Developing a simple software tool for the analysis of high resolution data from 'smart' utility meters

Graeme Stuart
Institute of Energy and Sustainable Development
De Montfort University
United Kingdom
gstuart@dmu.ac.uk

Dr Simon Taylor
Institute of Energy and Sustainable Development, De Montfort University

Professor Paul Fleming
Institute of Energy and Sustainable Development, De Montfort University

Keywords
analysis, buildings, energy management, metering, empirical modelling, utilities, water

Abstract
Energy management is used by organisations to monitor and control the environmental and financial costs of energy usage. It involves monitoring performance, identifying potential interventions and assessing their effectiveness. A key element of an effective regime is access to good quality data. High resolution utility consumption data (typically electricity, gas and water) is becoming widely available. Leicester City Council has access to such data in over 200 municipal buildings.

The data contains details of the effects of both interventions and previously unidentified events. Appropriate data analysis techniques can extract these details to provide a simple and intuitive view of the usage behaviour, although the process can be time consuming and complicated.

A typical analysis follows three main steps.

1. Producing an initial model of consumption using all the available data.
2. Subdividing the model into internally consistent, contiguous periods.
3. Assessing each distinct period of consistent consumption.

Developing a simple software tool to perform an analysis using this framework involves making compromises between complexity, resolution and generality. Designing the user interface introduces compromises between usability, training requirements and effectiveness.

A case study shows that implementing the analysis framework in a simple tool provides significant insight into consumption history and can quickly quantify the effects of interventions and identify hitherto unknown events.

Such a tool has the potential to improve the effectiveness of energy management by providing a rapid assessment of energy usage patterns. Identifying unexpected increases in consumption can trigger investigations that can lead to timely remedial action, quantifying the effect of interventions is critical in the assessment of their success. Future possibilities include exception reporting, ensuring high quality results with minimum input of time.

The tool is presented as a proof of concept for the basic analysis framework in a way that can be implemented easily to provide detailed analysis of data wherever it is available in an appropriate format. The potential to provide energy management information covering significant consumption is promising.

Introduction
With the growing need to reduce global greenhouse gas emissions, we need to reduce energy consumption from existing buildings. Traditionally, various energy management techniques have been used by organisations to monitor and control the environmental and financial costs of energy usage. Information revealed about the energy performance of buildings can inform investment in energy efficiency measures and is a critical input to strategic energy management decisions in large, multi-site organisations. Monitoring energy performance over time can quantify the effects of energy interventions and highlight faults as they occur. Comparison of the energy per-

formance of different buildings serves to highlight those where savings may be most cost effective.

Energy management information systems rely on good quality data and suitable analysis techniques to provide information on the energy performance of buildings. Energy savings can be achieved through acting on information provided by analysis of data[1]. Traditional energy management techniques are based on monthly billing data. Modern energy metering can produce data at far higher than monthly resolution.

Analysis techniques typically use the degree day [2], a measure of the heating (or cooling) requirement for a building based on heat transfer through building fabric due to a temperature gradient. Degree days for a period are determined by multiplying the number of degrees by which the average external temperature falls short of a "base temperature" (usually 15.5°C) by the number of days for which the deficit occurs. The base temperature is the internal temperature above which no heating is needed and may vary with circumstances. For example, less mobile people (e.g. the old or infirm) often require a higher internal temperature, and a higher base temperature might be appropriate. Degree days provide a means to compare energy performance in buildings under different conditions [3]. Analysis techniques also use degree days to produce empirical models of consumption [4].

This paper describes the key design considerations when designing a tool to implement degree day based analysis using high resolution data from automatic meter reading (AMR) systems, also referred to as "intelligent" or "smart" metering. It begins with the data requirements along with how they are satisfied by the output of a comprehensive monitoring system installed by Leicester City Council. Analysis of the data is covered next, and finally the design of the tool itself, with particular reference to the choice of analysis model and user interface design.

Data

Energy consumption data for energy management are typically sourced from billing records. Supply companies record electricity and fuel consumption to produce accurate and timely bills. In the UK, heating degree day data (base 15.5°C) are published monthly and can be collected from many sources. The UK Carbon Trust publishes degree day data for all UK regions [5]. Consumption data from billing systems and monthly degree day data are the recommended raw materials for an assessment of energy performance[6].

In practice billing data can suffer from significant inaccuracies and estimated readings are common where meter readings are not possible for one reason or another. Meters are not consistently read on the first day of every month so it is rare for "raw" billing data to be perfectly aligned with degree day data. There is thus a requirement for some pre-processing of the billing data to apportion consumption into monthly bins resulting in well-aligned but artificial, monthly consumption and degree day data.

Modern metering can produce electronic data that can be automatically logged at any interval. Typically in the UK data from large electricity supplies is produced at half hourly intervals due to more accurate billing arrangements following market deregulation. These data can be downloaded directly from the meters and suffer far less from the inaccuracies historically found in manually read billing data. When coupled with communications hardware, it is possible to develop systems [7] that transfer data from multiple meters to a central point where they are logged in a database and made accessible through software. Standard half-hourly data files are produced that hold data in 49 columns, one for date and one each for every half hour period in a day making each row represent a day of consumption.

Leicester City Council has installed a 'smart' or 'intelligent' metering system to monitor electricity, gas and water consumption in over 200 municipal buildings in the city of Leicester in the East Midlands region of the UK. The system is also connected to outside air temperature monitors so it is possible to calculate degree days for any period and to any base temperature. Half-hourly resolution data from this system is made available through, and analysed using, proprietary software. These data are easily exported in standard half-hourly data file format.

Analysis

High resolution gas consumption data can be used to identify the detailed energy consumption patterns in a building[8]. Detailed analysis can also reveal alterations in these patterns due to changes to building fabric, installed equipment, automatic controls or building occupant behaviour. Changes may be the result of active energy efficiency interventions, side effects from other building work or unidentified events such as faults. A data analysis methodology has previously been developed for use with half hourly electricity data [9]. A typical analysis follows three main steps:

1. Producing an initial model of consumption using all the available data
2. Subdividing the model into (relatively) internally consistent, contiguous periods
3. Modelling each distinct period of consistent consumption individually.

A plot of daily energy consumption against daily degree days provides a suitable initial model of consumption using all the available data. Where the relationship is linear, the y-intercept and gradient of this line can be used to assess the non-heating related consumption (e.g. for hot water) and the heating related consumption respectively. The scatter of the points around the line indicates the reliability of the relationship. The relationship can be nonlinear, and useful clues about problems with the system can be obtained from the general shape [xi, xii]. The degree day method is usually appropriate for measured gas consumption in continually heated buildings. Where a cooling load is present the method can be adapted to use cooling degree days but is less reliable. In intermittently heated buildings the simple model provides a poor fit, though more complex models are possible they are outside the scope of this paper.

The second stage, subdivision into sub-models for different time periods, can be carried out using CUSUM. CUSUM (CUmulative SUM) is a method of determining systematic patterns in the residuals resulting from the fit of an initial model to a set of data [10]. The cumulative sum of the residuals is plotted as a function of time and indicates systematic differences from the initial model. For example, a steady increase at constant slope corresponds to a heating load that exceeds the modelled one by a constant factor. Typically the plot consists of several sections of this type separated by distinct "kinks" which represent "events" at which changes have taken place. Details of the interpretation of consumption – degree day plots and their use to establish a prediction for use with the CUSUM event detection algorithm can be found in many publications [11, 12].

In the third stage, a similar model is defined for each identified period and parameters are compared between neighbouring periods to provide insight into the nature and magnitude of each change in consumption pattern.

The outcomes of such an analysis include model parameters that describe the data as a whole, a series of dates that represent significant events in the site's history and equivalent model parameters for each period between events. This information can prove useful to verify and quantify the effects of energy efficiency interventions, to identify and diagnose unexpected events as they occur and to gain detailed insight into the changing consumption patterns in a building. Application of this process to high-resolution data files can be time-consuming and complicated. A flexible software tool is the ideal solution since it can provide a simple interface and consistent analysis with replicable results.

Tool design

An analysis tool has been developed to facilitate the rapid analysis of high resolution gas consumption data. Details of the underlying models and event detection algorithms will be reported elsewhere. Here the emphasis is on the design of the tool. Developing a simple tool to perform an analysis using this framework involves making compromises between complexity, resolution and general applicability. The user interface design requires further compromises between usability, training requirements and effectiveness. This initial version of the tool was developed as a Microsoft Excel spreadsheet as a proof of concept.

A key decision made at the early stages was the resolution at which the tool would operate. The available data (gas consumption and outside air temperature) is at half-hourly resolution. The resolution of the analysis has an influence on the complexity of the analysis, the resolution of the results and the number of calculations required. Operating at daily rather than half-hourly resolution reduces the quantity of consumption data from 17,520 values to 365 values per year, thus reducing the computational resources required. It also simplifies the level of analysis required, by essentially averaging consumption over each day. The main disadvantage of choosing lower resolution is that the analysis provides less information. Models will be simple – there will be no need to distinguish between day-time and night-time consumption, for example – and so have fewer parameters, and events will be detected to the nearest day rather than to the nearest half hour. In the case of energy management in buildings it is perfectly acceptable to be notified of a problem after a number of days and there is rarely a case where notification is required at a sub-daily resolution. In practice data is transferred by Leicester City Council on a daily basis as 48 values and timestamps and so will not be available to the system until the following day anyway. With this in mind, a daily resolution is most appropriate and any loss in accuracy or insight has been accepted as the result of a trade-off for simplicity of the underlying calculations and simplicity of the tool itself. The tool accepts half-hourly consumption and temperature data in standard, 49 column format and calculates daily total consumption and degree days to enter into the all-data model.

CHOICE OF MODEL

Daily consumption and daily degree days are in many cases linearly related. However, buildings are not occupied 365 days per year. Many buildings are closed or at least use far less energy on weekends. In order to produce a software tool capable of analysing a wide variety of the buildings for which Leicester City Council collects data, it was decided that the underlying model would accommodate two "modes" of operation, "occupied" and "unoccupied". The tool presents an "all-data" model as a scatter chart of daily consumption against degree days. Degree day base temperature is variable and user-defined. Model parameters are calculated for both occupied and unoccupied periods. Our early work has, however, concentrated on the simplest possible cases to build a firm foundation for future studies. In the example shown in Figure 1 the building is a residential care home. The building is occupied 365 days a year so there are no "unoccupied" data.

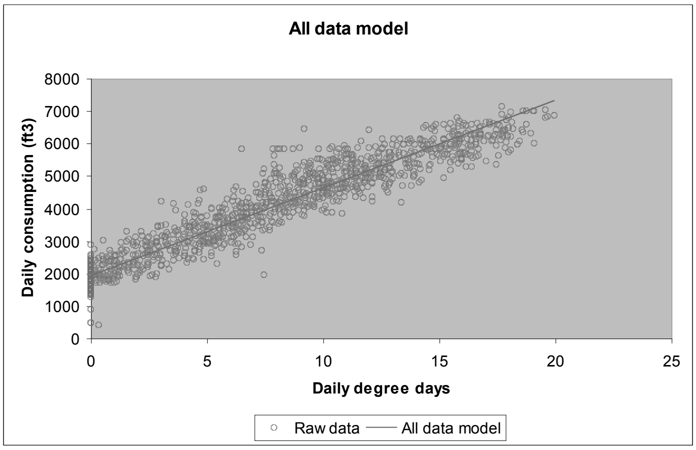

Figure 1: Basic, "all-data" degree day (base 15.5) model

USER INTERFACE

User interaction is critical in building a complete analysis. Firstly the user can choose a degree day base temperature appropriate to the building by either selecting the default value of 15.5°C or by inspecting the effect a change has on the all-data model. The event detection module is based on a user-driven visual inspection of a CUSUM chart produced from the model residuals. Users are presented with the chart and invited to log the dates of significant and abrupt changes in the CUSUM gradient. As they log the events, the event module calculates new model parameters for each inter-event period. The tool also plots straight lines through the CUSUM plot showing how the selected dates match the CUSUM kinks. An example is shown in Fig. 2.

To inspect the precise effect of events the user is invited to select an event to investigate. In the example of figure 2 the 7th event on the 21st November 2005 is highlighted. The selected

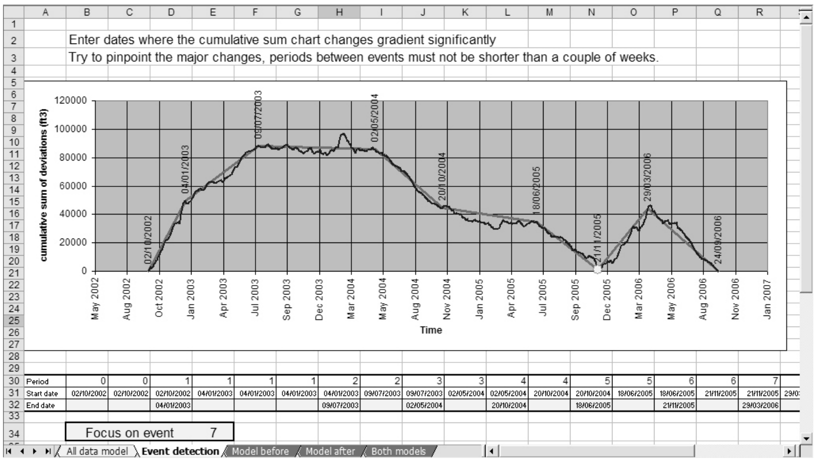

Figure 2: CUSUM chart and user interface

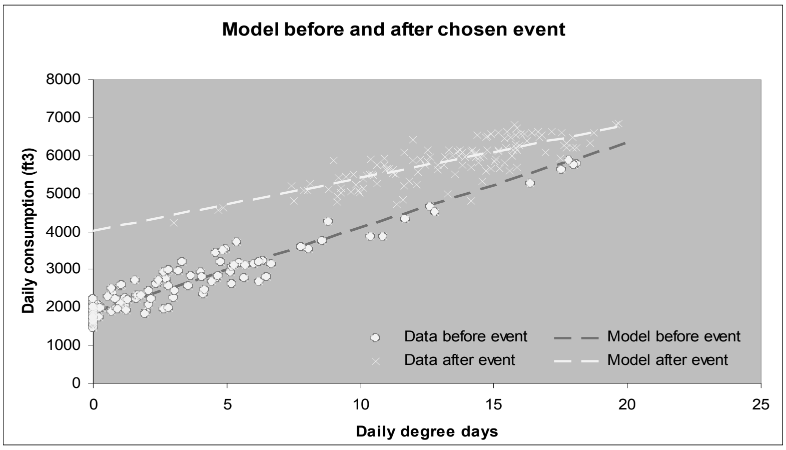

Figure 3: Degree day models either side of an event

event is presented as a scatter plot similar to the all-data model in figure 1. Only the data between the previous event and the next event are shown. Data from before the chosen event are distinguished from that after the chosen event by using different symbols and colours. Figure 3 shows how the event identified in figure 2 appears to be connected with a significant increase in y-intercept (from just below 2000 ft^3 to nearly 4000 ft^3) and a decrease in gradient (from about 200 ft^3/degree day to about 125 ft^3/degree day). Inspection of the following period shows consumption returning to the original pattern.

These parameter changes are directly related to performance. A change in intercept is related to the non-heating related consumption (including fixed losses in the heating system) and a change in the gradient is related to the systems response to changing degree days. Increases in system efficiency will lower the gradient but reductions in building internal temperature will have a similar effect. A coincident change in both parameters implies an event that affects both the fixed losses and the system response to changing degree days.

Conclusions

Implementing the analysis framework described earlier in a tool such as this can provide rapid access to useful results. It offers significant insight into consumption history and can quickly quantify the effects of interventions and identify hitherto unknown events that may require remedial action. Though data analysis alone cannot produce energy savings, the information provided through such an analysis has the potential to improve the effectiveness of energy management systems.

Such a tool has the potential to improve the effectiveness of energy management by providing a rapid assessment of energy usage patterns. As a proof of concept the tool provides easy access to information locked in high volumes of energy consumption data. Future possibilities include exception reporting ensuring high quality results with minimum input of time, implementing the analysis using a web-based interface to automatically connect to multiple datasets and provide results to multiple audiences including building users and energy managers and applying this type of analysis to large existing collections of data such as may be available through utility companies.

Further development of the analysis methodology and underlying energy models will enable the tool to analyse a wider range of building types and circumstances.

(Endnotes)

[1] V Ferreira, Fleming P.D and Stuart G. Intelligent energy and water performance assessment in municipal buildings. International Conference Improving Energy Efficiency in Commercial Buildings (2006). IEECB'06 conference proceedings. Germany.

[2] Department of the Environment, Transport and the Regions (DETR, 2006). Good practice guide 310: Degree days for energy management – a practical introduction. London: HMSO; 2006

[3] Department of the Environment, Transport and the Regions (DETR, 1998). Energy consumption guide 73: saving energy in schools – a guide for headteachers, governors, premises managers and school energy managers. London: HMSO; 1998.

[4] CIBSE (2006), Degree-days: Theory and application CIBSE TM41. London: Chartered Institution of Building Services Engineers, 2006.

[5] Carbon Trust. [WWW], (http://www.carbontrust.co.uk/KnowledgeCentre/degree_days/default.htm) [site accessed 19/01/2007].

[6] Carbon Trust (2006), CTG004: Degree days for energy management – a practical introduction. London: HMSO; 2006.

[7] Department of the Environment, Transport and the Regions (DETR, 1998). General information leaflet 49: low-cost automatic meter reading system using low-power radio. London: HMSO; 1998.

[8] V Ferreira, Fleming P.D and Ajiboye P. Identifying energy saving opportunities in buildings by the analysis of time series data. ECEEE summer study proceedings (2003). ECEEE. France.

[9] G. Stuart, Fleming P, Ferreira V, Harris P. Rapid analysis of time series data to identify changes in electricity consumption patterns in UK secondary schools. Building and Environment 42 (2007) 1568 – 1580.

[10] Page ES. Continuous inspection schemes. Biometrika 1954;41:100 – 115.

[11] Harris P. Energy monitoring and target setting using CUSUM, Cambridge: Cheriton Technologies Publication; 1989.

[12] Department of the Environment, Transport and the Regions (DETR, 1998). Good practice guide 112: monitoring and targeting in large companies. London: HMSO; 2006

Acknowledgements

This research is supported by the INREB Faraday partnership through EPSRC, Pilkington Energy Efficiency Trust and Leicester City Council.