COMMUNITY SCALE DYNAMIC BUILDING SIMULATIONS USING PROBABILISTIC BEHAVIOURAL PATTERNS

P. McCallum[1], D.P. Jenkins[1], and A.D. Peacock[1]

[1]Heriot-Watt University, Edinburgh, UK

ABSTRACT

The upscaling of energy models of buildings has entered mainstream discussion within the subject of energy system modelling (ESM), the primary role of such models being to generate policy; however, policy holds the potential to stimulate significant changes in energy demand, especially within the residential building sector. This cyclic system is not modelled fully at administrative levels in the UK. The mechanisms to infer useful insight into future demand caused by millions of new heat-pumps and electric vehicles, for example, must be implemented as an integrated part of the modelling process.

The move away from data-derived demand curves is exemplified in the present work. This focuses on a small community and explores methodologies to provide scalable solutions to characterise residential thermal demand. Despite being inherently deterministic, the tools employed have been configured here to run sequences of probabilistic inputs to deliver aggregate loads for a diverse building stock. The dwellings in this study have been classified into their respective archetypes based on building form and construction. Smart meter data have then been used to generate behavioural patterns which describe how the dwellings are used. Finally, probability distributions have been applied to the behavioural patterns to consider variability across the sub-groups within the stock.

INTRODUCTION

Scalable physics-integrated demand models of building stock is widely recognised as a key ambition within the 'whole-system' approach to ESM. Moreover, temporally precise dynamic loads are of critical importance when attempting to address questions around variability in supply from renewable generators, as well as utilisation of energy storage and demand response. The potential for unpredictable scenarios involving variation in supply on the electricity network, along with new demand patterns emerging from widespread use of Electric Vehicles (EVs) and electric heat pumps, has raised serious questions around energy security. Maintaining a balance with environmental and socio-economic factors requires careful planning, along with new approaches to modelling energy demand.

These scenarios are of concern at different scales. National-scale systems are typically assessed using Energy System Optimisation Models (ESOMs). These are designed to ensure that peak demands are met under various scenarios driven by environmental and costs factors (i.e. security is unconditional). Their ability to do this, however, is undermined when demand is characterised using over-simplified temporal schemes. Even when intra-day variation in demand is considered, this is derived from historical data, and cannot provide any insight into ongoing energy security for uncertain demand scenarios.

Community-scale systems can take advantage of more refined Energy System Simulation Models (ESSMs), due to their reduced size and complexity. Because the domain of such systems is significantly smaller, operational behaviour can be modelled with greater detail, implying that events which could cause system failures can be planned for and avoided. However, the success of this predictive method is once again closely linked to the applicability of the input demand timeseries.

The present work examines community-scale thermal demand modelling, using a framework intended for upscaling to much larger regions. Two variants of the procedure were tested and compared against each other, one of which was orientated towards accuracy, the other towards computational efficiency. An archetype system was adopted to model the Findhorn Ecovillage community system, which incorporates both large and small-scale renewable generators, a private-wire electricity network and community-owned utility provider. Due to the established nature of this community system, parallels can be drawn with Government interests in energy planning, economics, security and sustainability. Details of the site are discussed in the Methodology section below, along with descriptions of the Dynamic Building Model (DBM) and simulation routines. Diversified, aggregate outputs from the models are provided in the Results section. Finally, the Discussions and Conclusions section follows.

This study involving thermal demands from residential buildings at Findhorn Ecovillage is part of wider work being undertaken for the National Centre for Energy Systems Integration project (CESI, grant EP/P001173/1).

BACKGROUND AND LITERATURE

The issues introduced above have been discussed by a number of authors. These problems concern space heating, hot water loads and non-thermal loads (which are generally classed as electrical demands). Furthermore, they also span across domestic and non-domestic building stock. The existing literature generally approaches the same set of challenges, which is to characterise demands in a temporally precise manner and to aggregate these demands whilst incorporating diversity.

Simulation and validation of temporally refined aggregate electrical demand has received significant attention over a number of years. Examples of previous studies include application of statistical methods to generate synthetic profiles (Jenkins et al., 2014, Patidar et al., 2016). Other earlier work has also examined the usage of specific household appliances (Richardson et al., 2010, Widén and Wäckelgård, 2010). A further area of study which is closely related is the research surrounding occupant behaviour, which again features work examining rational behaviour in terms of interaction with devices, as well as statistical processes (Richardson et al., 2008, McKenna et al., 2015, Flett and Kelly, 2016, Liang et al., 2016, Flett and Kelly, 2017, Ramírez-Mendiola et al., 2017, Zhou et al., 2017).

Other work directly addresses aggregate and diversified thermal demand, either in isolation, or as part of a multi-vector model. Cipriano et al. (2015) carried out a study using validation data from an unoccupied office building, by performing batch simulations using a DBM, and applying probability distributions on unknown parameters. Good et al. (2015) examined the operational behaviour of systems, including that of heat pumps for residential use. Miller et al. (2015) demonstrated a method for clustering measured data to infer thermal requirements for a school campus and office. McKenna and Thomson (2016) present a thermal-electrical model, with comparisons against gas smart meter data representative of UK building stock. Other interesting developments include a reduced-order method by Heidarinejad et al. (2017), which was applied to a number of university campuses in the US.

METHODOLOGY

Site topology of community

Site data was available for Findhorn Ecovillage, located in Moray in north east Scotland (57.65°N, 3.59°W). The data included electrical mains power for around 30% of the dwellings on site, along with smaller subgroups incorporating sub-metered power for heating/hot water use, measurements associated with distributed PV and solar hot water generation, demands on a number of small district heating systems, outputs from a community owned wind turbine array, loads from community/commercial buildings, on-site weather and aggregate substation data for the entire community. Monitoring began in mid-2014; a number of sensors are still active, although in general, the dataset is partially complete.

Figure 1 provides details of the site in terms of building use (residential and non-residential) and classification of the thermal envelope. The present work applies only to the residential portion of the building stock, with the principle aim being to establish and test methodologies to simulate aggregate heat demand, which are also scalable for future application to larger geographical regions. An archetype system was adopted to represent the physical form of each dwelling, which respected the construction type and approximate floor area. The archetype groups were also generated in a way that ensured that at least one of the dwellings had corresponding field measurements.

Figure 1: Findhorn Ecovillage, site map showing (a) building use in terms of residential and community / commercial and (b) thermal envelope construction type.

Simulation model

The dynamic thermal simulations were carried out using the commercial software IES-VE (2018). This work was built specifically around the Python VEScript Application Programming Interface (API). A single model file was constructed in the usual manner (incorporating all geometric features, construction properties and infiltration rates, weather / location settings and simulation options), with this processed using the API through a large number of control profiles. Reference profiles were constructed and stored outside IES-VE and were adapted automatically within Python during the batch runs.

A simplified model topology was adopted to represent the community; simplified volumes were generated for each archetype. One building was constructed to represent each archetype, using one room per level to respect floor area. This overlooked the real topology of the site and placed the modelled volumes on a widely-spaced uniform grid, designed to prevent self-shading. Glazing areas were estimated based on typical percentages. No roof voids were included.

Simulation routine

The make-up of the aggregate thermal demand timeseries was built around a sequence of stages:

1. A series of archetypes were constructed to represent the diversity in the building, similar to that which might be used in other stock models (based on construction information, heating technology etc). Each archetype was assigned a weighting.
2. Diversity in behaviour was accounted for using a set of input profiles, which characterised the control routines dictated by the households of each dwelling. This took the form of thermostat set-point schedules, which automatically control the heating systems (if temperatures dictated that heating was necessary). It should be borne in mind that the intention is to produce an aggregated demand profile of many dwellings. This allows for some flexibility in the choice of behavioural profiles; where, for example, the use of more time-averaged (and therefore smoother) profiles can still be justified – see also note on validation in Discussions/Conclusions.
3. Each archetype-profile pairing was simulated repeatedly using the specifically configured batch routine, each applying a modified succession of input control profiles (i.e. offset in time, in 10-minute increments). Aggregate, synthesised thermal demands were generated through compound weightings, applied to the large set of simulation results. The weightings incorporated the stock weighting, profile weighting and probability weighting associated with the temporal diversity in the collective behaviour of the community.

Model calibration

The relevant transient building simulation inputs comprise the various weather timeseries (a single case has been set to suit the current location, AberdeenEWY) and the heating set-point control profiles associated with each dwelling. Appropriate calibration of the latter is essential for delivering meaningful thermal demand timeseries. For the present work, information regarding such cycles has been extracted from smart meter data to calibrate the model inputs. This includes identification of typical heating set-point control patterns, along with a corresponding likelihood of use. The procedure adopted generated a very large number of input profiles, before carrying out iterations of the simulation using those profiles.

Figure 2 provides a description of the procedure used to extract calibrated simulation inputs from the available smart meter data. Both heat and electricity meter data were used. On the left of Figure 2, heat meter data was studied to extract typical control profile shapes (stage 1), which were then used to generate the large set of input profiles (stage 5). Intervening procedures were carried out on the electricity meter data, which first generated logs of activity information from the data (stage 2), then extracted statistical information regarding the day-to-day activity variation within each dwelling (stage 3) and the corresponding dwelling-to-dwelling variation across the community (stage 4).

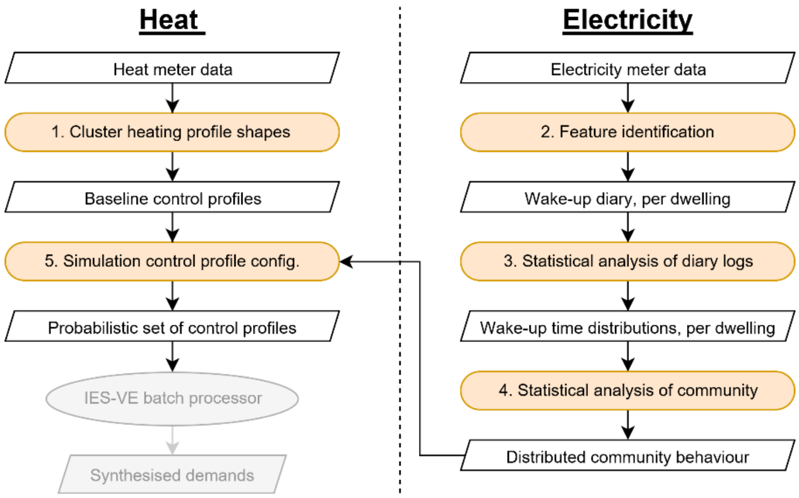

Figure 2: Process for the extraction of probabilistic thermal control profiles from smart meter data.

It should be noted that in the present usage, no attempt was made to model or synthesise electrical household demands, even though electrical smart meter data was utilised. This was purely to infer building occupant behaviour, by proxy, to establish diversity across the community in the temporal tense.

Weekday heat meter records from the Findhorn site are provided in Figure 3, each covering continuous periods of 1-3 months. The six profiles represent the dominant thermal control cycles observed in the available data. These clearly follow automatic control schemes, with very few manual override exceptions. Profiles very similar to those shown in Figure 3 were found across multiple dwellings; profile weightings for the building stock were derived from the frequency of occurrence (the collective number of days) for each profile. This process corresponds to stage 1 in Figure 2.

Six weekday/weekend heating control profiles pairs were generated using the meter data in Figure 3 (consistent data was also available for weekend/holiday behaviour; in some cases, the same control patterns emerged for both weekdays and weekends). The resulting control schemes are provided in Figure 4. Each profile has been realigned to t_0 – the timing of the first heating set-point change (i.e. the start of the morning heating cycle), as obtained from the reference smart meter data.

Behavioural diversity

In addition to the use of heat meter data (see Figure 3), electricity meter data was analysed to determine: (a) how building occupant activity varied day-to-day for a given dwelling, and (b) how the corresponding variation between different dwellings could be characterised and accounted for in the model inputs. Before carrying out this two-stage statistical assessment (represented by calibration stages 3 and 4 in Figure 2), a search procedure was applied to each electricity meter dataset to identify the first daily event signifying the onset of activity, or wake-up time (t_{wakeup}), within that household (calibration stage 2). An example of this feature identification scan is provided in Figure 5(a).

Timestamp logs were gathered for each metered dwelling, with these used to generate a set of distributions describing day-to-day activity variation. A closest-fit distribution (Cauchy) was identified, returning location (x_0) and scale (γ) parameters for every dwelling. Figure 5(b) provides the distribution of t_{wakeup} for one of the dwellings studied. Despite this continuous variation in t_{wakeup}, the corresponding automatic heating control profiles were regimental and unchanged over many months.

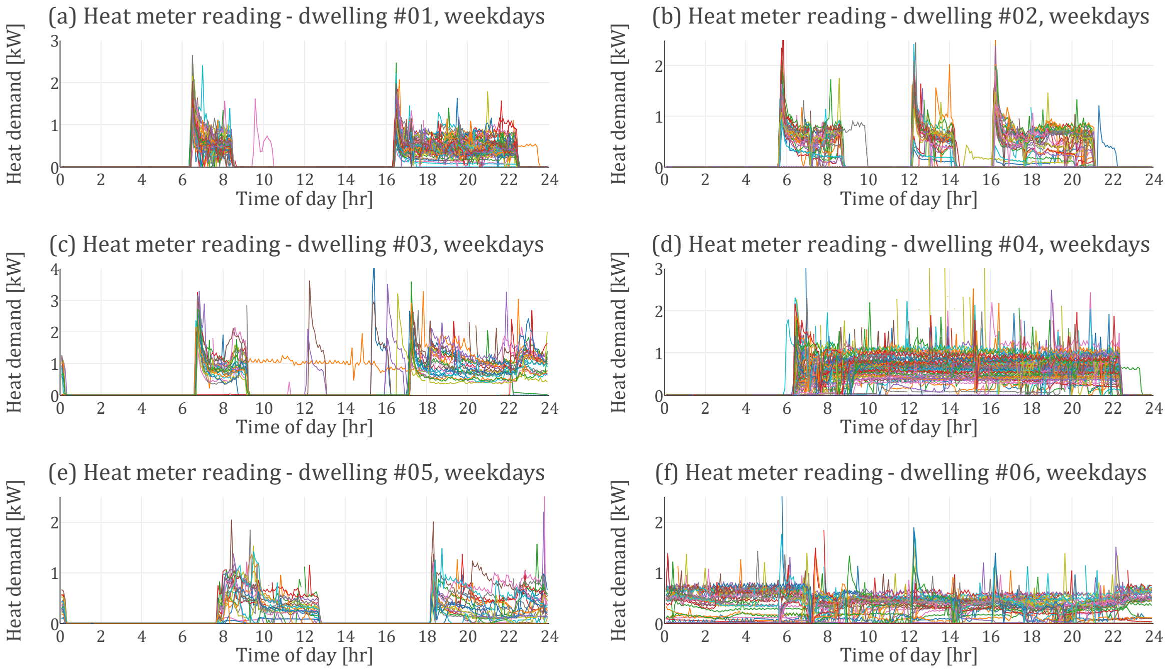

Figure 3: Weekday heat meter data from selected properties, continuous data covering 1 to 3 months.

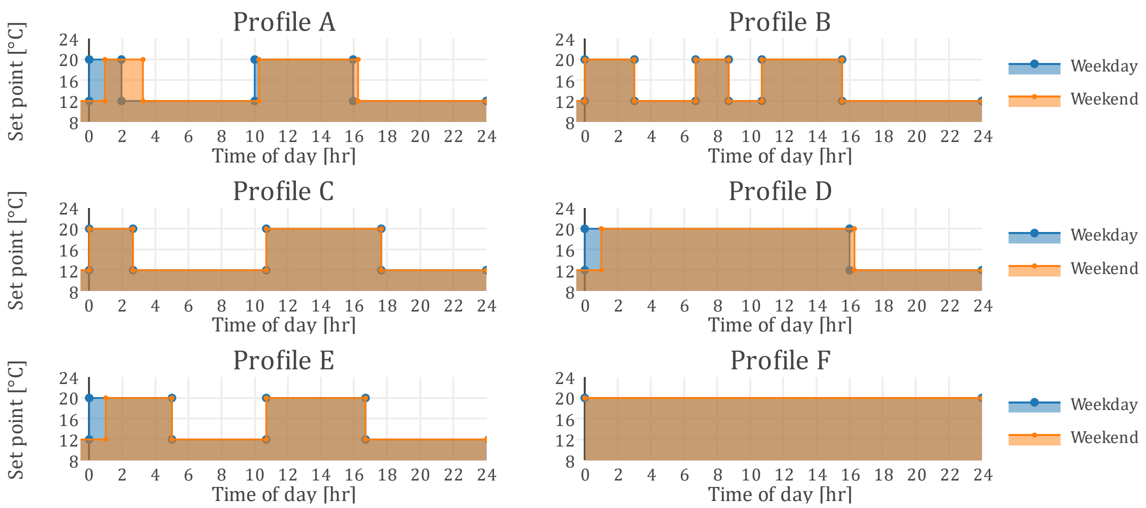

Figure 4: Input profiles accounting for building occupant behaviour. Six overall behavioural schemes were adopted (A to F), each using two different profiles for weekdays and weekends/public holidays. Set points were selected at random for each profile, within the range 18-24°C.

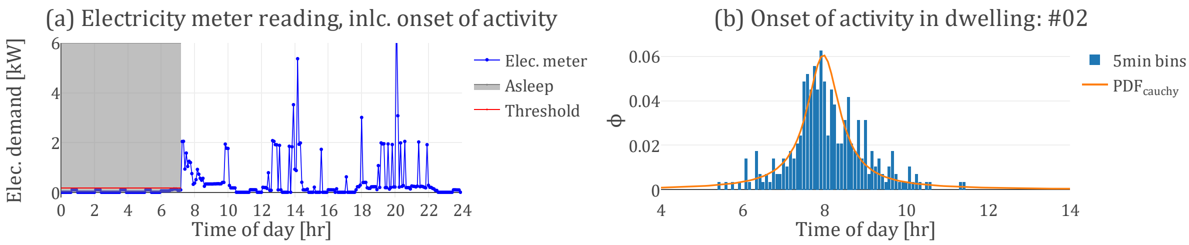

Figure 5: (a) Sample electricity meter reading for one day (calibration stage 2); (b) distribution of t_{wakeup} over captured over a 28-month period, for one of the assessed buildings (calibration stage 3).

The above set of x_0 parameters were adopted as characteristic descriptors for typical activity in the corresponding dwellings – effectively a mean wake-up time, per dwelling. It was observed that for dwellings where both the heating and electricity meter data was available, an approximate measure of the difference between x_0 and t_0 across the dwellings was 60 minutes. It follows that the distribution of x_0 across all dwellings can also be used to provide an approximate description of the variation in t_0, once adjusted according to this relationship. The resulting analysis of x_0 for households across the community is summarised in Figure 6, showing the closest-fit (normal) distribution, along with histograms based on 5-minute and 30-minute bins. An equivalent normal distribution (with an adjusted mean parameter, μ) has been used to diversify the input heating control profiles for the batch simulation.

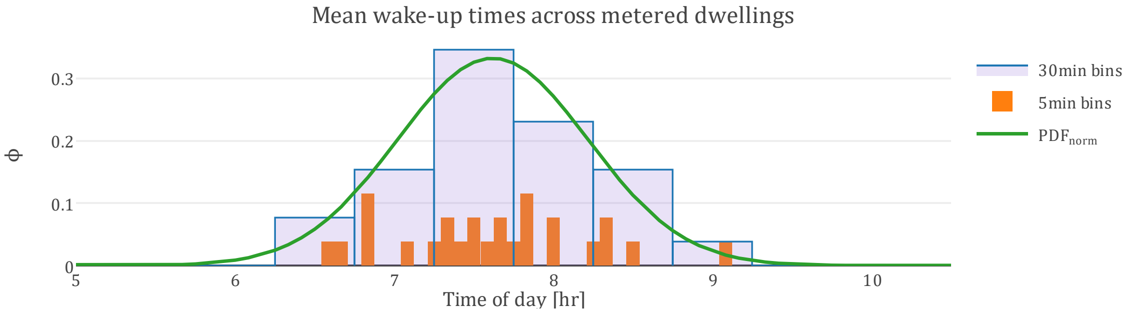

Figure 6: The distribution of wake-up times across the metered dwellings (from features identified in electricity data).

RESULTS

The output from the overall aggregation procedure provides synthesised thermal demands for the whole community (10-minute resolution timeseries for one year). All of the underlying simulation results, when considered individually, represented standard application of DBM, and were both deterministic and regimental in their behaviour on all days characterised by the same control profile. To highlight this, Figure 7(a) and 7(b) shows a single simulation result for one archetype, using Profiles A1 and A2 (7(a) provides weekdays and 7(b) the corresponding weekend, for a selected week during winter). Characteristic features include precise times for heating set-point changes, restricted heating system output to ensure realistic heat input rates, and variable heat output at other times, resulting from transient heat fluxes across the building fabric.

Figure 7(c) and (d) illustrates the probabilistic result following the diversity procedure described in the above methodology. Descriptively, this represents a medium to large set of houses corresponding to a single archetype, occupied by households which all follow an identical routine whereby the dwelling is empty during normal working hours. Very small proportions of these households are expected to wake extremely early or extremely late; the highest probability of occurrence is around the mean wake-up time (x_0). The procedure used to generate this diversity is repeated for all other archetypes and all other behavioural profiles. All the results in Figure 7 are scaled in order to correspond to a single dwelling within the aggregated building stock; however, the results in Figure 7(c, d) cannot be used to represent a single, stand-alone dwelling.

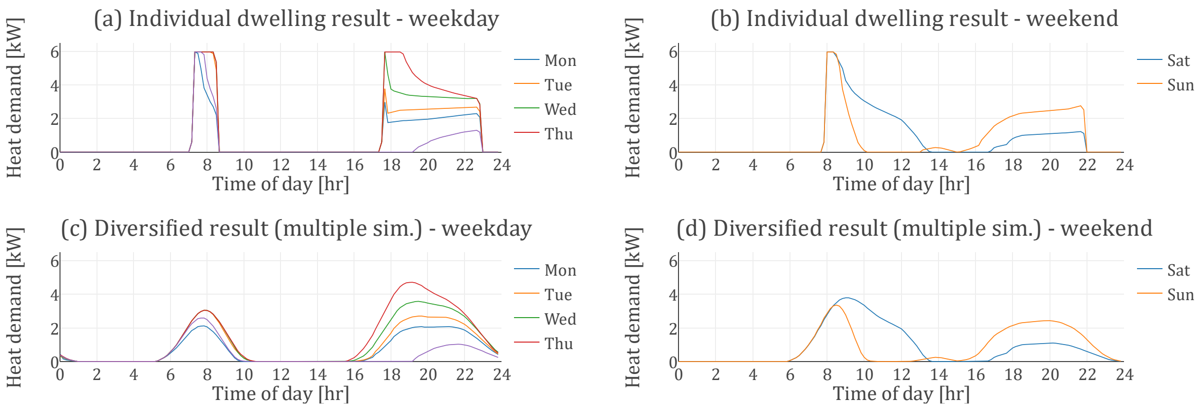

Figure 7: Result sample for a typical winter week: (a, b) Raw simulation results without temporal diversification. (c, d) Diversified results from arrays temporally offset variants of the input Profiles A1 and A2.

Figure 8(a) provides aggregated thermal demand results for the entire case-study (i.e. across all archetypes and behavioural profiles), with the corresponding external dry-bulb given in Figure 8(b) for reference. The typical UK two-peak demand shape is the most dominant pattern visible in the timeseries. During Saturday and Sunday (the two days on the farthest right), the middle part of the day shows quite different characteristics when the weather is particularly cold, which again reflects the input profiles. In general throughout the week, the severity of the trough during midday (and to a lesser degree the trough overnight) is a result of the narrow and specific focus on residential stock. Typical daily cycles of aggregated gas demand across a region, for example, will tend to include other types of non-domestic loads, such as from offices, commercial buildings, catering/service needs, and in the education sector, as well as 24-hour facilities such as healthcare and other critical services, and potential industrial loads.

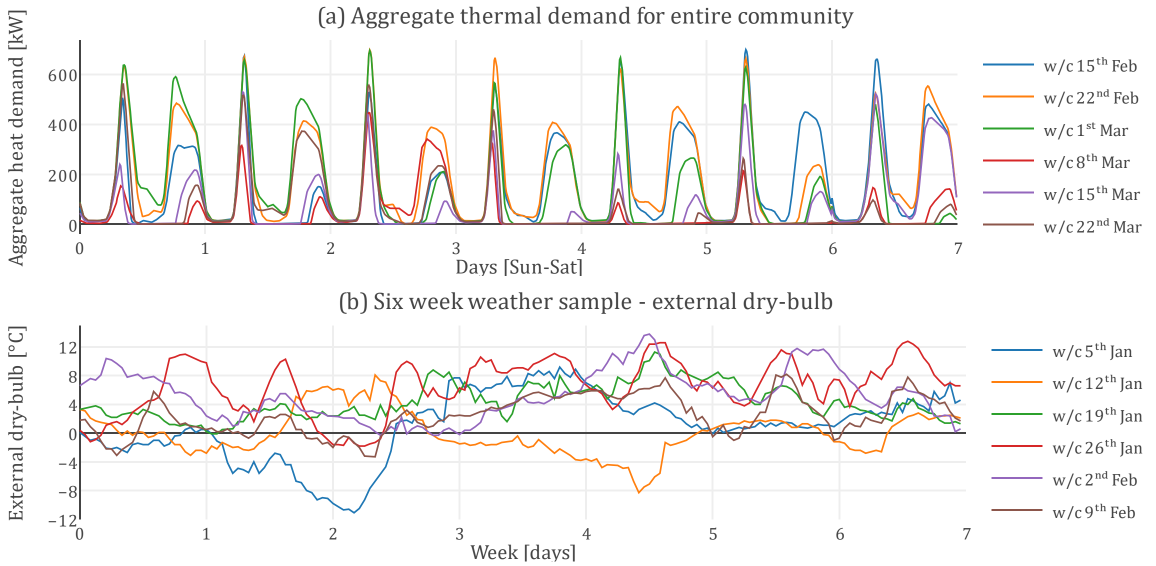

Figure 8: Sample of weather data (external dry-bulb) and corresponding aggregate heat demand for all residential building stock representing the Findhorn community (a six-week period during the winter: 5th Jan to 15th Feb).

DISCUSSIONS AND CONCLUSION

The study presented in this paper has demonstrated application of a framework and automated routine for diversification of thermal loads using dynamic building simulations. Synthesised thermal demand timeseries were generated using arrays of deterministic simulations, each characterised by a compound weighting which was linked to building stock size, behavioural profile group and a synthesised probability weighting to account for temporal diversity.

A major difficulty in carrying out such work is the validation of both input functions and the resulting aggregate output from the overall simulation routine (both are the subject of ongoing work). Regarding inputs, there is a reliance on treating the behaviour of people using repetitive, predictable and exact control functions, which defy true human nature. Stochastic or probabilistic behaviour can also be applied to DBM (the latter being the subject of this study), however, these are still constructed from large numbers of deterministic simulations which suffer from the aforementioned flaws. Furthermore, limited statistical data is available to substantiate the necessary assumptions. For widespread application of the presented framework, key data sources would include house condition surveys, census data and smart meter data; however, at present, it is extremely difficult to generate useful input functions for stochastic or probabilistic simulations.

Regarding validation of output data, challenges associated with the present work included lack of gas data, limited data describing thermal demand and widespread use of unmonitored secondary heating systems (fires and wood-burning stoves). General challenges that are non-specific to the current study include scarcity of high resolution data for space heating (not including hot water or cooking purposes). Furthermore, to provide a full validated study, validation of the output results and input behavioural functions discussed above, must occur simultaneously. There is also the well-known-issue of Performance Gap in building modelling, where purely theoretical building models rarely have a good match to empirical energy data for the same building. Such a validation exercise must therefore be taken with care, where the comparison between modelled and measured results is about returning similar overall characteristics of energy demand, rather than identical hourly demand values throughout the year. The authors intend to achieve this comparison through the use of regional gas demand data, which will have to be appropriately processed to be suitable as a proxy for thermal demand.

ACKNOWLEDGMENT

This work was funded by the National Centre for Energy Systems Integration (CESI). The authors would like to acknowledge the support of the Findhorn Foundation in their cooperation through use of measured data and information surrounding the site.

REFERENCES

Cipriano, J., Mor, G., Chemisana, D., Pérez, D., Gamboa, G., and Cipriano, X., 2015, "Evaluation of a Multi-Stage Guided Search Approach for the Calibration of Building Energy Simulation Models," Energy Build., 87, pp. 370–385.

Flett, G., and Kelly, N., 2016, "An Occupant-Differentiated, Higher-Order Markov Chain Method for Prediction of Domestic Occupancy," Energy Build., 125, pp. 219–230.

Flett, G., and Kelly, N., 2017, "A Disaggregated, Probabilistic, High Resolution Method for Assessment of Domestic Occupancy and Electrical Demand," Energy Build., 140, pp. 171–187.

Good, N., Zhang, L., Navarro-Espinosa, A., and Mancarella, P., 2015, "High Resolution Modelling of Multi-Energy Domestic Demand Profiles," Appl. Energy, 137, pp. 193–210.

Heidarinejad, M., Mattise, N., Dahlhausen, M., Sharma, K., Benne, K., Macumber, D., Brackney, L., and Srebric, J., 2017, "Demonstration of Reduced-Order Urban Scale Building Energy Models," Energy Build., 156, pp. 17–28.

Jenkins, D. P., Patidar, S., and Simpson, S. A., 2014, "Synthesising Electrical Demand Profiles for UK Dwellings," Energy Build., 76, pp. 605–614.

Liang, X., Hong, T., and Shen, G. Q., 2016, "Occupancy Data Analytics and Prediction: A Case Study," Build. Environ., 102, pp. 179–192.

McKenna, E., Krawczynski, M., and Thomson, M., 2015, "Four-State Domestic Building Occupancy Model for Energy Demand Simulations," Energy Build., 96, pp. 30–39.

McKenna, E., and Thomson, M., 2016, "High-Resolution Stochastic Integrated Thermal-Electrical Domestic Demand Model," Appl. Energy, 165, pp. 445–461.

Miller, C., Nagy, Z., and Schlueter, A., 2015, "Automated Daily Pattern Filtering of Measured Building Performance Data," Autom. Constr., 49(PA), pp. 1–17.

Patidar, S., Jenkins, D. P., and Simpson, S. A., 2016, "Stochastic Modelling Techniques for Generating Synthetic Energy Demand Profiles," Int. J. Energy Stat., 4(3), p. 1650014.

Ramírez-Mendiola, J. L., Grünewald, P., and Eyre, N., 2017, "The Diversity of Residential Electricity Demand – A Comparative Analysis of Metered and Simulated Data," Energy Build., 151, pp. 121–131.

Richardson, I., Thomson, M., and Infield, D., 2008, "A High-Resolution Domestic Building Occupancy Model for Energy Demand Simulations," Energy Build., 40(8), pp. 1560–1566.

Richardson, I., Thomson, M., Infield, D., and Clifford, C., 2010, "Domestic Electricity Use: A High-Resolution Energy Demand Model," Energy Build., 42(10), pp. 1878–1887.

Widén, J., and Wäckelgård, E., 2010, "A High-Resolution Stochastic Model of Domestic Activity Patterns and Electricity Demand," Appl. Energy, 87(6), pp. 1880–1892.

Zhou, Y., Yu, Z., Li, J., Huang, Y., and Zhang, G., 2017, "The Effect of Temporal Resolution on the Accuracy of Predicting Building Occupant Behaviour Based on Markov Chain Models," Procedia Eng., 205, pp. 1698–1704.