A Limited-Data Model Of Building Energy Consumption

Andrew Rice
acr31@cam.ac.uk

Simon Hay
sjeh3@cam.ac.uk

Dan Ryder-Cook

Computer Laboratory
University of Cambridge

Abstract

We present a model targeted at practical, wide-scale deployment which produces an ongoing breakdown of building energy consumption. We argue that wide-scale deployment is practical due to its reliance only on commonly available sensor information and crowd-sourced inventory data. The results for our own building over the previous 10 months show many of the trends seen in the building's true, metered energy consumption and we find our model predicts long term averages within 10% of the true value in some scenarios. We further use our model to estimate the potential impact of some energy saving scenarios.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Misc.

General Terms

Measurement, Economics

Keywords

Personal Energy Meter

1 Introduction

Organisations with large and diverse estates face many challenges as they attempt to reduce building energy consumption. One key issue is that it is difficult to know where to apply limited resources for maximum benefit. We are seeking ways to provide information for planners in this situation by deploying a practical sensor-driven system which produces an itemised, model-based breakdown of an individual building's energy consumption.

Many current modelling systems focus on exploring design options or performing one-off examinations of buildings. Although capable of producing accurate results they require expert users and detailed building survey information. We feel that the resources to model and profile every building in a large estate are beyond many organisations.

BuildSys '10, November 2, 2010, Zurich, Switzerland.

Display Energy Certificates (DECs) fall at the other end of the spectrum and are cheap and easy to deploy. However, the very lack of operational detail which makes them easy to deploy also makes them uninformative for detailed planning. In this paper we study the energy consumption of the Computer Laboratory at Cambridge University which has an electricity use of 180 kWh/m^2/year. This fares particularly unfavourably in the DEC comparison to a standard building of type 'University Campus' with a reference consumption of 80 kWh/m^2/year. The DEC efficiency rating gives little information to estate manager as to what proportion of this discrepancy is due to inefficiency and what is due to building use.

In this paper we present a middle-ground tool for modelling the energy consumption of a building. We argue that the our scheme produces useful information about building energy consumption whilst relying only on data that is easily gathered and sensing which could plausibly be done on a large scale. We first explain the principles of our model and differentiate it from existing solutions (Section 2). We use a tool called OpenRoomMap to crowd-source inventory data in a scalable fashion (Section 3) and provide a number of means to estimate the energy traces of particular building aspects (Section 4). We present our results (Section 5) which predict an average consumption within around 10% of metered consumption. We believe that the fact that we can reproduce many of the dynamic features of the building's energy consumption is a good indicator that the model is giving some valuable insight. We therefore use the model to examine energy consumption under a number of potential energy saving scenarios (Section 6).

This work is part of our investigations into the concept of a Personal Energy Meter under our Computing for the Future of the Planet research theme [4]. We envisage a system that collects information about our daily consumption and provides breakdowns of the energy costs of activities to help us target areas for reduction. In previous work we outlined a mechanism for desegregating the total energy consumption of a building by user and apportioning it to an individual [3]. Here, we explore disaggregation by function, which is valuable not only to building managers concerned with reducing its overall cost but also to individuals who seek to understand why they have been allocated a certain share.

2 Building modelling tools

Our model operates by estimating the energy consumption for categories of devices in the building inventory and we use a variety of estimation methods to model different energy use patterns. The summation of energy consumption for each category is then compared with the recorded energy consumption from the building's electricity and gas meters.

There are dozens of existing energy modelling packages, though most target the design stage and require significant measurement and input data to function [1]. Perhaps the best known example is the DOE-2 software.[1] produced by the US Department of Energy DOE-2 uses hourly weather data to calculate the hour-by-hour performance and response of a building with a known description; heat gains to building spaces are converted to cooling or heating loads on the air using pre-calculated "weighting factors". An accessible interface to DOE-2 is provided through the eQuest package. eQuest is designed to make it easy for a single user to capture a building design and parameters in contrast to our approach in which we collect data from a large number of users.

An alternative "heat balance" method uses a detailed heat model of the thermal transfer processes in the rooms to calculate loads from heat gains; this is generally slower but more accurate. The best known example is the Building Load Analysis and System Thermodynamics (BLAST) system, also supported by the US government. It was developed for predicting energy consumption and systems performance and costs of new or retrofit building designs. We use a very simplified form of this method for our predictions.

EnergyPlus[2] combines many of the features from these programs. It uses a modular system to permit the construction of detailed building models. Many new building technologies and building and systems simulation models are accessible which represents a significant step forwards in terms of both computational techniques and program structure [2]. TRNSYS[3] is a general simulation package which makes use of modules to model a wide range of systems. The modular and extensible nature of these two systems provides a huge degree of flexibility and both would be candidates for hosting modules implementing the various aspects of our model.

3 Crowd-sourcing inventory information

Modern buildings contain a large number of widely-distributed electrical devices. This makes collecting detailed inventory information a tedious and time-consuming task. Furthermore, keeping abreast of changes to maintain the correctness of this information over time is implausible in many circumstances.

Fortunately, a highly accurate inventory is not necessary for our purposes. The vast majority of energy consuming devices are low power (electrical) devices consuming less than 100 W. Individually, each device is less than 0.05% of the total building consumption. However, their combined consumption does amount to a significant figure.

[1] http://www.doe2.com/

[2] http://apps1.eere.energy.gov/buildings/energyplus/

[3] http://sel.me.wisc.edu/trnsys/

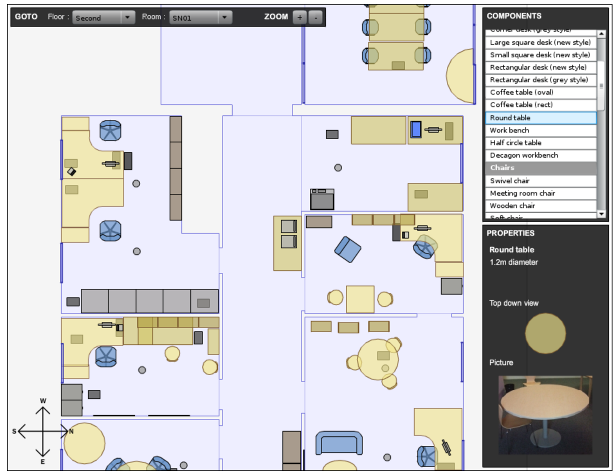

Figure 1. Building occupants can use the OpenRoomMap web interface to edit the building map.

From this we conclude that our inventory need not contain every single device as long as it contains a representative sample and an estimate of its coverage. Obviously, our confidence in the sample will increase with its size.

We use our OpenRoomMap system [6] to collect the building inventory data from the building's users themselves. OpenRoomMap displays through a web-based interface a plan view of the building contents which is editable by the occupants themselves (Figure 1). This data has found various uses ranging from assisting when locating printers and copiers to displaying layout information for room bookings. These varied uses have encouraged occupants to map public and shared spaces in addition to private offices. We estimate that around 70% of the building has been mapped.

OpenRoomMap provides a crucial part of our argument about the practicality of our system. The crowd-sourcing aspect of our data collection means it is much more likely that a representative selection of the building inventory will be collected than if we attempted to persuade every building manager to perform a survey manually. Examples such as OpenStreetMap highlight the effectiveness of recruiting just a small percentage of a large population for such participatory sensing tasks.

4 Energy estimation methods

We developed a number of energy estimation methods to model different device usage characteristics.

Constant rate The constant rate method simply assumes a continuous consumption for a device. This is appropriate for always-on devices such as safety lighting, VOIP telephones, and printer standby power. We measure the energy consumption of an example of each device using a plug-in power meter and assume this to apply to all devices of the same type.

Timed The building management system in our building controls the lighting in public areas (approximately 12 kW) according to a timer. This method applies our measured energy consumption of each device type at a constant rate during the programmed on periods.

Sub-metered It is increasingly common to install sub-metering to monitor the consumption of large energy consumers within a building. In our building for example we used sub-meters to profile the energy consumption of our machine rooms and associated air-conditioning. These account for an average of 89 kW which is a significant proportion. One should note that sub-metering by building region (e.g. corridor) is not directly useful here because many different types of device will be connected to the same circuit. This helps provide a spatial breakdown of energy consumption but not an itemised one. Similarly, the problem of providing an itemised breakdown cannot be fully solved by sub-metering large energy consumers. For example, we estimate our office-style lighting consumes more than 100 kW (when all switched on). Direct measurement of this would require metering every lighting circuit in the building and then removing the consumption of all the other (timed) lighting from this total.

4.1 Occupancy

Some devices in the building are switched on and off by occupants and so we provide a method to modulate the power estimate by the number of people in the building. This method scales the total power consumption of a set of devices by the proportion of the maximum expected building users currently present. We apply this method to devices such as computer monitors and office lighting.

As one might expect, building occupancy varies significantly over time. We expect low occupancy over weekends and holiday periods but also due to less predictable causes such as travel disruption (the UK transport infrastructure copes poorly with snow for example). For this reason we expect some form of sensor data will be required for occupancy estimation.

The method we use is based on the access logs from the security system [3]. This is a somewhat broad approximation in our building because 1) the system is based around electronic door locks and so many people can pass through a door for a single access entry; and 2) the system only authenticates ingress events and so we only see generic unlock events rather than (anonymised) identifiers when someone exits. Alternative or complementary approaches might be to determine occupancy based on wireless traffic from smart phones or workstation activity or to use GPS to determine when a user is en-route to or from the building.

Figure 2 shows our estimated occupancy trace for 2008. Dark colours correspond to a large number of people in the building and lighter colour correspond to fewer people. It shows quiet days which correspond to UK public holidays and a general ebb and flow corresponding to term and vacation periods. We note a number of exceptions such as public holidays during term time (which the university does not observe) and the especially quiet period for the two weeks surrounding Christmas.

4.2 Heating, Cooling and Ventilation (HVAC)

The final method provides an simple estimate of the energy consumption of the building's HVAC system. Our approach is to estimate the amount of energy required to keep the interior of the building at a desired set-point temperature (T_i) given the heat input in to the building from device energy use (including computers) and the heat loss (or gain) due to the outdoor temperature (T_e). We model the system as thermal energy movement between three bodies: 1) the interior, 2) the exterior walls, 3) the exterior (Figure 3). Thus, the exterior walls are acting as a buffer between the interior and exterior temperatures. We measure outdoor temperature with a weather station on the roof of our building.[4]

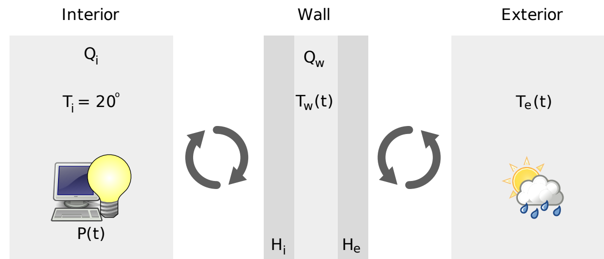

Figure 3. Thermal model for the building HVAC

$$\frac{dQ_i}{dt} = -H_i(T_i - T_w(t)) + P(t)$$
$$\frac{dQ_w}{dt} = C_w\frac{dT_w}{dt} = -H_e(T_w(t) - T_e(t)) + H_e(T_i - T_w(t))$$

The first of these equations gives the HVAC load—the energy required to maintain the internal fixed point temperature T_i (which for our building is an average of 21 °C). The (numerical) integral of the second equation tracks the temperature of the wall over time. A negative HVAC load indicates that energy must be put into the building to maintain the temperature (heating demand), and a positive HVAC load indicates that energy must be removed (cooling demand). Thus, we can interpret the load in three different ways:

1. $\left|\frac{dQ_i}{dt}\right|$ is the total power required to maintain the building temperature;
2. $\mathbf{max}(0, \frac{dQ_i}{dt})$ is the total power required to cool the building;
3. $\mathbf{max}(0, -\frac{dQ_i}{dt})$ is the total power required to heat the building.

H_i and H_e correspond to the leakiness of the wall towards the interior and exterior of the building respectively. These are derived from multiplying the surface area (m^2) by the thermal transmittance or U-value ($W/m^2/K$). We manually estimated the surface area (40,000 m^2) but such information could also be derived from the floor area in OpenRoomMap and an estimate of ceiling height and roof-pitch. U-values are normally quoted for a single surface and we adopt a typical value suggested by MacKay for best building methods of 0.15 [5]—our building won an architectural award for its heating and cooling.[5] Our model utilises a U-value for the inner shell (U_i) of the wall and a U-value for the outer shell (U_e) and so we further assume that the outer shell has 2.5

[4] http://www.cl.cam.ac.uk/research/dtg/weather/

[5] http://www.cabe.org.uk/case-studies/william-gates-building/design

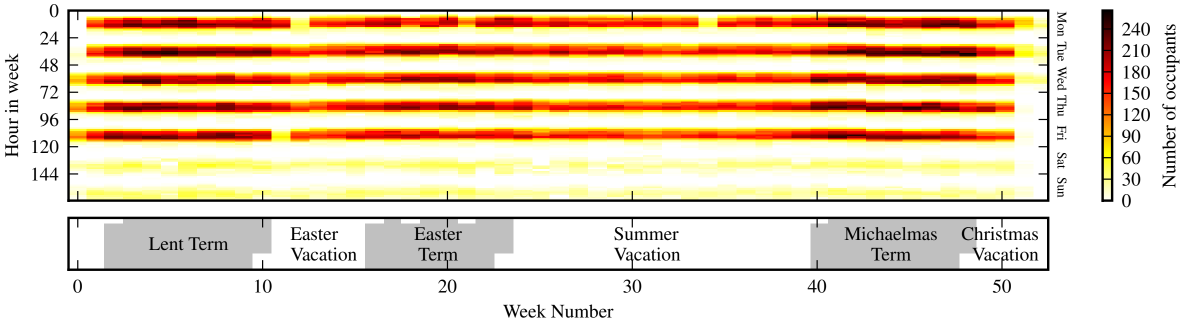

Figure 2. Estimated occupancy trace for 2008

times the thermal resistance of the inner shell. Given that U-values combine in the same manner as resistors in parallel $1/U = 1/U_i + 1/U_e$ and substituting $U_e = 2.5U_i$ we find $U_i = 3.5U = 0.53$ and $U_e = 3.5U/2.5 = 0.21$.

5 Results

Figure 4 shows the model output and recorded electricity consumption in kW for November 2009 to August 2010. The dark 'metered consumption' line is the half-hourly measurements of electricity consumption as recorded by the electricity company. The categories in the breakdown are as follows:

- **HVAC** is the output of the heat model for the building. We initially consider only cooling to account for electricity usage.
- **Lights** includes lighting within offices (modulated according to the occupancy of the building) and in public areas (modulated according to a timer function).
- **PCs** covers the energy use of personal computers and monitors in offices. We assume that the PC itself is left on continuously whereas the monitors are switched on or off according to the occupancy of the building. Both are assumed to consume 70 W.
- **Machine rooms** considers servers, uninterruptible power supplies and air conditioning units in our machine rooms. This is a mixture of sub-metered readings and manual estimates.
- **Other** contains minor items from the OpenRoomMap inventory such as printer idle power, telephones and a small number of electric heaters.

Notable from the graph is that the predicted consumption displays similar trends to the true measured value. Over the annual period we notice the load on the HVAC system increase during the summer months and fall to nothing over the Christmas period when the building is quiet and the exterior temperature is low. Figure 5 shows a two week period in January. The peaks in consumption during working weekdays are clear in the model and we can see from the breakdown that this is mostly due to lights being switched on (in offices). Figure 6 shows a two week period in July. In this case the HVAC energy usage is significantly higher due to higher outdoor temperatures.

We now consider the effect of including heating in our model. We assume that our heating system is 70% efficient

Scenario	Av. Power	Change	Saving
Metered	275 kW		
Current conditions	213 kW		
Normal comp.	118 kW	95 kW	£83,000
PCs off	206 kW	7 kW	£6,100
LED lighting	192 kW	21 kW	£18,000

Figure 7. Predicted annual energy savings

and from the gas consumption over the summer when no space-heating is needed we derive an additional cost of 1.4 kW for water heating which we include in the 'Other' category. We see what seems to be a more significant deviation from the measured trace (Figure 8, showing both electricity and gas). However, this is due in part to the fact that the gas consumption data for our building is measured monthly and interpolated linearly so day-scale changes in consumption as predicted by our HVAC model are not reflected in the measured consumption trace. There are many factors which we could alter to obtain a better fit, such as changing the U-value of the building, the efficiency of the heating system or the fixed point temperature but we refrain from doing so for fear of over-fitting what is a very simple model.

6 Energy saving scenarios

The itemisation produced by the model suggests three big areas in which we could save energy: machine rooms, PCs and lighting. We now use our model to consider these scenarios (Figures 7 and 9).

- **Normal computing** We estimate the energy consumption of our building if (like many other buildings on the estate) it contained no significant server infrastructure and a single workstation per occupant.
- **PCs off** We estimate the impact of building occupants switching off all workstations when not in the building.
- **LED lighting** We estimate the impact of switching to LED lighting, replacing our current 50 lm/W lighting with LED equivalents achieving 160 lm/W[6].

Finally, we return to our goal of producing a modelling tool which can be automatically applied across many build-

[6] The US Department of Energy estimates the 160 lm/W LED lighting will be market-ready by 2025: `http://www1.eere.energy.gov/buildings/ssl/efficacy.html`

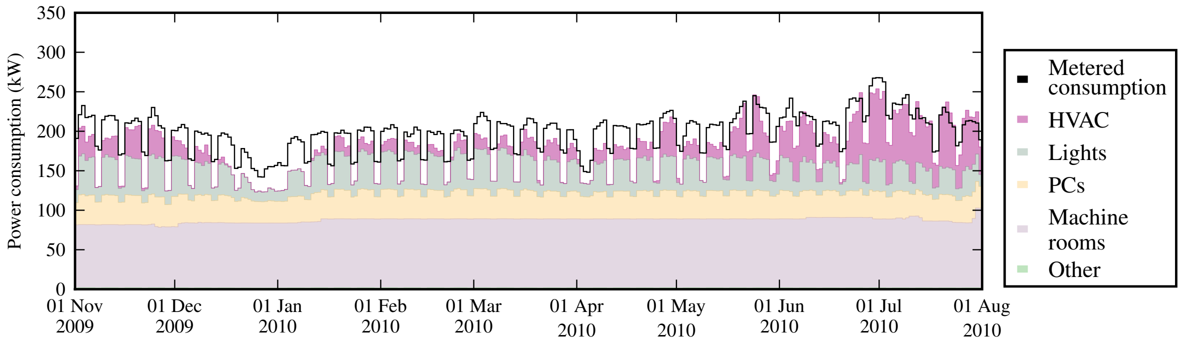

Figure 4. Daily breakdown (Nov 09 to Aug 10) shows trends in electricity consumption are correctly estimated

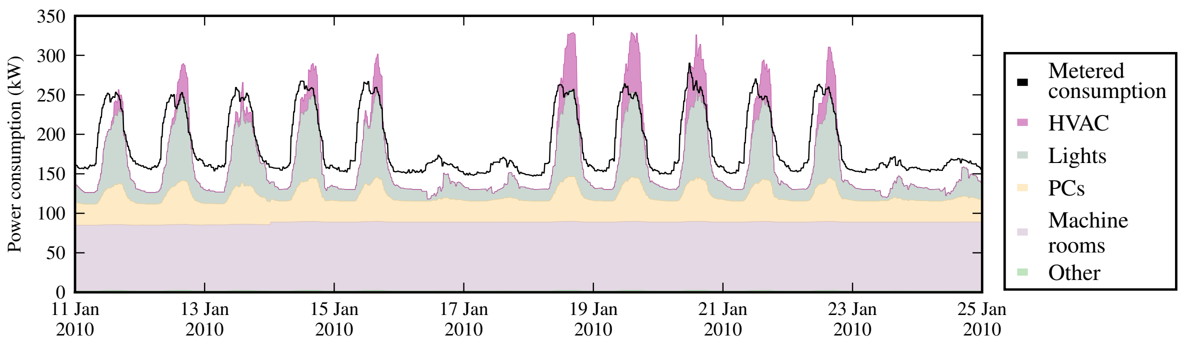

Figure 5. Half-hourly breakdown (Jan 10): electricity requirements during winter vary mostly due to lighting needs

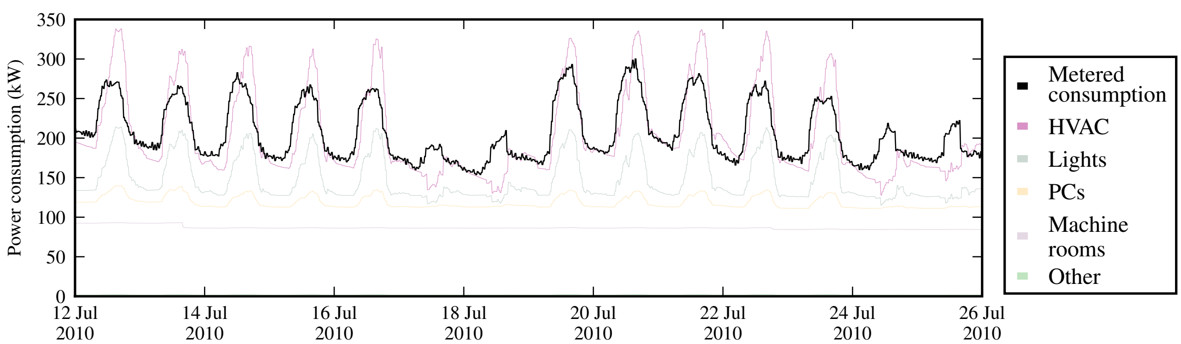

Figure 6. Half-hourly breakdown (Jul 2010): cooling dominates the electricity requirements during summer

ings and consider the sensitivity of the model to the building U-value. Figure 10 shows the result of running the simulation with 4 different U-values. These results show that a good choice of value is probably around our choice of 0.15. Its clear that the wrong choice of U-value can have a significant impact on the quality of fit. However, it is easy to notice that the fit is incorrect. One technique might be to collect data as to the point in the year when the building's heating system is first switched on for a significant period of time and to adjust the building U-value to produce a similar effect.

7 Conclusions and Future Work

Producing an explanation for a building's energy consumption is an important first step to improving efficiency. However, for large estates it is impractical to produce detailed engineering models of every building's energy consumption. We have described a modelling technique which we believe could be practically applied across many buildings. Minimising the effort involved in initial data collection is important to this goal and so we use OpenRoomMap to crowd-source this information from building users. We have also described the different ways in which we provide a system which can operate with a minimal amount of live sensing information but could still extend to accommodate more sources as they are installed.

There are many potential ways to improve the modelling and we are most interested in those techniques which require little additional data input. Useful examples might be to modulate building lighting with reference to natural light levels using more detailed weather data and information from OpenRoomMap about the positions of windows in offices. OpenRoomMap data could also provide more as-

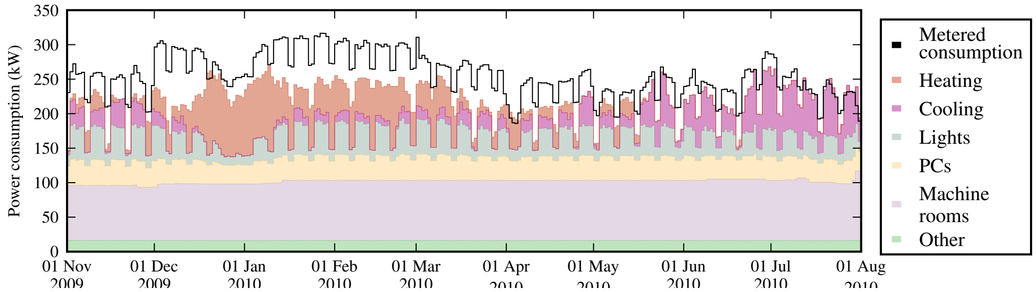

Figure 8. The model underestimates combined heating and cooling energy consumption during winter

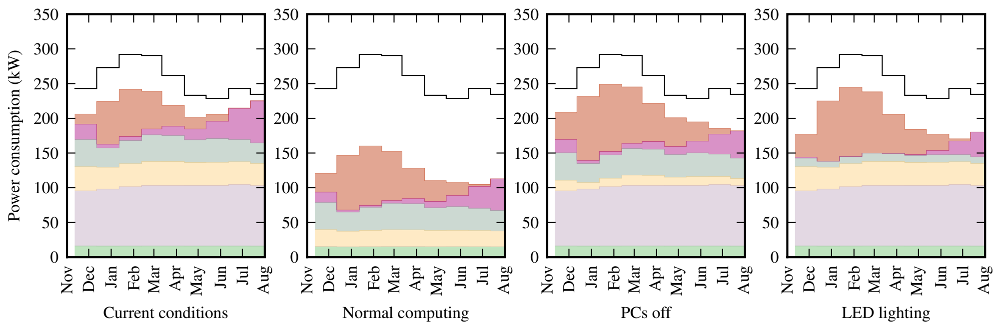

Figure 9. Increases in heating load and decreases in cooling load follow from energy savings

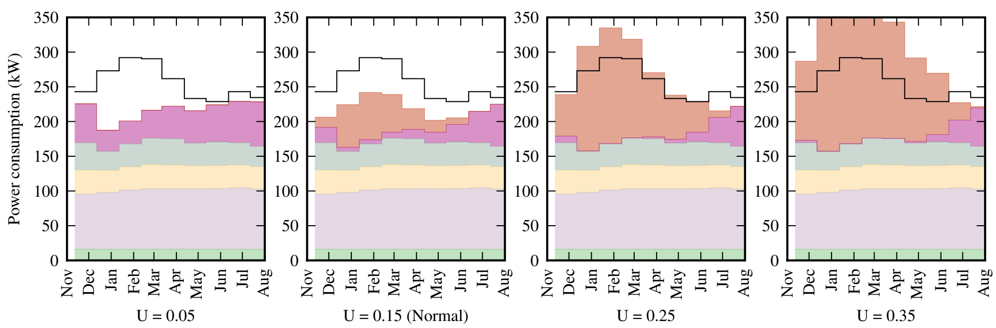

Figure 10. Varying the choice of U-value has a significant impact on the model prediction

sistance in estimating building parameters by providing estimates of building surface area and the relative ratios of walls, windows and roofing.

We are hopeful that we can make further extensions to the model as required by other buildings without compromising our goal of a scalable, practical system and intend to release our code either as a standalone package or integrated into an existing modelling framework.

8 Acknowledgements

We would like to thank Camille Baldock, Ian Burton-Palmer, Piete Brookes, Martyn Johnson and Paul Hasley for helping with various data and information, Oliver Woodman for his assistance with OpenRoomMap and Andy Hopper for his insight and support. This research was supported by a Google Focused Research Award and an EPSRC PhD studentship.

9 References

[1] D. Crawley et al. Contrasting the capabilities of building energy performance simulation programs. *Building and Environment*, 43(4):661–673.

[2] D. Crawley et al. EnergyPlus: creating a new-generation building energy simulation program. *Energy and Buildings*, 33(4):319–331.

[3] S. Hay and A. C. Rice. The case for apportionment. *In Proc. BuildSys '09*.

[4] A. Hopper and A. C. Rice. Computing for the future of the planet. *Phil. Trans. R. Soc. A*, 366(1881):3685–3697.

[5] D. J. C. MacKay. Sustainable energy - without the hot air. *UIT, Cambridge*.

[6] A. C. Rice and O. J. Woodman. Crowd-sourcing world models with OpenRoomMap. *In Proc. PerCom '10*.