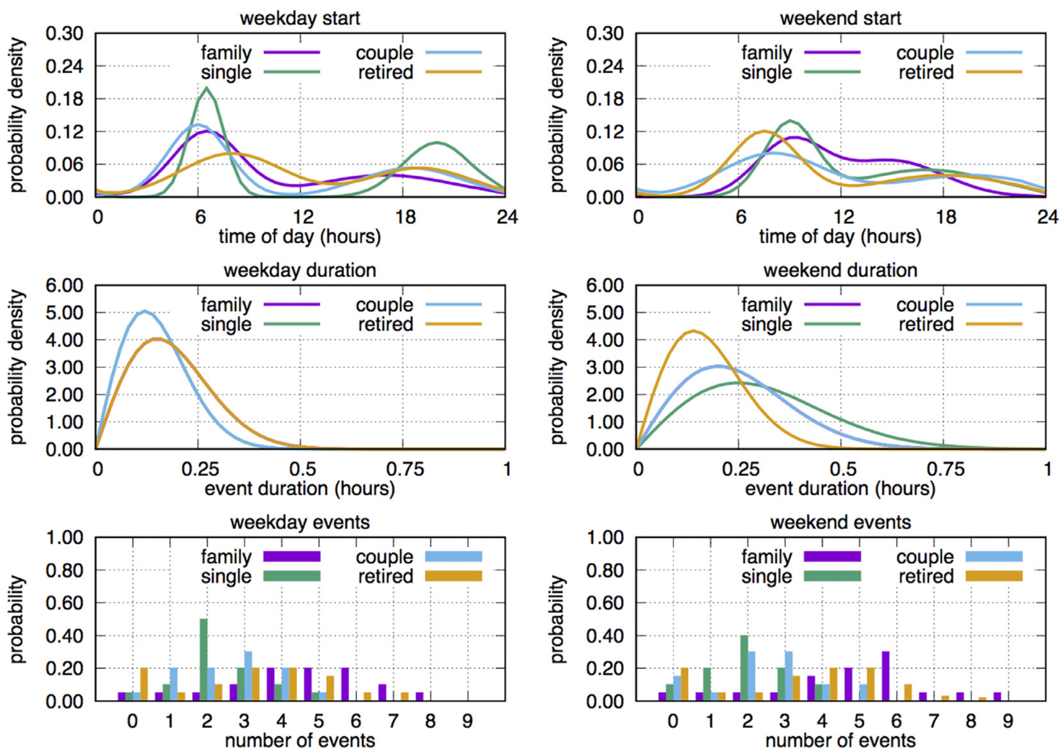

Fig. 7. Statistical characterization of the domestic hot water draw schedules, by start time, duration and number of events.

are continuous in value and slope, and normalized so their integral over the day is unity. The duration density functions are modeled using Weibull distributions with varying shape parameters and scale parameters. The probability mass functions for event number (where an event is defined as the change of state from OFF to ON) are defined manually, except for the case where many event numbers exist, in which case discretized Weibull distributions are used.

The rationale behind the design of the PDFs for each demographic cluster / appliance combination is exemplified by the case for domestic hot water (Fig. 7). For households occupied by single working people, hot water draws on weekdays are fairly regimented, with a strong morning peak corresponding to showering and an evening peak corresponding to cooking and dishwashing. The pattern is similar for working couples, with slightly broader and earlier morning and evening peaks. Morning and evening peaks also characterize draws for families, however peaks are broader and the probability density in the middle of the day is non-trivial. The pattern for retired people is similar to the pattern for families, but with later morning and evening peaks due to less constrained life schedules. During weekends, the patterns for all categories are modified by later morning peaks and non-zero density in the middle of the day. Weekday duration distributions for family and retired user clusters are equal to each other, as are duration distributions for singles and couples.

The former groups on average utilize longer water draws. The duration distribution for retired people on the weekend is the same as on weekdays, while durations are extended for the other groups due to more relaxed schedules. Finally, the number of events is largest for families, and lowest for singles. Retired people have a more uncertain pattern, with widely distributed number of events, due to their less constrained lifestyle.

The rationale for building statistical distributions is similar for other energy uses, in this case AC, refrigeration, clothes dryer, cooking range and lighting. These are not explained in detail here for brevity. Similarly, distributions for other smaller appliances are also not listed. Of course, the classes of behaviors models considered here do not represent an exhaustive list. For example, it would be possible to characterize behaviors related to holidays, vacations, sporting events and others, by simply adding the corresponding distributions, while retaining exactly the same structure. Similarly, it would be possible to extend the list of appliances modeled.

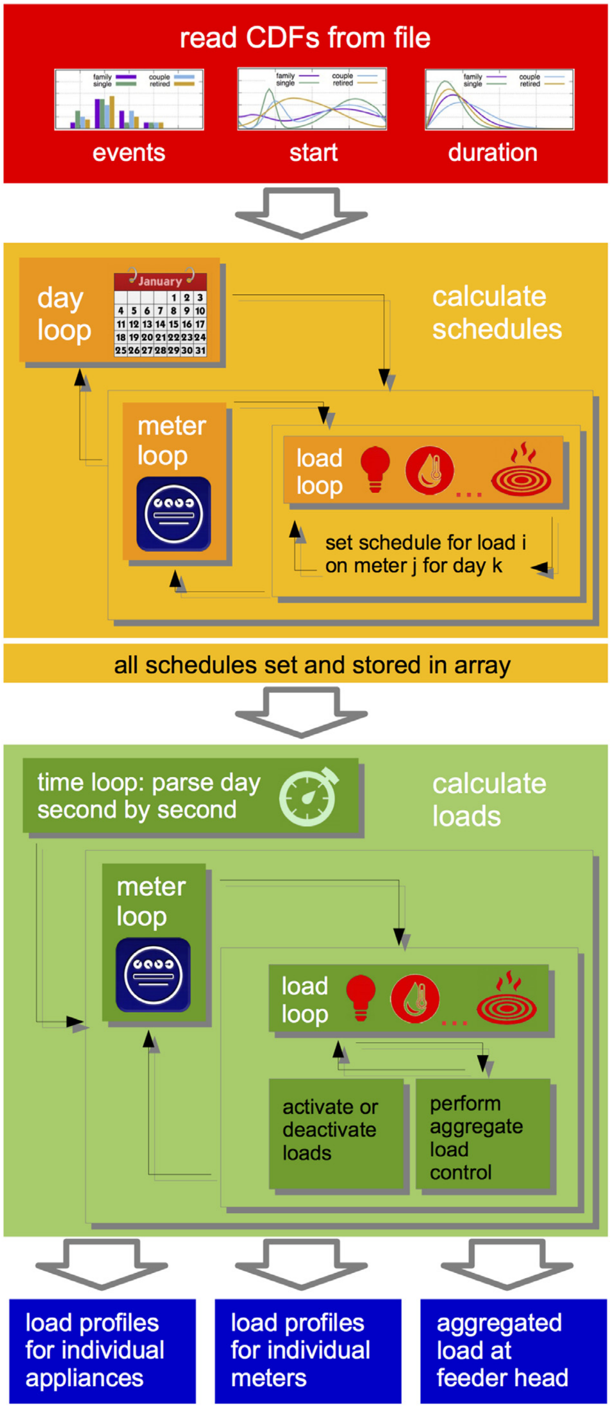

Fig. 8. Structure of the grassroots load simulator, showing the three principal phases of statistical data acquisition, schedule generation and load synthesis.

2.4. Structure of simulation engine

The grassroots load simulation is organized as shown in Fig. 8. In the first stage, the cumulative density functions (CDFs) associated with the start, duration and number of events probability density functions discussed in Section 2.3 are read from data files, and stored in array structures for quick access.

In the second stage, schedules for each appliance / meter combination are drawn from the CDFs, for the entire duration of the simulation period, specified in terms of number of days. The schedule generation follows a relatively simple set of rules. On a given day in the simulation period, for a given meter, a schedule is generated for each appliance, as appropriate to the appliance (see Section 2.1), following the steps:

1. Obtain number of events, by sampling the probability mass function for the event number for the relevant appliance and demographic cluster;
2. Obtain start time and duration by sampling the relevant CDF, ensuring that the event does not overlap with other events (except for the case of lights);
3. If an event extends into the next day, mark the end time of this event as the minimum start time for events on the following day.

All schedules generated are stored in memory, for later access during the load generation phase.

In the load generation phase, the algorithm simulates the real operation of the appliances, and produces the corresponding electrical load. Accordingly, time is the outer loop variable. At each timestep (typically one second) all loads on all meters are calculated, using the appropriate equations discussed in Section 2.1.

Appliances are turned on and off according to two criteria:

1. the physical / human interaction model, for example the combination of house / AC unit characteristics and outside air temperature (physical) and the schedules (human interaction);
2. a control signal from an energy aggregator service.[1]

For the present case, aggregator control signals are only applied to the house AC unit, and to the water heater, for real-time load control and resource scheduling respectively. However, the same logic could be applied to other schedulable or deferrable loads, for example to refrigerators.

Load files are produced for individual meters, and for the aggregated load from all meters, by appliance. So, for example the load history for the water heater behind meter *i* can be produced, as well as the aggregated load from all water heaters that the feeder substation meter observes. Disaggregated loads are seldom available for such large number of meters, but they could be useful to train machine learning algorithms used for appliance load disaggregation (see e.g. [30,31])

3. Characteristics of individual and aggregated loads

The load profiles over the course of a week, for each of the six appliances or devices modeled for a typical, belonging to the 'family' demographic cluster, is shown in Fig. 9.

It is evident that many of the loads (dryer, air conditioner, range) are very intermittent in nature, while others (lights, refrigerator) turn on and off very frequently, reflecting the event probability defined in the statistical distributions discussed previously. Some of these loads are large (e.g. dryer and air conditioner), and present opportunities for aggregated control. Others, such as LED lighting, are small and their external control would result in unacceptable loss of quality of service. However, they must be considered as part of the overall load when used within power flow calculations.

The aggregated load, as would be measured at the distribution feeder but showing the contribution of each appliance category, over the course of one week, is shown in Fig. 10. The total load for a typical meter is also shown for comparison. The difference in nature between loads at individual meters and aggregated feeder loads is evident. Aggregated loads are consistent and relatively

[1] An aggregator is any organization or individual that brings energy customers together as a group.

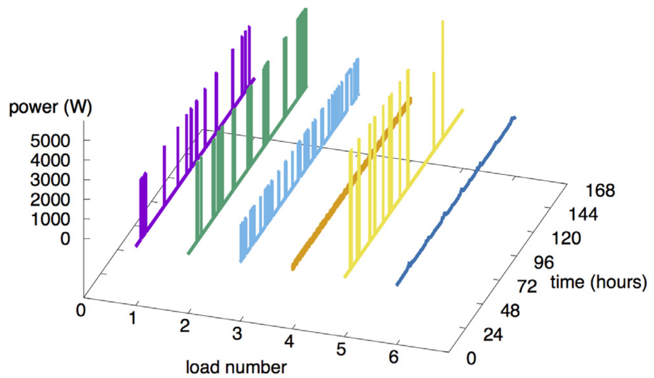

Fig. 9. Loads on a typical meter, with loads 1–6 associated with dryer, AC, domestic hot water, refrigerator, cooking range and lights respectively. Days 1 (0–24 h) and 7 (144–168 h) correspond to Sunday and Saturday respectively.

predictable from day to day (when categorized by weekday / weekend), while individual meter loads are not, as is the case in real life. This comparison highlights the necessity of the grassroots approach to capture the potentially highly localized power flow conditions along the feeder, especially for cases with a relatively small number of customers. It should also be noted that not all appliances are considered here. For example, electronics, miscellaneous loads, TVs, dishwashers, cloth washers are missing, but the overall result would be similar to what is shown here, because the underlying life patterns are similar. Moreover, these generally tend to be smaller loads than the ones considered here.

Another interesting feature is the ratio between the feeder peak (approximately 2 MW) and the individual meter peak (approximately 12 kW) of 167. Given that there are 1000 meters, it is clear that, at least over the course of one week, individual meter peaks do not coincide. However, this also highlights the fact that if all meter peaks could be made to coincide (e.g. by a cyber-attack on smart appliance controls or by poorly designed DR algorithms), severe problems could arise. The model could be used to simulate such scenarios.

The proportions of electricity load due to each class of appliance are consistent with the data reported in the recent Residential Energy Consumption Survey (RECS) [32]. Discrepancies may be due to climatic differences for our small sample, demographic difference, and our assumption of high electrification of energy uses.

Finally, inspection of the load duration curve reveals that the peak 12.5% of the load could be eliminated by shifting loads by a total of 3 h over the course of one week. 25% of the peak load could be removed by shifting loads for a total of 14 out of 168 h. Both measures are easily achievable without loss of quality of service, given adequate controls.

4. Controls for aggregated resources

Demand-response (DR) programs that use residential resources, primarily TCLs such as AC units or water heaters, are relatively commonplace [33,34]. These typically involve simple actions such as turning off devices for a specified period. In some cases, override capabilities allow the customer to stop the DR event at their premises, usually with some form of penalty, if the quality of service (QoS) reduction is unacceptable. Such traditional DR programs are based on pre-Internet-of-Things communications. New connectivity between devices and networked computing and optimization resources make it possible to improve on existing DR actions to make them useful to distribution system operators, utilities and to customers themselves. More recently, DR schemes that take advantage of IoT capabilities have been introduced, but still the prevailing operating principle is to reset the temperature [35], sometimes at the expense of comfort.

The approach used here to achieve real-time control of the aggregated load of many AC units and water heaters is similar to that proposed by Mathieu et al. [36] for a variety of TCLs. In this approach, comfort is not sacrificed, because the TCLs always operate within their user-defined deadband. Here we show that the approach performs as intended in a quasi-experimental setting, by combining it with the bottom-up load model. Air conditioners are attractive because they are high-powered, meaning that useful aggregated response is possible, and because they possess an inherent storage capacity associated with the effective heat capacity of the space they are conditioning, which however is relatively small. The combination of high power and low energy capacity makes AC units suitable for responding to high-frequency events (with characteristic times on the order of seconds to tens of seconds), like cloud-driven intermittency of photovoltaic (PV) generation. Water heaters, specifically of the heat pump type, are attractive because of the large storage capacity, combined with daily periods in which hot water use is limited. Water heaters are particularly well suited for shifting loads, for example to reduce peak load on a distribution transformer, or to reduce demand on a resource-constrained islanded feeder as is the case here.

Consider first a fleet of 1000 AC units, with a mean electric demand of 3 kW, and assume that all are available to participate in a real-time DR program (with response time on the order of seconds to a few minutes). The aggregated uncontrolled simulated load of all AC units for the period from 12:00 (noon) to 16:00 on day 1 (a Sunday) is shown in Fig. 11.

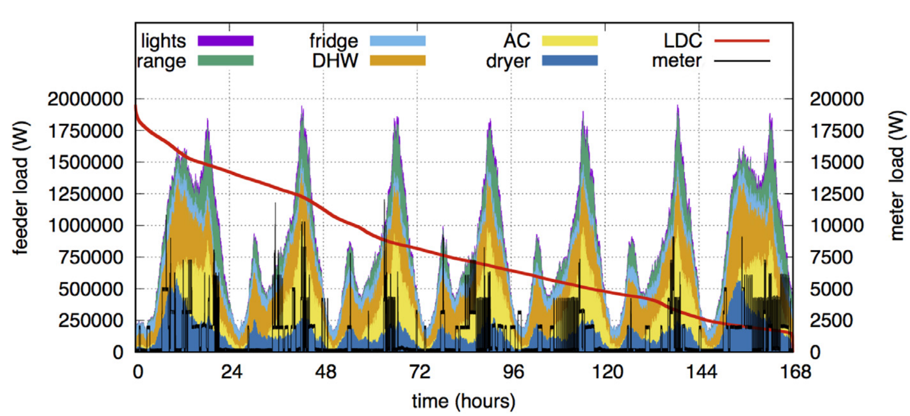

Fig. 10. Aggregated loads on feeder for one week, and associated load duration curve (LDC). The load for a typical meter is also shown for comparison with the aggregated load.

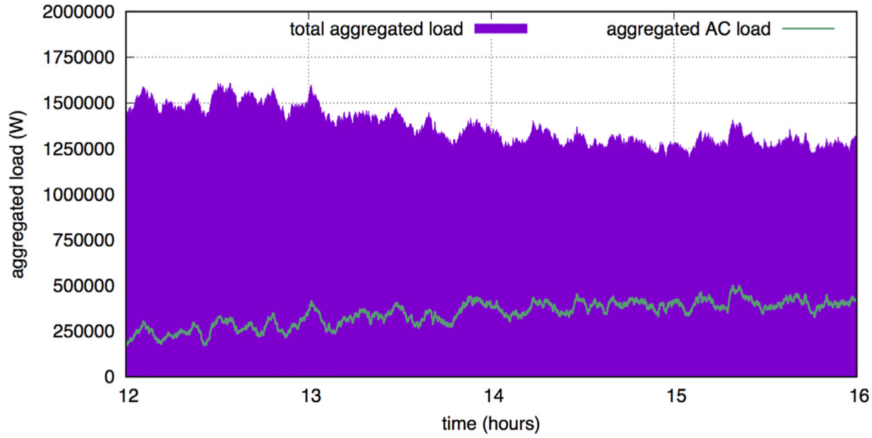

Fig. 11. Aggregated total simulated loads on feeder for noon to 16:00 on day 1 (a Sunday, with mild temperatures ranging from 17 °C to 30 °C), and aggregated simulated AC load.

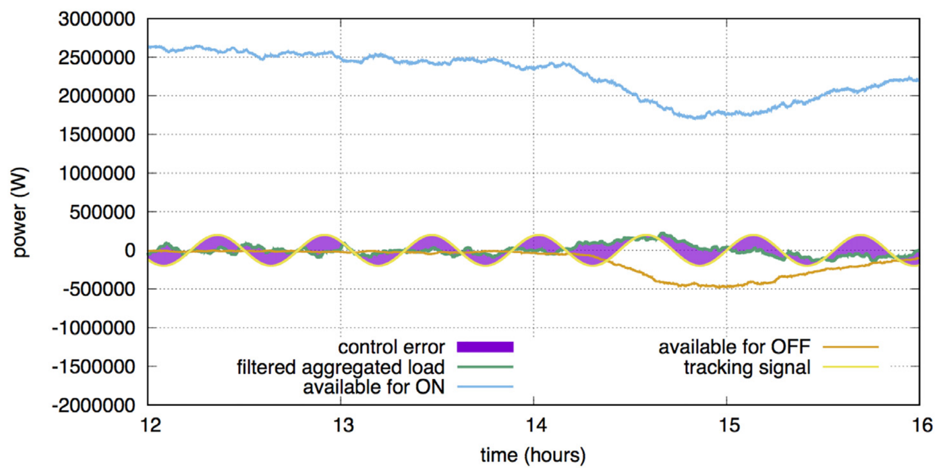

Fig. 12. Capacity of the aggregated AC loads to respond to a control signal, showing the total aggregated load, the load increase and reduction capacity of available AC units, and a sinusoidal control signal.

To achieve load control without affecting the QoS, it is necessary to control AC units associated with a space where the temperature is within the comfort deadband. AC units that are active (ON) can be turned off to reduce load, while units that are inactive (OFF) can be turned on to increase load. An additional constraint is that a compressor cannot be switched unless it has been in the current state for a minimum amount of time (the deadtime), and the space temperature is within the deadband margins.

The objective is to shape the total feeder load to respond to some external control signal, for example to the power production of a substation-sited PV array. It is unreasonable to expect that a given resource will be able to provide unlimited frequency response, since the resource may not have adequate energy, power or ramping capacity. Therefore, first the total feeder load is filtered using an infinite impulse response (IIR) recursive bandpass time-series digital filter [37]. The filtered load is shown in Fig. 12, along with a sinusoid representing the desired load shape. The available response capacity, in terms of the total power of units available to switch in either direction, is also shown.

The control is implemented following previous work by the authors and others [38–40], and reflects a simple proportional control strategy. The assumption is that each AC unit is associated with a smart thermostat capable of reporting its status to the outside world in real-time, of receiving external signals from a load aggregator, and of performing some relatively simple calculations. The control strategy is implemented according to the following steps:

1. At timestep k the substation reports its control error $\mathcal{E}$ - namely the difference between a desired load and the current load;
2. The load aggregator calculates the maximum positive and negative response capacities (i.e. the capacity to increase or reduce load respectively), $\mathcal{P}$ and $\mathcal{N}$ respectively, from information about the reported state of each smart thermostat (temperature of the space, deadband, ON/OFF status and time in present state);
3. Noting that $\mathcal{R}$ is the maximum ramping rate of the system (i.e. the rate of change of load), obtained by multiplying the aggregated power of all the AC units operating at a given time by the high frequency cutoff, the load aggregator calculates the fraction $\mathcal{F}$ of AC compressors that must be switched according to:
$$\mathcal{F} = \begin{cases} \frac{\min(\mathcal{R},\mathcal{E})}{\mathcal{P}}, & \text{if } \mathcal{E} > 0 \\ \frac{\max(-\mathcal{R},\mathcal{E})}{-\mathcal{N}}, & \text{otherwise} \end{cases} \tag{13}$$
and broadcasts this number to all thermostats;
4. Each thermostat calculates a random number $\mathcal{I}$ between 0 and 1 from a uniform distribution. If $\mathcal{I} < \mathcal{F}$ then the compressor is switched from ON to OFF at timestep $k+1$ if the error is positive, or from OFF to ON otherwise. This process guarantees that the expected fraction of ACs turned or on off is $\mathcal{F}$. The variance from $\mathcal{F}$ is proportional to the

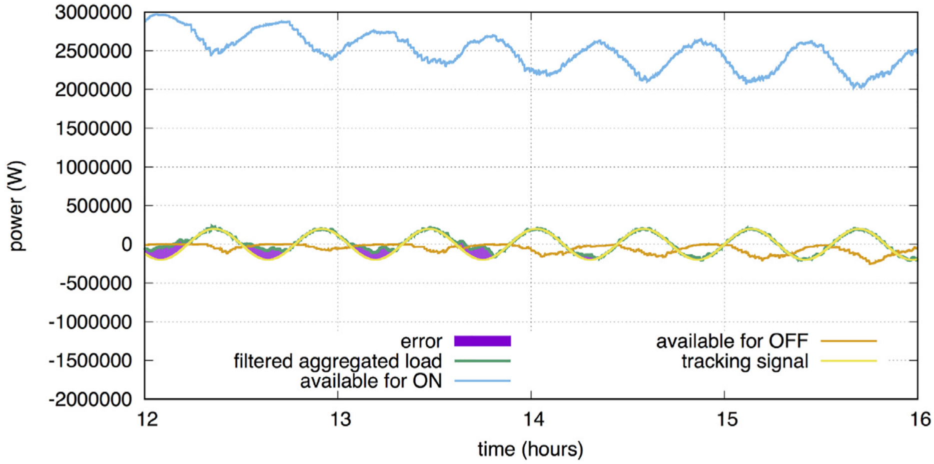

Fig. 13. System response to the probabilistic control strategy for aggregated loads.

inverse of the number of ACs. For example, for 1000 units, the highest possible standard deviation is 1.6×10^{-2}, so the typical contribution to the control error of this random process is negligible. Moreover, this can be calculated easily by having access to the number of thermostats participating.

The response of the system to this control strategy is shown in Fig. 13, for a signal with frequency of 0.0005 s^{-1}, corresponding to a period of 2000 s. Noting that the response of a system to a sinusoidal input of varying frequency is the basis of the construction of a transfer function, the sinusoid exemplified here is representative since its frequency is on the order of magnitude of a typical renewable resource such as wind or PV. Higher and lower frequencies were also tested successfully, with good magnitude and phase response.

The filtered controlled power matches the control signal well between 14 and 16 h, while between 12 and 14 h the system cannot match the load exactly when a load reduction is called for. It is interesting to note that the capacity for load increase is significantly higher than that for load reduction, which occasionally is close to zero. This is a consequence of the fact that the call for cooling is lower than the design maximum for the systems. It is also interesting to note that there is a periodicity in the power increase or reduction capacity, induced by the use and release of AC compressors.

Now consider a fleet of 1000 water heaters of the heat pump variety. With a typical heat storage capacity of 42 MJ (obtained using a typical capacity of 250 l of water with a 40 °C temperature difference between charged and discharged state), and assuming a COP of 2.2 [41], 1000 units are equivalent to an electric storage system of 5.3 MWh capacity, and therefore useful for distribution-level load shifting. For example, it could be useful to move electric load from the early part of the morning (i.e. when heaters would normally recharge after morning showers, to the middle of the day, when maximum PV-derived electricity is usually available. To achieve this, in addition to the aggregated control capability, it is also necessary to prevent the heaters from recharging as soon as the SOC reaches the lower deadband setpoint of 0.9. The simplest way to do this is to schedule the lower deadband setpoint as a function of time of day. The operation of the water heaters, in the context of the other appliances, is shown in Fig. 14, for the case of constand deadband and for the case of variable deadband, with no other external control applied.

The variable deadband case presents several clear differences in comparison to the constant deaband case, notably:

- at 4:00 h, when the heaters are charged in preparation for morning showers,
- in the period from 8:00 to 10:00 h, when water heating loads are reduced by a factor of over two,
- in the period from 16:00 to 20:00 h, when loads are reduced by a factor of approximately two,
- in the period from 20:00 to 22:00 h, when loads are increased.

The rebound observed at 4:00 h in the variable deadband case is a well-known effect that can be dealt with using well-established methods [42–44].

Operating in the variable deaband mode, control is then implemented in a similar way as for the AC units, for the purpose of adding a controlled load of up to 500 kW in the period from 11:00 to 14:00 h, and of removing a late-afternoon load between the hours of 14:45 and 18:15. Two strategies are are adopted: in the first strategy, an additional hot water control signal is ramped up linearly to 500 kW between 11:00 and 11:30 h, then maintained constant between 11:30 and 13:30, then ramped down to zero linearly between 13:30 and 14:00 h; subsequently, a load control signal is ramped down to -250 kW between 15:45 and 16:00 h, maintained constant until 18:00 h, and ramped up to 0 kW between 18:00 and 18:15 h. In the second strategy, the initial load addition is identical, while the late afternoon load reduction is implemented by simply by adding a control signal of 0 kW, which has the effect of removing intermittency in the frequency range considered. The outcome of the two strategies is shown in Fig. 15.

The second strategy appears to be more successful than the first, in the sense that the late afternoon load reduction of 250 kW only lasts approximately 30 min, rather than the desired 2.5 h. The reason is simple: there are no heaters that can be turned off. On the other hand, the second strategy appears to be largely successful in reducing intermittency for the majority of the 2.5 h between 15:45 and 18:15 h. It should be noted, however, that the act of adding a controlled load in the middle of the day automatically has the effect of reducing load later in the afternoon, since water draws utilize previously stored energy.

5. Load shedding in critical conditions

Utility distribution feeders with microgrid-like characteristics are currently being considered as elements of a more resilient grid. Examples of such feeders are the one at Borrego Springs in California [45], and Ameren's Technology Applications Center microgrid in Illinois [46]. In addition to distributed generation and

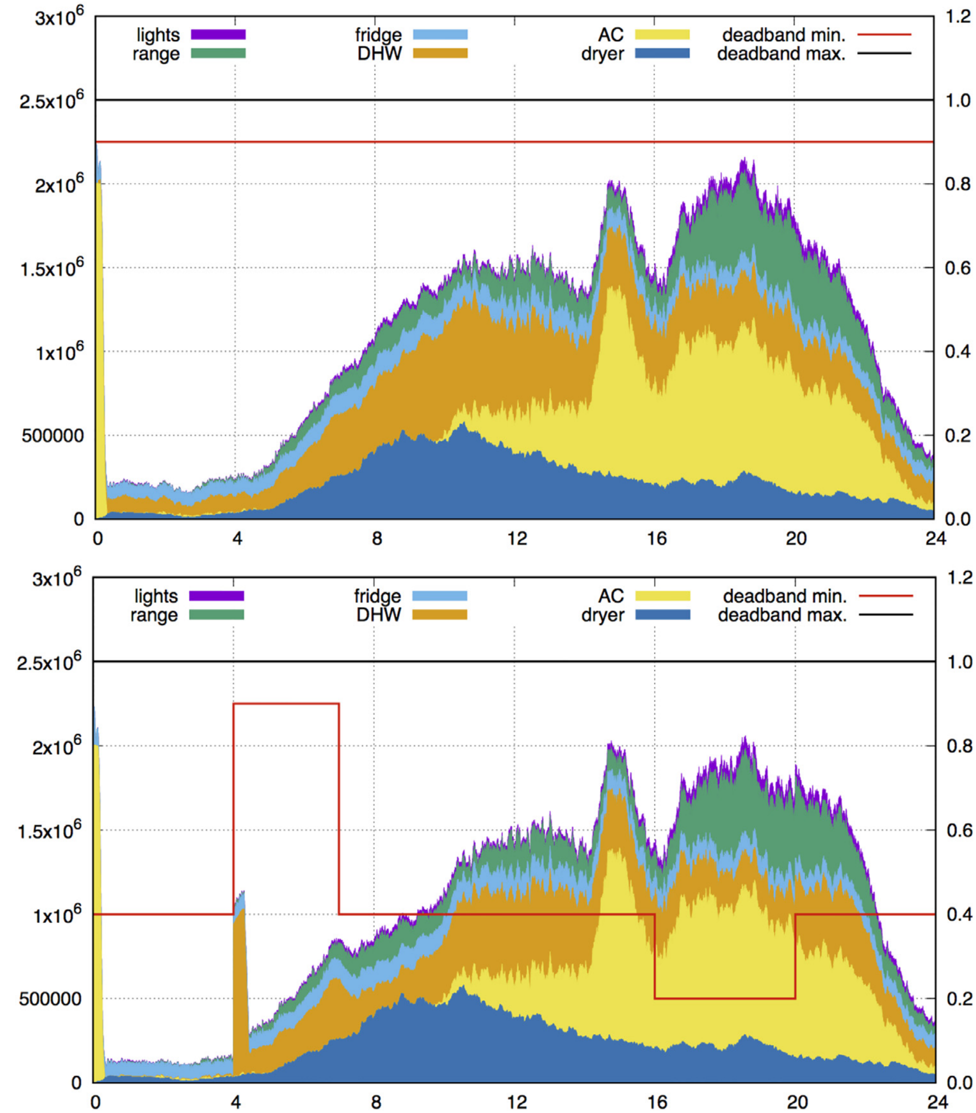

Fig. 14. Load profiles for a single weekday, with constant deaband (top) and variable deadband (bottom). The DHW load is shown in orange. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

storage, distribution-level controls manage loads via a combination of technologies including smart breakers, smart meters and home energy management systems. Utility distribution microgrids can island from the bulk grid and deliver services to customers solely from feeder-sited generation and storage resources, owned by customers, third parties or by the utility itself. In the case of extended periods of islanding from the bulk grid, it may be necessary to curtail some of the loads to enabled continued supply of power to critical loads, including medical facilities, food storage and distribution facilities or water services [25,47].

The results from the Texas and New York survey indicate that approximately 80% of respondents are either likely or certain to limit certain energy-consuming to assist the grid operator to preserve power service. Such curtailment could be enacted in several ways, depending on the energy service in question. For example, loads from air conditioning could be reduced either by automatically resetting a temperature setpoint on a smart thermostat, or by asking the customer to do so, for example via a smart phone app. Similarly, curtailment of hot water energy demand could be implemented by asking customers to reduce the amount of showers or other hot water uses, or by automatically changing the temperature setpoint of the hot water heater.

To illustrate the effect of load curtailment measures based on a combination of behavioral changes and automated schemes implemented on individual appliances, load curtailment requests were enacted between 8 and 11 and between 16 and 19 on day 1 (a Sunday). For participating households (assumed 75% of the total, based on survey results [48]), the load curtailment on each energy service during a curtailment event was implemented as follows:

1. Laundry (clothes dryer): use of the clothes dryer is reduced to 10% of the base value as a result of behavior modification prompted by a smart phone request from the distribution operator;
2. Space heating & cooling (AC unit): the temperature setpoint is raised or lowered by 3 °C for cooling or heating conditions respectively, implemented via direct command to smart thermostat or behavior modification;
3. Domestic hot water (heat pump water heater): hot water draws (primarily from showering) are reduced to 30% of their original value, as a result of a request from the distribution operator via smart phone resulting in behavior modification;
4. Food refrigeration (refrigerator): no change;

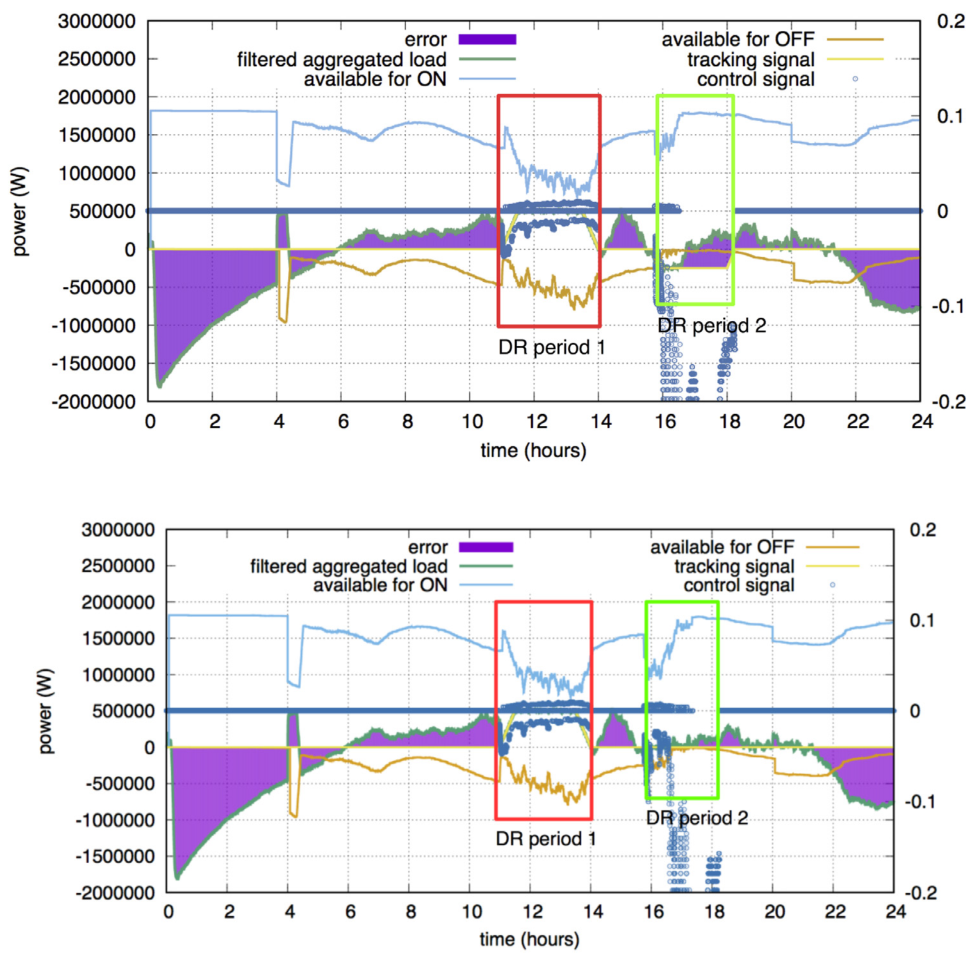

Fig. 15. Operating parameters for DHW load control for a single weekday, for the case of variable SOC deadband. The aggregated load is filtered using a time series digital filter with lower and upper setpoints of 0.00005 s^{-1} and 0.02 s^{-1}. The control gain is set to zero except of the periods of interest, between 11:00 and 14:00 h, and between 15:45 and 18:15 h. Note that, while in both cases control in the first period is successful (the yellow control signal curve coincides with the green response curve and the error is zero), control in the second period is only partially successful due to the absence of heaters that can be turned off. The second strategy, in which only the intermittency is controlled in the second period, is more successful, as indicated by the smaller error (purple shading). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

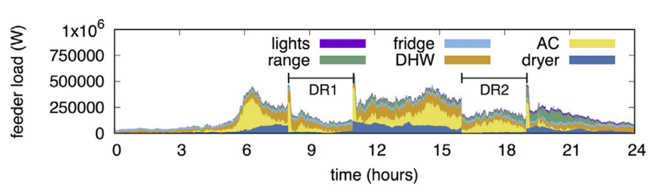

Fig. 16. Load profiles by energy service for curtailment events between 8 and 11 h and between 16 and 19 h.

5. Cooking (electric range): the frequency of cooking events is reduced to 50% of its original value as a results of behavior modification resulting from a request by the distribution system operator, delivered via smart phone;
6. Lighting (LED lights) the combined frequency and duration of 'lights on' events is reduced to 50% of its original value as a results of behavior modification resulting from a request by the distribution system operator, delivered via smart phone.

The curtailment in use of the energy services is based on a reasonable interpretation of survey and focus group results, explained in detail by Chen et al. [48] and by Abreu et al. [28]. The outcome of the load curtailment request is shown in Fig. 16, for two three-hour events, the first between 8 and 11 (DR1), and the second between 16 and 19 (DR2). During curtailment events, the clothes dryer power use is negligible, as expected. The AC load, resulting from cooling in both cases, is reduced substantially. The spikes at the beginning and end of the curtailment event are due to the forced sudden change in temperature deadbands, and could be eliminated simply by letting the temperature drift naturally at the beginning of curtailment, and by preventing all AC units from turning on together at the end of the event. The end of curtailment rebound is a well-known effect of DR schemes and can be reduced using well-established methods [44]. The DWH load is reduced in both DR events, but more so in the afternoon one, due to expected higher use. The refrigerator load is unchanged, as would be expected. The cooking load is reduced substantially as a result of behavioral change (customers use alternative food preparation strategies during curtailment events). Lighting load, although very small to start with, is reduced due to customers paying more attention to energy waste.

The simulations show that load can be reduced to approximately one half of the regular value by resorting to customer behavior change, assisted by automation where possible, for periods of time lasting on the order of hours. Whether such reductions could be sustained for several days is an open question, that requires more detailed human behavior models that account for the ability of users to go without energy services, as well as interactions with other users and the system interface, i.e. social factors. While validation is not practical in this case due to lack of infrastructure, the results have a basis in reality. The information that can be provided by models of this type can enable the designer of a community microgrid to explore scenarios quantitatively with better confidence than just by arbitrary assumptions.

6. Discussion and conclusions

A bottom-up model of residential electrical load resulting from uses of energy services, that includes both physical and behavioral models, is combined with realistic strategies for high-temporal-resolution load control and load shedding, in a complete framework designed for co-simulation using power flow simulation tools. The bottom-up load simulation methodology is similar other efforts, for example the SynPro tool developed by Fisher et al. [11]. Unlike the case of other studies, where the emphasis was to produce accurate load simulations for certain settings, the emphasis of the work presented here was to demonstrate the effectiveness of control of aggregated resources *at the scale of a small community* for the purpose of real-time control and emergency load shedding, in a realistic setting that incorporates user behavior. The ability to simulate individual loads realistically is critical to conducting simulations of power flow on a distribution feeder, especially in the case where advanced demand-response mechanisms are studied, or where the effects of high penetration of distributed resources are of interest. Both use cases, real-time control and load shedding, are critical for community microgrid applications. A detailed quantitative calibration of the models, followed by a validation step, is not the purpose of this work, as this has been done before. In contrast, it is shown here that it is possible to obtain realistic behavior predictions from well-designed surveys that include hitherto unexplored questions on willingness of energy users to forego full control of their appliances to benefit the performance of their community's grid, or even to allow its stability in the face of unusual circumstances. Moreover, it is also shown that the level of response of both real-time control and load shedding is substantial and can make a substantial impact on the controllability and stability of the system, at much lower cost than simply by using battery storage.

The ability to use air conditioners, controlled via smart thermostats, to offset variability due to high-penetration PV, or simply to improve the behavior of the load at the feeder, has been demonstrated in the context of operations of a variety of appliances. The energy contained in the air mass and the building structure provides enough reserve to control loads over a range of frequencies, ranging from $0.003\,s^{-1}$ to $0.0005\,s^{-1}$. Response at higher frequencies is limited by the switching dead-time of the AC compressors, while response at lower frequencies is limited by the ability of the structures and air mass contained within them to store sensible energy.

For the case of heat-pump water heaters, it was shown that the collective ability to absorb substantial amount of energy (over $1\,MWh_e$) in a highly controlled way can be achieved with relatively low control effort. This could be useful, for example, to absorb high levels of PV electricity production around solar noon, thereby reducing reverse power feed from the feeder into the transmission grid (the belly of the 'duck curve'). Control of the release of the accumulated energy in the afternoon is more difficult, due to the unavailability of devices that can be turned off. However, the afternoon load is naturally reduced compared to the case where peak energy absorption if not implemented. To achieve better control, more sophisticated algorithms could be implemented, however this is beyond the scope of the present study.

Finally, the framework presented here can easily be used in conjunction with distribution system simulations such as GridLAB-D or OpenDSS, or with real-time hardware-in-loop systems, for the purpose of sizing and optimizing system components and for designing protection systems. The model presented here applies to residential loads only, while in a typical situation, commercial and sometimes light industrial loads could also be present on the feeder. These could also be co-simulated alongside the residential loads using relevant models.

Conflict of interest

None.

Acknowledgments

This research was supported in part by the Mitsubishi Research Institute, who is entrusted by New Energy and Industrial Technology Development Organization, under contract MIRI/EEU 46-001, in part by the U.S. National Science Foundation, under awards 1541148 and 1541117, and in part by the Engineering Research Center Program of the NSF and the Department of Energy under NSF award EEC-1041877 and the CURENT Industry Partnership Program.

References

[1] Energy flow charts, 2015, [Accessed 16 February 2017]. https://flowcharts.llnl.gov/commodities/energy.
[2] J.A.P. Lopes, F.J. Soares, P.M.R. Almeida, Integration of electric vehicles in the electric power system, Proc. IEEE 99 (1) (2011) 168–183.
[3] J.H. Williams, A. DeBenedictis, R. Ghanadan, A. Mahone, J. Moore, W.R. Morrow, S. Price, M.S. Torn, The technology path to deep greenhouse gas emissions cuts by 2050: the pivotal role of electricity, Science 335 (6064) (2012) 53–59.
[4] M.-L. Chan, E. Marsh, J. Yoon, G. Ackerman, Simulation-based load synthesis methodology for evaluating load-management programs, IEEE Trans. Power Appar. Syst. (4) (1981) 1771–1778.
[5] T. Calloway, C. Brice, Physically-based model of demand with applications to load management assessment and load forecasting, IEEE Trans. Power Appar. Syst. (12) (1982) 4625–4631.
[6] M. Kamgarpour, C. Ellen, S.E.Z. Soudjani, S. Gerwinn, J.L. Mathieu, N. Müllner, A. Abate, D.S. Callaway, M. Fränzle, J. Lygeros, Modeling options for demand side participation of thermostatically controlled loads, in: Bulk Power System Dynamics and Control-IX Optimization, Security and Control of the Emerging Power Grid (IREP), 2013 IREP Symposium, IEEE, 2013, pp. 1–15.
[7] A. Capasso, W. Grattieri, R. Lamedica, A. Prudenzi, A bottom-up approach to residential load modeling, IEEE Trans. Power Syst. 9 (2) (1994) 957–964.
[8] C. Walker, J. Pokoski, Residential load shape modelling based on customer behavior, IEEE Trans. Power Apparat. Syst. (7) (1985) 1703–1711.
[9] M. Muratori, M.C. Roberts, R. Sioshansi, V. Marano, G. Rizzoni, A highly resolved modeling technique to simulate residential power demand, Appl. Energy 107 (2013) 465–473.
[10] J. Widén, M. Lundh, I. Vassileva, E. Dahlquist, K. Ellegård, E. Wäckelgård, Constructing load profiles for household electricity and hot water from time-use datamodelling approach and validation, Energy Build. 41 (7) (2009) 753–768.
[11] D. Fischer, A. Härtl, B. Wille-Haussmann, Model for electric load profiles with high time resolution for german households, Energy Build. 92 (2015) 170–179.
[12] Y. Ge, K. Qian, J. Dai, C. Zhou, Modelling of domestic load demand in the presence of microgrid with wind and photovoltaic resources, in: Smart Grid and Clean Energy Technologies (ICSGCE), 2015 International Conference on, IEEE, 2015, pp. 180–186.
[13] K. du Preez, H. Vermeulen, Development of a bottom-up methodology for modelling electrical residential load from event data, in: Domestic Use of Energy (DUE), 2016 International Conference on the, IEEE, 2016, pp. 1–7.
[14] P. Gotseff, B. Lundstrom, Data-driven residential load modeling and validation in gridlab-d, in: Green Technologies Conference (GreenTech), 2017 Ninth Annual IEEE, IEEE, 2017, pp. 20–25.
[15] M. Hayn, V. Bertsch, W. Fichtner, Electricity load profiles in Europe: the importance of household segmentation, Energy Res. Soc. Sci. 3 (2014) 30–45.
[16] I. Richardson, M. Thomson, D. Infield, C. Clifford, Domestic electricity use: a high-resolution energy demand model, Energy Build. 42 (10) (2010) 1878–1887.
[17] W.J. Cole, J.D. Rhodes, W. Gorman, K.X. Perez, M.E. Webber, T.F. Edgar, Community-scale residential air conditioning control for effective grid management, Appl. Energy 130 (2014) 428–436.
[18] A. Chrysopoulos, C. Diou, A.L. Symeonidis, P.A. Mitkas, Bottom-up modeling of small-scale energy consumers for effective demand response applications, Eng. Appl. Artif. Intell. 35 (2014) 299–315.

[19] S. Bashash, H.K. Fathy, Modeling and control of aggregate air conditioning loads for robust renewable power management, IEEE Trans. Control Syst. Technol. 21 (4) (2013) 1318–1327.
[20] S. Shao, M. Pipattanasomporn, S. Rahman, Development of physical-based demand response-enabled residential load models, IEEE Trans. Power Syst. 28 (2) (2013) 607–614.
[21] Eload database documentation and evaluation and use in nems, 2001, http://www.onlocationinc.com/LoadShapesReload2001.pdf.
[22] A.J. Collin, G. Tsagarakis, A.E. Kiprakis, S. McLaughlin, Development of low-voltage load models for the residential load sector, IEEE Trans. Power Syst. 29 (5) (2014) 2180–2188.
[23] J. Iria, F. Soares, A. Madureira, Advanced models and algorithms for demand participation in electricity markets, in: North American Power Symposium (NAPS), 2014, IEEE, 2014, pp. 1–6.
[24] A. Jambagi, M. Kramer, V. Cheng, Residential electricity demand modelling: activity based modelling for a model with high time and spatial resolution, in: Renewable and Sustainable Energy Conference (IRSEC), 2015 3rd International, IEEE, 2015, pp. 1–6.
[25] V. Ayon, M. Robinson, A. Mammoli, A. Fisher, J. Fuller, Integration of bottom-up statistical models of loads on a residential feeder with the gridlab-d distribution system simulator, and applications, in: 2017 IEEE PES Innovative Smart Grid Technologies Conference Europe, IEEE, 2017.
[26] R.K. Vieira, D.S. Parker, J.F. Klongerbo, J. Stone, J. Cummings, How contractors really size air conditioning systems, in: Proceedings of the 1996 ACEEE Summer Study on Energy Efficiency in Buildings, Vol. 8, 1996, pp. 207–213.
[27] B. Sparn, K. Hudon, D. Christensen, Laboratory Performance Evaluation of Residential Integrated Heat Pump Water Heaters, Technical Report, National Renewable Energy Lab. (NREL), Golden, CO (United States), 2014.
[28] J. Abreu, C. McIlvennie, N. Weingartz, A. Mammoli, What would you do to keep the lights on in your community? ACEEE Summer Study Energy Efficiency in Buildings, American council for an Energy Efficient Economy, 2018.
[29] X. Xu, C.-f. Chen, X. Zhu, Q. Hu, Promoting acceptance of direct load control programs in the United States: financial incentive versus control option, Energy 147 (2018) 1278–1287.
[30] D. Egarter, V.P. Bhuvana, W. Elmenreich, Paldi: online load disaggregation via particle filtering, IEEE Trans. Instrum. Meas. 64 (2) (2015) 467–477.
[31] S.M. Tabatabaei, S. Dick, W. Xu, Toward non-intrusive load monitoring via multi-label classification, IEEE Trans. Smart Grid 8 (1) (2017) 26–40.
[32] Residential energy consumption survey (RECS), Technical Report, U.S. Energy Information Administration, 1000 Independence Ave., SW Washington, DC 20585, 2015.
[33] M.H. Albadi, E. El-Saadany, A summary of demand response in electricity markets, Electr. Power Syst. Res. 78 (11) (2008) 1989–1996.
[34] J.S. Vardakas, N. Zorba, C.V. Verikoukis, A survey on demand response programs in smart grids: pricing methods and optimization algorithms, IEEE Commun. Surv. Tutor. 17 (1) (2015) 152–178.
[35] W. Zhang, J. Lian, C.-Y. Chang, K. Kalsi, Aggregated modeling and control of air conditioning loads for demand response, IEEE Trans. Power Syst. 28 (4) (2013) 4655–4664.
[36] J. Mathieu, M. Dyson, D. Callaway, A. Rosenfeld, Using residential electric loads for fast demand response: the potential resource and revenues, the costs, and policy recommendations, in: ACEEE Summer Study on Energy Efficiency in Buildings, Citeseer, 2012, pp. 189–203.
[37] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, Numerical Recipes in Fortran 77: The Art of Scientific Computing, Vol. 1, Second ed., Cambridge University Press, 1996.
[38] J.L. Mathieu, S. Koch, D.S. Callaway, State estimation and control of electric loads to manage real-time energy imbalance, IEEE Trans. Power Syst. 28 (1) (2013) 430–440.
[39] A. Molina-Garcia, M. Kessler, J.A. Fuentes, E. Gomez-Lazaro, Probabilistic characterization of thermostatically controlled loads to model the impact of demand response programs, IEEE Trans. Power Syst. 26 (1) (2011) 241–251.
[40] Y. Yasaei, A.A. Mammoli, A novel framework for characterizing the aggregated response of thermostatically controlled loads in distribution networks, in: Power and Energy Society General Meeting (PESGM), 2016, IEEE, 2016, pp. 1–5.
[41] B. Sparn, K. Hudon, D. Christensen, Laboratory Performance Evaluation of Residential Integrated Heat Pump Water Heaters, Technical Report, National Renewable Energy Laboratory, 2014.
[42] Y. Li, B.L. Ng, M. Trayer, L. Liu, Automated residential demand response: algorithmic implications of pricing models, IEEE Trans. Smart Grid 3 (4) (2012) 1712–1721.
[43] M. Muratori, G. Rizzoni, Residential demand response: dynamic energy management and time-varying electricity pricing, IEEE Trans. Power Syst. 31 (2) (2016) 1108–1117.
[44] A. Khan, N. Javaid, A. Ahmad, M. Akbar, Z.A. Khan, M. Ilahi, A priority-induced demand side management system to mitigate rebound peaks using multiple knapsack, J. Ambient Intell. Humaniz. Comput. (2018) 1–24.
[45] W. Feng, M. Jin, X. Liu, Y. Bao, C. Marnay, C. Yao, J. Yu, A review of microgrid development in the united states–a decade of progress on policies, demonstrations, controls, and software tools, Appl. Energy 228 (2018) 1656–1668.
[46] G. Vaidyanathan, Core concept: microgrids offer flexible energy generation, for a price, Proc. Natl. Acad. Sci. 115 (17) (2018) 4298–4300.
[47] V. Ayon, M. Robinson, A. Mammoli, Simulation of real-time demand-response to support a residential distribution feeder in microgrid mode, 2017 Cigrè Grid of the future Symposium, Cigrè, 2017.
[48] Q. Shi, C.-F. Chen, A. Mammoli, F. Li, Estimating the profile of incentive-based demand response (IBDR) by integrating technical models and social- behavioral factors, IEEE Trans. Smart Grid. (2019). *Under review*